Dynamic Choice Models with Conditional Gauss Markov Signals – Theory and Perturbation Approximation*

Elena Afanasyeva† Ivana Komunjer‡ Molin Zhong§

March 2025

[PRELIMINARY. PLEASE DO NOT CIRCULATE.]

Abstract

In this paper, we consider a well known class of infinite horizon dynamic stochastic choice models in discrete time, and seek answers to two questions: first, how to characterize optimal policies when the decision maker does not perfectly observe the model states? and second, whether the commonly employed perturbation method delivers a good approximation to the optimal policy in this model with learning? To answer those questions we focus our attention on a new class of information structures which we call *Conditional Gauss Markov* (cGM). In answer to the first question, we show that the Euler equation, which characterizes optimal decision rules under imperfect information, has under the cGM information structure a full information equivalent with finite dimensional states. In answer to the second question, we derive the deterministic limit of this full information equivalent model and show that the usual perturbation approximation applies.

1 Introduction

Dynamic stochastic choice models in which agents learn about hidden states are ubiquitous in Economics. It is well known that in such models optimal decision rules depend on the decision maker's posterior distribution over the hidden states given the observed history.[1] As a consequence, models with learning generally have state vectors that are infinite dimensional, which makes their study challenging. In certain

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

Acknowledgments: The authors would like to thank the participants at the Econometrics seminars at Penn State University, University of Chicago, University of Notre Dame, and University of Pittsburgh, as well as 18th Greater New York Metropolitan Area Econometrics Colloquium at the University of Rochester, for their valuable feedback.

†Federal Reserve Board, Washington D.C.; elena.afanasyeva@frb.gov.

‡Georgetown University; ik334@georgetown.edu.

§Federal Reserve Board, Washington D.C.; molin.zhong@frb.gov.

[1]Those results can be traced back to Striebel (1965), Rhenius (1974), and Yushkevich (1976), and Bertsekas and Shreve (1978).

special cases, however, the conditional distributions that make the new states can be summarized by a handful of sufficient statistics which are transitive (Bahadur, 1954), in the sense that their current values can be calculated from their values in the previous period and the newly observed variables. In this paper, we define and analyze a class of nonlinear dynamic choice models which have this property. We call them the *conditional Gauss Markov* models.

We embed our analysis in the well known framework of Sokey and Lucas (1989) but with one key difference: instead of requiring that the agent perfectly observes the exogenous model states, we allow the states to remain hidden. The information structure is that of a conditional Gauss Markov model: a nonlinear state space model with two key properties. First is the property that the conditional distributions of the hidden states given the observed history be Gsussian; second is the property that the statistics sufficient for control be Markovian. After setting up the model, we show that the well known optimality conditions—Euler equations and transversality conditions—characterize optimal decision rules even in the situations in which the decision maker does not observe all the relevant exogenous states. This result allows us to derive the first theoretical contribution of the paper: a full information model which we show is equivalent to the original incomplete information model. In the full information equivalent model, the agent observes all the model states whose dynamics we characterize. With the full information equivalent model in hand, we next examine the question of optimal policy approximation. By studying the asymptotic limit of the model as all the stochastic terms go to zero at the same speed, we show that the regular perturbation approximation à la Fleming (1971) applies. This result is the second theoretical contribution of the paper.

The remainder of this paper is organized as follows. Section 2 previews our key findings in the context of a simple example: a stochastic growth model with time varying parameters and uncertainty. Section 3 defines the class of dynamic choice models with conditional Gauss Markov information structure and studies the necessary and sufficient conditions for optimality. Section 4 then analyzes the perturbation approximation of such models. All the assumptions and proofs are relegated to an Appendix.

2 Motivating Example

Consider a simple a stochastic growth model with a representative agent choosing her consumption C_t and investment I_t so as to maximize

$$E\left[\sum_{t=0}^{\infty}\beta^t\frac{C_t^{1-\nu}}{1-\nu}\right] \quad (\text{case } \nu>1), \qquad \text{or} \qquad E\left[\sum_{t=0}^{\infty}\beta^t\ln C_t\right] \qquad (\text{case } \nu=1)$$

with $0<\beta<1$ subject to the output Y_t and capital K_t constraints: $Y_t = A_t K_{t-1}^{\alpha}$, $Y_t = C_t + I_t$, $K_t = (1-\delta)K_{t-1} + I_t$, $C_t \geqslant 0$, $I_t \geqslant 0$, $K_t \geqslant 0$, with $0<\delta\leqslant 1$, $0<\alpha\leqslant 1$, and $K_{-1}>0$ given. At each time period $t=0,1,\ldots$ the agent observes the last period capital stock K_{t-1}, the current level of (log) technology $a_t = \ln A_t$, whose law of motion is given by:

$$a_{t+1} = \theta_t a_t + \sigma_{a,t+1}\epsilon_{a,t+1}, \quad \ln\sigma_{a,t+1}^2 = \omega_a + \lambda_a \ln\sigma_{at}^2 + \alpha_a a_t, \quad \sigma_{a0}^2 \text{ given}, \tag{1}$$

where the latent drifting parameter θ_t evolves as:

$$\theta_t = \theta_{t-1} + \sigma_\theta \epsilon_{\theta t}, \quad \theta_0 \sim N(\overline{\theta}, \overline{P}). \tag{2}$$

We assume that the technology shocks $\{\epsilon_{at}\}_{t=1}^{+\infty} \sim$ iid $N(0,1)$ are Gaussian and independent of the parameter shocks $\{\epsilon_{\theta t}\}_{t=1}^{+\infty} \sim$ iid $N(0,1)$, which are also Gaussian. In this model, there is both a drifting unobserved parameter θ_t and a time varying uncertainty in productivity as captured by the time varying volatility σ_{at}. The model is similar to those studied in Cogley and Sargent (2005) and Primiceri (2005).[2]

The first question we ask is: how should one solve for the optimal consumption C_t? Our first result is to show that this incomplete information model has a *full information equivalent* in which all the model states are *observed* by the decision maker. Moreover, the optimal policy in the full information equivalent model equals the optimal policy in the original incomplete information model. Therefore, one can characterize optimal consumption C_t by considering the optimality conditions of the full information equivalent model. Specifically, for the model considered here, optimal consumption C_t is obtained by finding a policy $C_t = C(K_{t-1}, , a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$ which solves

$$E\left[\beta\left(\alpha \exp(a_t) K_{t-1}^{\alpha-1} + (1-\delta)\right)\left(\frac{C_{t+1}}{C_t}\right)^{-\nu} \,\middle|\, K_{t-1}, a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2\right] = 1 \quad a.s. \quad t = 0, 1, \ldots \tag{3}$$

where the exogenous states $(a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$, observed by the agent at time t, follow the law of motion:

$$\begin{aligned}
a_{t+1} &= \widehat{\theta}_t a_t + \sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}\,\epsilon_{t+1} \\
\widehat{\theta}_{t+1} &= \widehat{\theta}_t + \frac{a_t P_t}{\sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}}\epsilon_{t+1} \\
P_{t+1} &= \sigma_\theta^2 + \frac{\sigma_{a,t+1}^2}{a_t^2 P_t + \sigma_{a,t+1}^2} P_t \\
\sigma_{a,t+1}^2 &= \exp(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t)
\end{aligned} \tag{4}$$

for every $t = 0, 1, \ldots$, with $\{\epsilon_t\}_{t=1}^{+\infty} \sim$ iid $N(0,1)$, and where $\widehat{\theta}_0 = \bar{\theta}$, $P_0 = \bar{P}$, σ_{a0}^2 given.[3] In the equations above, $\widehat{\theta}_t = E(\theta_t \mid a_t, \ldots, a_0)$ is the conditional expectation of the latent parameter θ_t given the agent's information, while $P_t = E[(\theta_t - \widehat{\theta}_t)^2 \mid a_t, \ldots, a_0]$ is its conditional variance.

The second question we ask is: is it possible to approximate the optimal policy $C_t = C(K_{t-1}, , a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$ in (3) using the *perturbation* method? Specifically, suppose that in the original incomplete information model (1)-(2) we assume that

$$\{\epsilon_{at}\}_{t=1}^{+\infty} \sim \text{ iid } N(0, \sigma^2) \quad \text{and} \quad \{\epsilon_{\theta t}\}_{t=1}^{+\infty} \sim \text{ iid } N(0, \sigma^2), \quad \text{with} \quad \sigma \to 0. \tag{5}$$

[2]Cogley and Sargent (2005) impose further conditions that guarantee stationarity of $\{a_t\}_{t=0}^{+\infty}$. This stationarity requirement is left out for simplicity.

[3]Detailed derivations for the motivating example are in Appendix D.

The limit experiment in (5) corresponds to the one in which the noise becomes small and the variables $(a_t, \theta_t, \sigma_{at})$ become deterministic. Note however that θ_t is unobservable for any $\sigma > 0$ which raises the question of what are the state variables of the deterministic limit of the model (1)-(2)? The second main result of our paper is to show that under (5), the law of motion of the exogenous states of the full information equivalent keeps exactly the same form as in (4) but with innovations which now follow

$$\{\epsilon_t\}_{t=1}^{+\infty} \sim N(0, \sigma^2) \tag{6}$$

While the law of motion of the exogenous states does not change, their interpretation does: we still have $\widehat{\theta}_t = E(\theta_t \mid a_t, \ldots, a_0)$ but now $P_t = \sigma^{-2} E[(\theta_t - \widehat{\theta}_t)^2 \mid a_t, \ldots, a_0]$. Combing (4) and (6) immediately shows that in the deterministic limit $\sigma = 0$, the exogenous model states follow:

$$a_{t+1} = \widehat{\theta}_t a_t, \quad \widehat{\theta}_{t+1} = \widehat{\theta}_t, \quad P_{t+1} = \sigma_\theta^2 + \frac{\sigma_{a,t+1}^2}{a_t^2 P_t + \sigma_{a,t+1}^2} P_t, \quad \sigma_{a,t+1}^2 = \exp(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t). \tag{7}$$

The above result is non trivial: it shows that the deterministic limit of a model with learning is *not* the same as that of a model in which the agent observes all the states. Indeed, in a complete information model in which we assume that the agent observes the drifting parameter θ_t in (2), the deterministic limit is simply: $a_{t+1} = \theta_t a_t$, $\theta_{t+1} = \theta_t$ and $\sigma_{a,t+1}^2 = \exp(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t)$. Even if we were to set $\bar{\theta} = \theta_0$, the states of the deterministic model with observed θ_t differ from the ones in which the agent learns about θ_t: the difference is in the variable P_t which captures the evolving precision of the learning process.

We now outline how the two main results generalize to more complicated setups. First, we enlarge the class of dynamic stochastic choice models to include those considered in Stokey and Lucas (1989). While we broadly follow the setup defined in their Chapter 9.1 we pay particular attention to the assumptions on the stochastic elements of the model. Our key objective is to avoid making bounded support assumptions which would put the framework at odds with the Gaussian assumptions that follow.

Second, we broaden the information structure in equations (1)-(2) to a new class of models, which we call *conditional Gauss Markov*. As suggested by their name, the conditional Gauss Markov models have two important properties: they belong to the class of conditionally Gaussian models introduced by Liptser and Shiryayev (1978) and are thus characterized by a finite number of sufficient statistics. Moreover, they are Markovian in the sense that those sufficient statistics can be calculated from their own past values and the newly observed variables.[4] The conditional Gauss Markov family contains numerous information structures of interest to economists, including: time-varying parameter VARs à la Cogley and Sargent (2005) and Primiceri (2005), models with conditional heteroskedasticity à la ARCH/GARCH and their multivariate extensions, threshold models, smooth transition models à la Auerbach and Gorodnichenko (2012), growth-at-risk models à la Adrian, Boyarchenko, and Giannone (2019).

Similar to the optimality analysis in our motivating example, we show that any of our conditional Gauss Markov dynamic choice models (which are models with incomplete information) has a full infor-

[4] In statistical decision theory (see, e.g. Bahadur, 1954), such statistics are called *transient*.

mation equivalent, for which we provide necessary and sufficient optimality conditions (Euler equations and transversality conditions). Moreover, we show that the full information equivalent model has a well defined small noise limit which lends itself to regular perturbation analysis in these models. This result is to the best of our knowledge new to the perturbation literature which has up to now focused on full information models (see, e.g. Fleming, 1971; Judd, 1996; Jin and Judd, 2002).

3 Dynamic Choice Models with Gaussian Signals

To setup the framework we first consider the dynamic choice problems with learning. The setup is such that it can accommodate the models in which some components of the state vector are imperfectly observed, as well as the models in which the decision maker has signals about the future values of some components of the state.

3.1 Stokey-Lucas Setup

The optimal choice problems that we consider are similar to the ones described in Stokey and Lucas (1989), Ch 9.1 with one key difference: the structure of the information available to the decision maker at the time the decisions are made. We first recall the main features of Stokey-Lucas framework. Let x denote the *endogenous state variable* taking values in the set $X \subset \mathbb{R}^{n_x}$ and let z denote the *exogenous variable* taking values in the set $Z \subset \mathbb{R}^{n_z}$; realizations of both x and z are observed by the decision maker.[5] The decision problem is the one in which in each period t, the decision maker chooses the value for the endogenous state variable in the subsequent period.[6] The constraints on the choice of x_{t+1} are described by the *feasible set* $\Gamma(x,z) \subset X$ where $\Gamma : X \times Z \to X$. In what follows, let A be the graph of Γ, $A = \{(y,x,z) \in X \times X \times Z \mid y \in \Gamma(x,z)\}$, let $F : A \to \mathbb{R}$ be the per-period return function, and let $0 < \beta < 1$ be the discount factor. The goal of the decision maker is to

$$\text{maximize} \quad E\left[\sum_{t=0}^{\infty} \beta^t F(x_{t+1}, x_t, z_t)\right], \tag{8}$$

with x_0 given.

Time starts at $t = 0$ and we assume that x_0 is given (deterministic). Let us denote by I_t the information vector available to the decision maker at time $t = 0, 1, \ldots,$

$$I_t \equiv \{z_t, x_t, \ldots, z_0, x_0\}.$$

[5] Foreshadowing the discussion to follow, one can think of z as consisting of observed exogenous states as well as signals about latent exogenous states.

[6] Models in which optimality can be defined using the Euler equation require that the dimension of the action (or decision or control) variable c be the same as the dimension of the endogenous state variable x. For instance, if one starts with the state transition equation of the type $x_{+1} = g(x,c,z)$ where x_{+1} denotes the one period ahead value of the endogenous state variable x, and per-period utilities of the form $u(c)$, then implicit here is the assumption that the state transition equation can be solved for c in terms of the other variables, $c = h(x_{+1}, x, z)$, which can then be plugged into $u(c)$ to yield per-period returns $F(x_{+1}, x, z) = u(h(x_{+1}, x, z))$.

It is assumed here that the decision maker knows all the parameters of the decision problem in (8), such as β, F and Γ, even though we do not explicitly include these values in her information set. The starting point of optimality analysis is to reformulate the original choice problem in terms of policies. Rather than choosing a sequence of endogenous states $(x_1, x_2, \ldots)$, the decision maker chooses a (non-randomized) plan $\widetilde{\pi} = (\widetilde{\pi}_0, \widetilde{\pi}_1, \ldots)$, which is a sequence of maps $\widetilde{\pi}_t$, that assign a value for x_{t+1} to all the information I_t available to the decision maker at time t. Stating the decision maker's problem in terms of plans rather than sequences of controls ensures that she utilizes the full information available to her at every point in time. Starting at time $t = 0$, the decision maker takes her initial information set $I_0 = \{z_0, x_0\}$ and she chooses $x_1 = \widetilde{\pi}_0(I_0)$ which yields the return $F(\widetilde{\pi}_0(I_0), x_0, z_0)$. Time moves to $t = 1$, the decision maker observes z_1, adds it to her information set and chooses x_2 as a function of all the variables in $I_1 = \{z_1, x_1, z_0, x_0\}$. The return $F(\widetilde{\pi}_1(I_1), \widetilde{\pi}_0(I_0), z_1)$ is realized and its discounted value added to the previous reward. Time moves on and the entire sequence repeats. Thus, for each t, the policy $\widetilde{\pi}_t : \mathcal{I}_t \to X$ is a (Borel) measurable function that to each history $I_t \in \mathcal{I}_t = Z \times X \times \cdots \times Z \times X$ ($2(t+1)$ factors) assigns $x_{t+1} = \widetilde{\pi}_t(I_t)$.

We now need to discuss how the expectation in (8) is calculated. This requires setting up the probabilistic background. Assume then that the observed data are realizations of a stochastic process $\{z_t\}_{t=0}^{+\infty}$ defined on a complete probability space $(\mathcal{Z}^\infty, \mathcal{B}_{\mathcal{Z}}^\infty, P)$, where $\mathcal{Z}^t \equiv Z \times \cdots \times Z$ ($t+1$ factors) and $\mathcal{B}_{\mathcal{Z}}^t = \mathcal{B}(\mathcal{Z}_t)$ is the associated Borel σ-algebra, $t = 0, 1, \ldots, \infty$. Denote by $z^t = (z_0, \ldots, z_t)$ $(t = 0, 1, \ldots)$ the finite history of the process up to time t and by μ_t the corresponding finite dimensional distribution. Explicit assumptions on μ^t, which will be our point of departure from Stokey-Lucas setup, shall be made clear in the following sections.

The plan $\widetilde{\pi} = (\widetilde{\pi}_0, \widetilde{\pi}_1, \ldots)$, together with the distribution of $\{z_t\}_{t=0}^{+\infty}$, defines for each initial $x_0 \in X$ given, a distribution of the process $\{x_t^{\widetilde{\pi}}\}_{t=1}^{+\infty}$ given by:

$$\begin{aligned} x_1^{\widetilde{\pi}} &= \widetilde{\pi}_0(z_0, x_0) \\ x_{t+1}^{\widetilde{\pi}} &= \widetilde{\pi}_t(z_t, x_t^{\widetilde{\pi}}, \ldots, z_1, x_1^{\widetilde{\pi}}, z_0, x_0), \quad t = 1, 2, \ldots. \end{aligned}$$

It is clear from this construction that with $x_0 \in X$ given, the policies $\widetilde{\pi}_t$ are $\mathcal{B}_{\mathcal{Z}}^t$-measurable, $t = 0, 1, \ldots$. In other words, each $\widetilde{\pi}_t$ can be written as a function of z^t (and x_0) alone. Let us then define the sequence $\pi = (\pi_0, \pi_1, \ldots)$ where each $\pi_t : \mathcal{Z}^t \to X$ is given by:

$$\pi(z_0) \equiv \widetilde{\pi}_0(z_0, x_0) \quad \text{and} \quad \pi_t(z^t) \equiv \widetilde{\pi}_t(z_t, x_t^{\widetilde{\pi}}, \ldots, z_1, x_1^{\widetilde{\pi}}, z_0, x_0), \quad t = 1, 2, \ldots. \tag{9}$$

Put in words, the policies π_t are obtained by repeatedly substituting past endogenous variables as functions of previous z's. Note that in order to simplify the notation we have dropped reference to x_0 in (9). The definitions that follow are now stated in terms of the plan π. In subsequent sections, we shall however go back to the original definition of the plan $\widetilde{\pi}$ as keeping the reference to the past endogenous states will be important for reduction to stationary plans.

The sequence $\pi = (\pi_0, \pi_1, \ldots)$ is said to be a *feasible plan* (from x_0) if we have

$$\begin{aligned}
&\pi_0(z_0) \in \Gamma(x_0, z_0), \quad \text{for all } z_0 \in Z, \\
&\pi_t(z^t) \in \Gamma(\pi_{t-1}(z^{t-1}), z_t), \quad \text{for all } z^t \in \mathcal{Z}^t, \quad t = 1, 2, \ldots.
\end{aligned}$$

The set of all feasible plans is denoted by Π. Note that the set Π depends on x_0 as well as all the parameters (β, Γ, F) of the decision problem, though we do not make this dependence explicit. Hereafter, we assume that the set Π is nonempty.[7] Note that for any $\pi \in \Pi$, the period-t reward $F(x^\pi_{t+1}, x^\pi_t, z_t)$ in (8) is $\mathcal{B}^t_{\mathcal{Z}}$-measurable. Now for each $\pi \in \Pi$ define $v : \Pi \to \mathbb{R} \cup \{-\infty\}$ by

$$\begin{aligned}
v(\pi) &= \lim_{T\to\infty} \int_Z F(\pi_0(z_0), x_0, z_0)\mu^0(dz_0) + \sum_{t=1}^{T} \beta^t \int_{\mathcal{Z}^t} F\left(\pi_t(z^t), \pi_{t-1}(z^{t-1}), z_t\right) \mu^t(dz^t) \\
&= E\left[\sum_{t=0}^{+\infty} \beta^t F(x^\pi_{t+1}, x^\pi_t, z_t)\right]. \qquad (10)
\end{aligned}$$

That $v(\pi)$ is well defined requires assumptions (stated in Appendix A) that mirror those of Stokey and Lucas (1989), Ch 9.1. Finally, the decision problem in (8) then amounts to finding a feasible plan $\pi^* \in \Pi$ such that

$$v(\pi^*) = \sup_{\pi\in\Pi} v(\pi). \qquad (11)$$

Such π^* is called *optimal*.

3.2 Euler Equation

Up to this point, the setup is the one described by Stokey and Lucas (1989), Ch 9.1. We now describe the information structure which is the point of departure from the traditional Stokey-Lucas framework. The analysis in Stokey and Lucas (1989), Ch 9.1 proceeds under the additional assumption that the exogenous variable z follows a Markov process. This assumption is no longer tenable if, for example, one allows for some components of the exogenous state to be unobserved by the decision maker. The assumption also breaks down if the decision maker receives additional signals that help her predict the future values of the exogenous state. In both cases, the exogenous observed variables z follow a *hidden* Markov process.[8] The question then becomes: how can one analyze the decision problem in (8) when z is *not* Markovian?

To start, recall again that the decision problem in (8) is equivalent to the optimization problem in (11). There are two ways to approach the latter: first is to use the principle of Dynamic Programming. The idea here is to connect the supremum function $v(\pi^*)$ with solutions to the Bellman functional equation, then use the latter to gain insights into the properties of π^*. This is the approach taken in Stokey and Lucas (1989) for complete information problems. This is also the approach taken by

[7] For completeness, all our assumptions are formally stated in Appendix A.

[8] We make precise assumptions on how this hidden Markov process is generated in the following section.

Bertsekas and Shreve (1978) to analyze stochastic control problems with partial information, and derive their full information equivalents. The second approach to solving (11) is to use calculus of variations; see, e.g. Fleming and Rishel (1975), Section 9.5 in Stokey and Lucas (1989), Santos (2000), Section 16.3 in Acemoglu (2009), Gonzalez-Sanches and Hernandez-Lerma (2013).[9] The idea here is to connect the optimal policy π^* with solutions to the Euler equations. While both approaches are theoretically equally valid, it is worth to note that most works that analyze approximate solutions to decision problems such as (8), either via perturbation or projection methods (see, e.g. Judd (1996)) take the Euler equation as the starting point. For this reason, we pursue the Euler equation approach here.

To set up the calculus of variations, let Λ be the vector space of all plans (i.e. sequences of $\mathcal{B}_{\mathcal{Z}}^t$-measurable functions),

$$\Lambda = \{(\lambda_0, \lambda_1, \ldots) \mid \lambda_t : \mathcal{Z}^t \to \mathbb{R}^{n_x}, t = 0, 1, \ldots\}$$

with the standard addition and scalar multiplication. We need the following definition of an *internal plan* (see p.3 in Fleming and Rishel (1975)).

Definition 1. *Let $\pi \in \Pi$ and $\lambda \in \Lambda$. We say that π is an **internal plan** in the direction λ if there exists a real number $\epsilon_\lambda > 0$ such that $\pi + \epsilon\lambda \in \Pi$ for all $|\epsilon| < \epsilon_\lambda$.*

Next, we impose differentiability of the return function $F(\cdot, \cdot, z)$ in (8) as well as a condition that ensures that differentiating under the integral in (10) is legitimate; see our Assumptions 5 and 6 in Appendix A. Under those assumptions the function $v(\cdot)$ is Gateaux-differentiable (or directionally differentiable) at any internal point in Π. Specifically, let $\pi \in \Pi$ be an internal plan in the direction $\lambda \in \Lambda$, and let $\delta v(\pi, \lambda)$ denote the Gateaux differential of the function $v(\cdot)$ at π in the direction λ,

$$\delta v(\pi, \lambda) = \frac{dv}{d\epsilon}(\pi + \epsilon\lambda)\bigg|_{\epsilon=0}.$$

In calculus of variations, $\delta v(\pi, \lambda)$ is called the first variation. We then have the following lemma whose proof is in Appendix B.

Lemma 1. *Let Assumptions 1-6 hold. Consider the function $v(\cdot)$ in (10) and let $\pi \in \Pi$ be an internal plan in the direction $\lambda \in \Lambda$. Then the Gateaux differential $\delta v(\pi, \lambda)$ of the function $v(\cdot)$ at π in the direction λ exists and is given by*

$$\delta v(\pi, \lambda) = E\left[\frac{\partial F}{\partial x_{+1}}(x_1^\pi, x_0, z_0) \cdot x_1^\lambda + \sum_{t=1}^{+\infty} \beta^t \left(\frac{\partial F}{\partial x_{+1}}(x_{t+1}^\pi, x_t^\pi, z_t) \cdot x_{t+1}^\lambda + \frac{\partial F}{\partial x}(x_{t+1}^\pi, x_t^\pi, z_t) \cdot x_t^\lambda\right)\right] \quad (12)$$

With the Gateaux differential in hand, deriving the Euler equation is straightforward: if $v(\pi)$ has a supremum at an internal point π^* in the direction λ, then a first order condition $\delta v(\pi^*, \lambda) = 0$ holds (see, e.g. Theorem 2.1 in Fleming and Rishel (1975)). The Euler equation follows by considering directions

[9] Note that none of these works deal with the discrete time dynamic stochastic choice problems with incomplete information, which is the class of problems that we are interested here. In addition, the assumptions used Acemoglu (2009) (finite or compact Z), for example, rule out the linear Gaussian dynamics for the exogenous states that we are mostly interested here.

which are of the form $\lambda = (0, \ldots, 0, \nu_\tau, 0, \ldots)$ for $\tau = 0, 1, \ldots$ where $\nu_\tau : \mathcal{Z}^\tau \to \mathbb{R}^{n_x}$ is any measurable function. To ensure that π^* is an internal point in all such directions, we need Assumption 10 found in Appendix A. Imposing the Euler equation in every time period does not exhaust all the implications of the first order condition $\delta v(\pi^*, \lambda) = 0$; there remains a limit transversality condition. It is the combination of the Euler equation *and* the transversality condition that is equivalent to $\delta v(\pi^*, \lambda) = 0$. That the latter is also a sufficient condition for optimality of π^* follows by further restricting our attention to convex feasible sets and concave functions $v(\cdot)$; the precise conditions can be found in Assumptions 7-9 in Appendix A. The main result of this section can then be stated as follows.

Theorem 1. *Let Assumptions 1-9 hold. If a plan $\pi^* \in \Pi$ satisfies the Euler equation*

$$\frac{\partial F}{\partial x_{+1}}(x_{t+1}^{\pi^*}, x_t^{\pi^*}, z_t) + \beta E\left[\frac{\partial F}{\partial x}(x_{t+2}^{\pi^*}, x_{t+1}^{\pi^*}, z_{t+1})\middle| z^t\right] = 0 \quad a.s. \quad t = 0, 1, \ldots$$

and the transversality condition

$$\lim_{T\to\infty} E\left[\beta^T \frac{\partial F}{\partial x_{+1}}\left(x_{T+1}^{\pi^*}, x_T^{\pi^*}, z_T\right) \cdot x_{T+1}^{\pi - \pi^*}\right] = 0, \quad \textit{for any } \pi \in \Pi,$$

then π^* is the unique optimal plan in Π. Conversely, if $\pi^* \in \Pi$ is an optimal plan for which Assumption 10 holds, then π^* satisfies the above Euler equation and transversality condition.

3.3 Conditional Gauss Markov Information Structure

First, let us assume that by appropriate transformations, the exogenous variable z in our decision problem can be made to take values in the entire Euclidean space $\mathbb{R}^{n_z}$. For instance, if (a component of) z is always non-negative (such as e.g. the TFP or gross return on assets) then use the transformation $z \mapsto \ln z$. We shall then assume that in the stochastic decision model described in Sections 3.1-3.2 the exogenous variable z follows a nonlinear state space model driven by an *unobserved exogenous state variable* ξ taking values in $\mathbb{R}^{n_\xi}$ whose *transition equation* is given by:

$$\xi_{t+1} = A(z^t)\xi_t + u(z^t) + B(z^t)w_{t+1}, \quad t = 0, 1, \ldots, \ \xi_0 \sim N(\bar{\xi}_0, \bar{\Sigma}_0) \tag{13}$$

with $A(z^t) \in \mathbb{R}^{n_\xi \times n_\xi}$, $u(z^t) \in \mathbb{R}^{n_\xi}$, $B(z^t) \in \mathbb{R}^{n_\xi \times n_w}$, and with a covariance matrix $\bar{\Sigma}_0 \in \mathbb{R}^{n_\xi \times n_\xi}$ that is positive semidefinite. The variable z, which is observed by the decision maker, is related to ξ through an *observation equation* given by:

$$z_t = C(z^{t-1})\xi_t + v(z^{t-1}) + D(z^{t-1})w_t, \quad t = 0, 1, \ldots \tag{14}$$

with $C(z^{t-1}) \in \mathbb{R}^{n_z \times n_\xi}$, $v(z^{t-1}) \in \mathbb{R}^{n_z}$, $D(z^{t-1}) \in \mathbb{R}^{n_z \times n_w}$.[10] The n_w-vector w_t collects all the structural shocks and signal errors in the system and we assume it to be *Gaussian*,

$$w_t \sim \text{ iid } N(0, \mathrm{Id}).$$

[10] By convention $z^{-1} = \emptyset$ so for $t = 0$, $(C(z^{-1}), D(z^{-1}), v(z^{-1})) = (C, D, v)$ are constant matrices.

Note that setting the covariance matrix of w_t to be an identity matrix is without loss of generality if one is willing to assume that all the components of w_t are linearly independent; in this case, the Cholesky factor of the covariance matrix of w_t can simply be subsumed in the matrices $B(z^t)$ and $D(z^{t-1})$. No restrictions are placed on the dimensions n_ξ, n_z and n_w. Moreover, not all components of w_t are required to appear in every equation.

The system matrices $(A(z^t), B(z^t), C(z^t), D(z^t), u(z^t), v(z^t))$ are all allowed to be random and adapted to the observation process, so long as all their entries are finite with probability one. The sequence $\{w_t\}_{t=1}^{+\infty}$ and ξ_0 are assumed to be independent; Assumptions 11 and 12 in Appendix A summarize the restrictions just stated. As before, we assume that the decision maker knows all the parameters of the model, so that the values of $(\bar{\xi}_0, \bar{\Sigma}_0)$ in the equations above are deterministic and known. Similarly, once z^t (resp. z^{t-1}) is observed, $A(z^t), B(z^t), u(z^t)$ (resp. $C(z^{t-1}), D(z^{t-1}), v(z^{t-1})$) are known to the decision maker. Systems of this type are known as *conditionally Gaussian* systems (Liptser and Shiryayev, 1978; Chen, Kumar, and van Schuppen, 1989).

It is important to emphasize that the equations (13)-(14) are inherently *nonlinear* and that they generate a stochastic process $\{(\xi_t, z_t)\}_{t=0}^{+\infty}$ which is generally *not* Gaussian. Still, this process has the important property that for each $t = 0, 1, \ldots$, the conditional distribution of ξ_t given z^t is Gaussian whose parameters can be calculated recursively using the familiar Kalman-Bucy filtering equations, formally stated below. To further emphasize the nonlinear character of the system in (13)-(14), let us compare it to the well known case in which the system matrices (A, B, C, D, u, v) follow a deterministic sequence (i.e. $A(z^t) = A_t$ deterministic, etc for all the other system components); this is the original Kalman-Bucy setup (see, e.g. Harvey, 1989) and also a special case of the setup considered here. Whereas in the Kalman-Bucy scheme the process $\{(\xi_t, z_t)\}_{t=0}^{+\infty}$ is Gaussian, and the resulting filtering equations linear, in the conditionally Gaussian case considered here, the filtering equations are generally speaking nonlinear.

Another important difference between the Gaussian (i.e. Kalman-Bucy) setup and the conditionally Gaussian setup is in their ability to accommodate several desirable features commonly found in macroeconomic models, such as (i) conditional heteroskedasticity of shocks or errors in the transition and/or observation equations, so long as the heteroskedasticity is of known form (e.g. GARCH-type heteroskedasticity is allowed but stochastic volatility is ruled out), (ii) parameter switching driven by the observables, or (iii) learning about constant but unknown parameters of the exogenous process. While all those features are compatible with the conditionally Gaussian system in (13)-(14), they are not permitted in the original Kalman-Bucy setup.

Lastly, note that our setup also nests the case in which the exogenous state variable $\xi = z$ is perfectly observed by the decision maker: it suffices to let $C(z^{t-1}) \equiv \text{Id}$, $v(z^{t-1}) \equiv 0$, and $D(z^{t-1}) \equiv 0$ in equation (14). Beyond the case with complete information, we allow for two additional sets of variables in the model: (i) exogenous state variables ξ that remain latent to the decision maker at the time the decisions are made, and (ii) signals s related to ξ that the decision maker observes prior to taking the decision.

We can now give an explicit expression for the distribution μ^t of z^t for $t = 0, 1, \ldots$. For this, start at time $t = 0$ with initial state ξ_0 as in (13), and note that the initial observation equation (14) implies

that (ξ_0, z_0) are jointly Gaussian,

$$\begin{pmatrix}\xi_0\\ z_0\end{pmatrix} \sim N\left(\begin{pmatrix}\bar{\xi}_0\\ C\bar{\xi}_0 + v\end{pmatrix}, \begin{pmatrix}\bar{\Sigma}_0 & \bar{\Sigma}_0 C'\\ C\bar{\Sigma}_0 & C\bar{\Sigma}_0 C' + DD'\end{pmatrix}\right) \tag{15}$$

This implies in particular, that the conditional distribution of ξ_0 given z_0 is Gaussian with parameters $\widehat{\xi}_0 \equiv E(\xi_0 \mid z_0)$ and $\Sigma_0 \equiv E[(\xi_0 - \widehat{\xi}_0)(\xi_0 - \widehat{\xi}_0)' \mid z_0]$ given by

$$\widehat{\xi}_0 = \bar{\xi}_0 + \bar{\Sigma}_0 C'(C\bar{\Sigma}_0 C' + DD')^+(z_0 - C\bar{\xi}_0 - v) \text{ and } \Sigma_0 = \bar{\Sigma}_0 - \bar{\Sigma}_0 C'(C\bar{\Sigma}_0 C' + DD')^+ C\bar{\Sigma}_0, \tag{16}$$

where for any matrix P we use P^+ to denote the pseudoinverse (or Moore-Penrose inverse) of P.[11]

Next, one proceeds by induction: assuming that for some t, (ξ_t, z_t) is conditionally Gaussian given z^{t-1}, the equations (13)-(14) then imply that (ξ_{t+1}, z_{t+1}) is conditionally Gaussian given z^t. Thus, ξ_{t+1} given z^{t+1} is also conditional Gaussian with conditional mean $\widehat{\xi}_{t+1} \equiv E(\xi_{t+1} \mid z^{t+1})$ and conditional covariance $\Sigma_{t+1} \equiv E[(\xi_{t+1} - \widehat{\xi}_{t+1})(\xi_{t+1} - \widehat{\xi}_{t+1})' \mid z^{t+1}]$ which can be calculated recursively from

$$\widehat{\xi}_{t+1} = A(z^t)\widehat{\xi}_t + u(z^t) + K(z^t, \Sigma_t) \underbrace{\left(z_{t+1} - C(z^t)A(z^t)\widehat{\xi}_t - C(z^t)u(z^t) - v(z^t)\right)}_{e_{t+1} \equiv z_{t+1} - E(z_{t+1}|z^t)} \tag{17}$$

where

$$K(z^t, \Sigma_t) \equiv \left[A(z^t)\Sigma_t A(z^t)'C(z^t)' + B(z^t)B(z^t)'C(z^t)' + B(z^t)D(z^t)'\right]\Sigma_e(z^t, \Sigma_t)^+ \tag{18}$$

$$\Sigma_e(z^t, \Sigma_t) \equiv C(z^t)A(z^t)\Sigma_t A(z^t)'C(z^t)' + \left[C(z^t)B(z^t) + D(z^t)\right]\left[C(z^t)B(z^t) + D(z^t)\right]' \tag{19}$$

and

$$\Sigma_{t+1} = A(z^t)\Sigma_t A(z^t)' + B(z^t)B(z^t)' - K(z^t, \Sigma_t)\Sigma_e(z^t, \Sigma_t)K(z^t, \Sigma_t)' \tag{20}$$

for $t = 0, 1, \ldots$ with $\widehat{\xi}_0$ and Σ_0 as given in (16).[12]

Several remarks are in order. First, note that the coefficients of the matrices K, Σ_e, Σ_t do not depend explicitly on the system vectors u, v. Second, unlike in the Kalman-Bucy case, those matrices now depend on the observations z^t; this is because the system matrices A, B, C, D in (13)-(14) are themselves functions of z^t. In particular, both the Kalman gain $K(z^t, \Sigma_t)$ as well as Σ_{t+1} the conditional covariance of ξ_{t+1} are now functions of the observations z^t. Note that the latter is predetermined, in the sense that the covariance conditional on z^{t+1} only depends on data up to time t; thus Σ_{t+1} can be computed at time t. Third, note that in (17), the conditional expectation $\widehat{\xi}_t$ is a nonlinear function of z^t; this again is in contrast with the Kalman-Bucy scheme in which $\widehat{\xi}_t$ is a linear function of the observations z^t. Last, unless the system matrices A, B, C, D are deterministic, the unconditional distribution of the

[11] Recall that any matrix $P \in \mathbb{R}^{n \times n}$ has a unique pseudoinverse (or Moore-Penrose inverse) P^+. The updating formulas that use pseudoinverses can be found e.g. in Ch 13 of Liptser and Shiryayev (1978) or Chen, Kumar, and van Schuppen (1989).

[12] For reasons that will become apparent in Section 3.4, we keep explicit the dependence of K, Σ_e on Σ_t.

latent exogenous variable ξ_t is not Gaussian.

With the filtering formulas (17)-(20) in hand, we now proceed to calculate the conditional distributions of z_{t+1} given z^t for $t = 0, 1, \ldots$. Combined with the marginal distribution of z_0 in (15), they will then uniquely define the finite dimensional distributions μ_t of z^t, for $t = 0, 1, \ldots$. For this, consider the innovations e_{t+1} defined in (17) and assume that the system matrices B, C, D are such that the covariance matrix $\Sigma_e(z^t, \Sigma_t)$ is positive definite with probability one, for $t = 0, 1, \ldots$ (see Assumption 13 in Appendix A). Let

$$\epsilon_{t+1} \equiv \Sigma_e(z^t, \Sigma_t)^{-1/2} e_{t+1} \tag{21}$$

where for any positive definite $n \times n$ matrix Σ we use the notation $\Sigma^{1/2}$ to denote the Cholesky square root of Σ, that is, the unique lower diagonal $n \times n$ matrix with positive diagonal entries such that $\Sigma^{1/2}(\Sigma^{1/2})' = \Sigma$; see Appendix C for details.

The properties of $\{\epsilon_t\}_{t=1}^{+\infty}$ are given in the lemma below.

Lemma 2. *Let Assumptions 11-13 hold. Then, for $t = 0, 1, \ldots$*

$$\begin{aligned} \widehat{\xi}_{t+1} &= A(z^t)\widehat{\xi}_t + u(z^t) + K(z^t, \Sigma_t)\Sigma_e(z^t, \Sigma_t)^{1/2}\epsilon_{t+1} \\ z_{t+1} &= C(z^t)A(z^t)\widehat{\xi}_t + C(z^t)u(z^t) + v(z^t) + \Sigma_e(z^t, \Sigma_t)^{1/2}\epsilon_{t+1}, \end{aligned} \tag{22}$$

where $\{\epsilon_t\}_{t=1}^{+\infty} \sim$ iid $N(0, Id)$. Moreover, letting $\mathcal{B}_{\mathcal{E}}^t$ denote the σ-algebra generated by $(z_0, \epsilon_1, \ldots, \epsilon_t)$, we have $\mathcal{B}_{\mathcal{Z}}^t = \mathcal{B}_{\mathcal{E}}^t$ for $t = 0, 1, \ldots$.

The representation above is called the *innovations* representation of the original state space system (13)-(14). While its validity is well known under the Kalman-Bucy scheme with deterministic system matrices (see, e.g., Anderson and Moore, 1979), the fact that the representation still holds with random system matrices as in (13)-(14) is less well understood; thus we formally state the above results as a lemma whose proof can be found in Appendix B. Lemma 2 has several interesting implications that are worth pointing out. First, it shows that the innovations ϵ_{t+1} are not necessarily an affine transformation of z^{t+1} which is the case obtained under the Kalman-Bucy scheme. Here, ϵ_{t+1} is a nonlinear function of past z^t; what matters is that ϵ_{t+1} is a linear transformation of z_{t+1} *conditional* on the past z^t. Second, the equivalence of σ-algebras means that the innovation process $\{\epsilon_t\}_{t=1}^{+\infty}$ and z_0 contain the same information as the observable process $\{z_t\}_{t=0}^{+\infty}$. In particular, this implies that for every $t = 1, 2 \ldots$, ϵ_t is $\mathcal{B}_{\mathcal{Z}^t}$-measurable, i.e. ϵ_t is a function of z^t and so is observed by the decision maker at time t. The equivalence of σ-algebras also implies that ϵ_t is independent of z^s for any $s < t$. Thus, ϵ_t consists of that part of z_t containing new information not carried in z^{t-1}, which is the usual interpretation of an innovation (see, e.g. Ch 5.3 in Anderson and Moore, 1979). This last property is important as it allows us to calculate the conditional distributions of z_{t+1} given z^t, which equal

$$z_{t+1} \mid z^t \sim N\left(C(z^t)A(z^t)\widehat{\xi}_t + C(z^t)u(z^t) + v(z^t), \Sigma_e(z^t, \Sigma_t)\right), \quad \text{for } t = 1, 2, \ldots$$

Combining the above with the distribution of z_0 given in (15) then determines the finite dimensional

distributions μ^t of z^t for every $t = 0, 1, \ldots$.

Now, let us consider the structure of the recursive equations (22) and (18)-(20). It is clear from these equations that if the system matrices A, B, C, D, u, v are deterministic, then $(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1})$ has a *Markov* property, i.e. $P\big(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1} \mid z^t\big) = P\big(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1} \mid z_t, \widehat{\xi}_t, \Sigma_t\big)$. This Markov property will play an important role in the developments that follow as it will allow us to reformulate the original decision problem in (8) into an equivalent problem with fully observed states that are Markovian. In order to preserve the Markov property when the system matrices A, B, C, D, u, v are random, we hereafter restrict our attention to a subfamily of nonlinear state space models (13)-(14) which we call *conditionally Gauss Markov*.

Definition 2 (Conditionally Gauss Markov System)**.** *The nonlinear state space system in* (13)-(14) *is said to be conditionally Gauss Markov if there exists a sequence of random variables $m_t \in \mathbb{R}^{n_m}$ such that for $t = 1, 2, \ldots$, (i) m_t is $\mathcal{B}_{\mathcal{Z}}^{t-1}$ measurable, (ii) m_t is* recursive *in the sense that*

$$m_t = f(m_{t-1}, z_{t-1}), \tag{23}$$

for some function $f : \mathbb{R}^{n_m} \times \mathbb{R}^{n_z} \to \mathbb{R}^{n_m}$ and some m_0 given (deterministic), and (iii) such that A, B, C, D, u, v depend on z^t only through (z_t, m_t), i.e.

$$\begin{aligned} &A(z^t) = A(z_t, m_t),\ u(z^t) = u(z_t, m_t),\ B(z^t) = B(z_t, m_t), \\ &C(z^t) = C(z_t, m_t),\ v(z^t) = v(z_t, m_t),\ D(z^t) = D(z_t, m_t). \end{aligned}$$

Put in words, the state space system (13)-(14) is conditionally Gauss Markov if it is conditionally Gaussian *and* Markov. The restrictions imposed in Definition 2 are sufficient to obtain Markovianity, as we now have $P\big(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, m_{t+1} \mid z^t\big) = P\big(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, m_{t+1} \mid z_t, \widehat{\xi}_t, \Sigma_t, m_t\big)$, that is $(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, m_{t+1})$ has a Markov property. Note that Definition 2 leaves open the choice of m_0. Since for $t = 0$, the system matrices are only functions of z_0, we can set m_0 to any value.[13]

Before moving on, it is useful to give some examples of conditionally Gauss Markov state space models.

Example 1. ARCH models Consider an RBC model with log-productivity $a_t = \ln A_t$ that follows a random walk with drift, $\Delta a_t = g_{t-1} + \sigma_{at}\varepsilon_{at}$ where $\Delta a_t = a_t - a_{t-1}$ is observed, $g_t = (1-\rho)g^* + \rho g_{t-1} + \sigma_g \varepsilon_{gt}$ is latent, and where $\sigma_{at}^2 = \omega + \alpha(\Delta a_{t-1})^2$. When $\alpha \neq 0$, the log-productivity increments exhibit conditional heteroskedasticity of ARCH-type. To fit this model into our setup, let $\xi_t = g_t$, $z_t = \Delta a_t$, $w_t = (\varepsilon_{gt}, \varepsilon_{at})'$, and define the system matrices as

$$A = \rho,\ B = \begin{pmatrix} \sigma_g & 0 \end{pmatrix},\ u = (1-\rho)g^*,\ C = 1,\ D(z_t) = \begin{pmatrix} 0 & \sqrt{\omega + \alpha z_t^2} \end{pmatrix},\ v = 0$$

Note that in this example, the system matrix D depends only on the contemporaneous value of z and not its past. All the other system matrices are time invariant.

[13] In specific applications, this value can be chosen so as to make m_0 disappear from the formulas above; see, for instance our Example 2 below.

Example 2. GARCH models Consider a similar setup as above, but assume that the conditional heteroskedasticity is of GARCH-type, $\sigma_{at}^2 = \omega + \beta\sigma_{a,t-1}^2 + \alpha(\Delta a_{t-1})^2$. Now let $m_t = \sigma_{at}^2$ and note that m_t satisfies our restrictions (i) and (ii) above. In this case, let the system matrices A, B, u, C, v be as before and set

$$D(z_t, m_t) = \begin{pmatrix} 0 & \sqrt{\omega + \beta m_t + \alpha z_t^2} \end{pmatrix}.$$

This then satisfies condition (iii). Note that because of the recursive nature of GARCH-type volatility, the system matrix D now depends on the additional random variable m. We can set $m_0 = 0$ so that at time $t = 0$, D is only a function of z_0.

Example 3. Threshold models Consider a setup in which $\Delta a_t = g_{t-1} + \sigma_{at}\varepsilon_{at}$ is observed and $g_t = (1 - \rho_t)g^* + \rho_t g_{t-1} + \sigma_{gt}\varepsilon_{gt}$ is latent. We assume that the parameters in the model switch regimes via the following rule:

$$\{\sigma_{gt}, \rho_t, \sigma_{at}\} = \begin{cases} \{\sigma_{g1}, \rho_1, \sigma_{a1}\} \text{ for } \frac{1}{S}\sum_{s=1}^{S} \Delta a_{t-s} > \overline{a} \\ \{\sigma_{g2}, \rho_2, \sigma_{a2}\} \text{ for } \frac{1}{S}\sum_{s=1}^{S} \Delta a_{t-s} <= \overline{a} \end{cases}$$

where $\overline{a}$ is a specified threshold rule. This model can be put into our setup where $\xi_t = g_t, z_t = \Delta a_t$, $m_t = \left(z_{t-1}, z_{t-2}, z_{t-3}\right)'$, $w_t = \left(\varepsilon_{gt}, \varepsilon_{at}\right)'$:

$$A(z_t, m_t) = \rho(z_t, m_t),\ B(z_t, m_t) = \begin{pmatrix} \sigma_g(z_t, m_t) & 0 \end{pmatrix},\ u(z_t, m_t) = (1 - \rho(z_t, m_t))g^*$$
$$C = 1,\ D(z_t, m_t) = \begin{pmatrix} 0 & \sigma_g(z_t, m_t) \end{pmatrix},\ v = 0.$$

The transition rule for m_t is:

$$\begin{pmatrix} z_{t-1} \\ z_{t-2} \\ z_{t-3} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} z_{t-2} \\ z_{t-3} \\ z_{t-4} \end{pmatrix} + \begin{pmatrix} z_{t-1} \\ 0 \\ 0 \end{pmatrix}.$$

Example 4. Smooth Transition models Consider a setup similar to Example 3, but assume that the parameters in the model smoothly transition between regimes instead:

$$\sigma_{gt} = \left(1 - F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\right)\sigma_{a1} + F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\sigma_{g2}$$
$$\rho_t = \left(1 - F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\right)\rho_1 + F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\rho_2$$
$$\sigma_{at} = \left(1 - F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\right)\sigma_{a1} + F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\sigma_{a2}$$
$$F\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right) = \frac{\exp\left(-\gamma\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\right)}{1+\exp\left(-\gamma\left(\frac{1}{S}\sum_{s=1}^{S}\Delta a_{t-s}\right)\right)}, \quad \gamma > 0$$

This model can be put into our setup in the same way as Example 3. A prominent application (under full information) of this model can be found in Auerbach and Gorodnichenko (2012).

Example 5. Exogenously Evolving Time-Varying Parameters Consider the following model with an exogenously evolving time-varying autoregressive coefficient $\Delta a_t = (1-\rho_t)a^* + \rho_t \Delta a_{t-1} + \sigma_a \varepsilon_{at}$ and $\rho_t = (1-\lambda)\rho^* + \lambda\rho_{t-1} + \sigma_\rho \varepsilon_{\rho t}$. This can be put in our framework by defining $\xi_t = \rho_t, z_t = \Delta a_t, w_t = \begin{pmatrix}\varepsilon_{\rho t}, \varepsilon_{at}\end{pmatrix}'$:

$$A = \lambda,\ B = \begin{pmatrix}\sigma_\rho & 0\end{pmatrix},\ u = (1-\lambda)\rho^*,\ C(z_t) = z_t - a^*,\ D = \begin{pmatrix}0 & \sigma_a\end{pmatrix},\ v = a^*.$$

It is straightforward to generalize the above example to vector valued processes of the kind considered in Cogley and Sargent (2005) and Primiceri (2005), for example.

Example 6. Parametric Growth-at-Risk Consider the following model with correlated shifts in mean and volatility that generates growth-at-risk as in Adrian, Boyarchenko, and Giannone (2019): $\Delta a_t = \gamma x_{t-1} + g_{t-1} + \sigma_{at}\varepsilon_{at}$, $g_t = (1-\rho)g^* + \rho g_{t-1} + \sigma_{gt}\varepsilon_{gt}$, $x_t = (1-\lambda)x^* + \lambda x_{t-1} + \sigma_x \varepsilon_{xt}$, $\log\sigma_{at}^2 = \delta_{a0} + \delta_{a1}x_{t-1}$, and $\log\sigma_{gt}^2 = \delta_{g0} + \delta_{g1}x_{t-1}$. In this setup, Δa_t and x_t are observed and g_t is latent. This can be put into our framework by setting $\xi_t = g_t, z_t = \begin{pmatrix}\Delta a_t & x_t\end{pmatrix}', w_t = \begin{pmatrix}\varepsilon_{gt}, \varepsilon_{at}, \varepsilon_{xt}\end{pmatrix}'$ and defining the system matrices as

$$A = \rho,\ B(z_t, m_t) = \begin{pmatrix}\sqrt{\exp\left(\delta_{g0} + \delta_{g1}\begin{pmatrix}0 & 1\end{pmatrix}z_t\right)} & 0 & 0\end{pmatrix},\ u(z_t, m_t) = (1-\rho)g^*,$$
$$C = \begin{pmatrix}1\\0\end{pmatrix},\ D(z_t, m_t) = \begin{pmatrix}0 & \sqrt{\exp\left(\delta_{a0} + \delta_{a1}\begin{pmatrix}0 & 1\end{pmatrix}z_t\right)} & 0\\ 0 & 0 & \sigma_x\end{pmatrix},\ v(z_t, m_t) = \begin{pmatrix}\gamma\begin{pmatrix}0 & 1\end{pmatrix}z_t\\ (1-\lambda)x^* + \lambda\begin{pmatrix}0 & 1\end{pmatrix}z_t\end{pmatrix}$$

Example 7. Parameter Estimation Consider a time invariant linear Gaussian state space model $\xi_t = a\xi_{t-1} + Bw_t$, $z_t = c\xi_t + Dw_t$, with scalar ξ that is latent, scalar z that is observed, scalar a, c and $B = (b\ 0)$ and $D = (0\ d)$ ($b \neq 0$, $d \neq 0$), and assume that the decision maker does *not* know the

parameters $\theta = (a, b, c, d)$ of the model. Instead, she is estimating θ using an estimator $\widehat{\theta}_t$ that has the property of (i) being based only on the data z^{t-1} up to time $t-1$, and (ii) being recursive in the sense that $\widehat{\theta}_t$ can be calculated from $\widehat{\theta}_{t-1}$ and z_t. We can then take $m_t = \widehat{\theta}_t$ and let the decision maker use the model $\xi_t = a(m_{t-1})\xi_{t-1} + B(m_{t-1})w_t$, $z_t = c(m_{t-1})\xi_t + D(m_{t-1})w_t$ for decision making.

3.4 Full Information Equivalence

The results of Sections 3.3 and 3.2 can now be combined to define a full information equivalent of the dynamic model in Section 3.1. For this, recall that solving the decision problem in (8) is equivalent to finding a feasible plan π in Π that satisfies the Euler and transversality conditions of Theorem 1. The plans in Π have the form given in equation (9),

$$\pi(z_0) \equiv \widetilde{\pi}_0(z_0, x_0) \quad \text{and} \quad \pi_t(z^t) \equiv \widetilde{\pi}_t(z_t, x_t^{\widetilde{\pi}}, \ldots, z_1, x_1^{\widetilde{\pi}}, z_0, x_0), \quad t = 1, 2, \ldots$$

with policies $\widetilde{\pi}_t : \mathcal{I}_t \to X$ which to each information set $I_t = \{z_t, x_t, \ldots, z_0, x_0\}$ assign $x_{t+1} = \widetilde{\pi}_t(I_t)$.

As defined in Section 3.3, the exogenous process $\{z_t\}_{t=0}^{+\infty}$ is generated by the conditionally Gauss Markov state space system,

$$\begin{aligned} \xi_{t+1} &= A(z_t, m_t)\xi_t + u(z_t, m_t) + B(z_t, m_t)w_{t+1} \\ z_t &= C(z_{t-1}, m_{t-1})\xi_t + v(z_{t-1}, m_{t-1}) + D(z_{t-1}, m_{t-1})w_t, \quad t = 0, 1, \ldots \\ m_{t+1} &= f(m_t, z_t) \end{aligned} \tag{24}$$

with $\{w_t\}_{t=0}^{+\infty} \sim$ iid $N(0, \mathrm{Id})$ independent of $\xi_0 \sim N(\bar{\xi}_0, \bar{\Sigma}_0)$, initial m_0 given and deterministic, and $z_{-1} = \emptyset$ and $m_{-1} = \emptyset$ so the system matrices (C, D, v) are constant and deterministic in (24) at time $t = 0$ (see also footnote 10). Under the conditional Gauss Markov assumption (see Definition 2 and Assumption 14 in Appendix A), for every $t = 0, 1, \ldots$, the conditional distribution of ξ_t given z^t is Gaussian with parameters

$$\widehat{\xi}_t \in \mathbb{R}^{n_\xi} \quad \text{and} \quad \mathrm{vech}(\Sigma_t) \in \mathbb{R}^{n_\xi(n_\xi+1)/2} \tag{25}$$

which are determined recursively through the filtering equations (16)-(20).[14]

Now consider including in the information set I_t all the filtering information about the latent state process $\{\xi_t\}_{t=0}^{+\infty}$, i.e. let

$$\overline{I}_t \equiv \{z_t, \widehat{\xi}_t, \Sigma_t, m_t, x_t, \ldots, z_0, \widehat{\xi}_0, \Sigma_0, m_0, x_0\}, \quad t = 0, 1, \ldots \tag{26}$$

be the extended information set available to the decision maker at time t.[15] Denote by $\overline{\mathcal{I}_t}$ the space of all information sets $\overline{I}_t$, i.e. $\overline{\mathcal{I}_t} = \mathbb{R}^{n_z} \times \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\xi(n_\xi+1)/2} \times \mathbb{R}^{n_m} \times X \times \cdots \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\xi(n_\xi+1)/2} \times \mathbb{R}^{n_m} \times X$ ($5t + 5$ factors), and consider policies $\overline{\pi}_t : \overline{\mathcal{I}_t} \to X$ which to each $\overline{I}_t \in \overline{\mathcal{I}_t}$ assign $x_{t+1} = \overline{\pi}_t(\overline{I}_t)$. Let $\overline{\Pi}$

[14] Note that we a slight abuse of notation, we hereafter write that Σ_t belongs to $\mathbb{R}^{n_\xi(n_\xi+1)/2}$.

[15] As before, we do not keep reference to the model parameters $\beta, \Gamma, F, A, B, C, D, u, v, \bar{\xi}_0, \bar{\Sigma}_0$ in the definition of the information set $\overline{I}_t$ even though all those parameters are known to the decision maker.

denote the set of all plans $\overline{\pi} = (\overline{\pi}_0, \overline{\pi}_1, \ldots)$ which are feasible. It is clear that

$$\Pi \subset \overline{\Pi}$$

The following lemma, due to Rhenius (1974), shows that with any plan from $\overline{\Pi}$ the decision maker cannot obtain a greater value v than with the plans out of Π.

Lemma 3. *Let Assumptions 1-4 hold. Then for every $\overline{\pi} \in \overline{\Pi}$ there exists $\pi \in \Pi$ such that $v(\overline{\pi}) = v(\pi)$.*

In particular, lemma 3 implies that an optimal plan $\overline{\pi}^*$ in $\overline{\Pi}$ satisfies Theorem 1 with the same Euler equation and transversality condition. The key reason for considering plans in $\overline{\Pi}$ with policies $\overline{\pi}_t : \overline{I}_t \to X$, is that the information set $\overline{I}_t$ on which the policy $\overline{\pi}_t$ operates has a *Markov* structure with

$$P\big(z_t, \widehat{\xi}_t, \Sigma_t, m_t \mid \overline{I}_{t-1}\big) = P\big(z_t, \widehat{\xi}_t, \Sigma_t, m_t \mid z_{t-1}, \widehat{\xi}_{t-1}, \Sigma_{t-1}, m_{t-1}\big). \tag{27}$$

Specifically, using the innovations representation in Lemma 2 the law of motion for $\{z_t\}_{t=0}^{+\infty}$ can also be written as

$$\begin{aligned}
\widehat{\xi}_{t+1} &= A(z_t, m_t)\widehat{\xi}_t + u(z_t, m_t) + K(z_t, m_t, \Sigma_t)\Sigma_e(z_t, m_t, \Sigma_t)^{1/2}\epsilon_{t+1} \\
z_{t+1} &= C(z_t, m_t)A(z_t, m_t)\widehat{\xi}_t + C(z_t, m_t)u(z_t, m_t) + v(z_t, m_t) + \Sigma_e(z_t, m_t, \Sigma_t)^{1/2}\epsilon_{t+1} \\
\Sigma_{t+1} &= A(z_t, m_t)\Sigma_t A(z_t, m_t)' + B(z_t, m_t)B(z_t, m_t)' - K(z_t, m_t, \Sigma_t)\Sigma_e(z_t, m_t, \Sigma_t)K(z_t, m_t, \Sigma_t)' \\
m_{t+1} &= f(m_t, z_t)
\end{aligned} \tag{28}$$

for $t = 0, 1, \ldots$, with $\{\epsilon_t\}_{t=1}^{+\infty} \sim$ iid $N(0, \mathrm{Id})$, z_0 and $\widehat{\xi}_0$ as in (15) and (16), and where the matrices $K(z_t, m_t, \Sigma_t)$ and $\Sigma_e(z_t, m_t, \Sigma_t)$ are as defined in (18) and (19). Thus, the Markovian property in (27) follows. It is important to emphasize that Markovianity does not obtain with the information sets I_t, where, due to the presence of the latent states ξ, $P(z_t \mid I_{t-1}) \neq P(z_t \mid z_{t-1})$.

With the Markov structure in hand, we can now apply the results of Blackwell (1965) on the existence of optimal *stationary* plans $\pi = (\phi, \phi, \ldots)$ in $\overline{\Pi}$ where $\phi : \mathbb{R}^{n_z} \times \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\xi(n_\xi+1)/2} \times \mathbb{R}^m \times X \to X$ is a function that maps states $(z_t, \widehat{\xi}_t, \Sigma_t, m_t, x_t)$ into choices x_{t+1}, $t = 0, 1, \ldots$. For a stationary plan π, let $\{x_t^\phi\}_{t=1}^{+\infty}$ denote the corresponding endogenous variable process,

$$x_1^\phi = \phi(z_0, \widehat{\xi}_0, \Sigma_0, m_0, x_0), \quad x_{t+1}^\phi = \phi(z_t, \widehat{\xi}_t, \Sigma_t, m_t, x_t^\phi), \quad t = 1, 2, \ldots \tag{29}$$

with $\{\widehat{\xi}_t\}_{t=0}^{+\infty}$ and $\{\Sigma_t\}_{t=0}^{+\infty}$ as defined in equations (16)-(20), and where $\{m_t\}_{t=1}^{+\infty}$ is as in (23) with m_0 fixed. Note that the set of feasible stationary plans, which we denote by Π^∞, is a subset of $\overline{\Pi}$,

$$\Pi^\infty \subset \overline{\Pi}.$$

The following theorem shows two results: (i) if there is an optimal π^* in Π, there is one which is stationary; and conversely (ii) if a stationary plan $\pi^\infty = (\phi^*, \phi^*, \ldots)$ with $\phi^* : \mathbb{R}^{n_z} \times \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\xi(n_\xi+1)/2} \times \mathbb{R}^{n_m} \times X \to X$ and the corresponding $\big\{x_t^{\phi^*}\big\}_{t=1}^{+\infty}$ in (29) satisfies the *stationary* Euler equation and transversality conditions stated below, then π^∞ is optimal in Π.

Theorem 2. *Consider the decision problem in equation* (8) *with exogenous variables* ξ *(latent) and* z *(observed) that follow the conditionally Gauss Markov state space system in* (24)*, and let Assumptions 1-14 hold. If* $\pi^\infty \in \Pi^\infty$ *satisfies the stationary Euler equation*

$$\frac{\partial F}{\partial x_{+1}}\left(x_{t+1}^{\phi^*}, x_t^{\phi^*}, z_t\right) + \beta E\left[\frac{\partial F}{\partial x}\left(x_{t+2}^{\phi^*}, x_{t+1}^{\phi^*}, z_{t+1}\right)\middle| z_t, \widehat{\xi}_t, \Sigma_t, m_t, x_t^{\phi^*}\right] = 0 \quad a.s. \quad t = 0, 1, \ldots$$

with x^ϕ *as in* (29)*,* $z, \widehat{\xi}, \Sigma, m$ *as in* (28) *with* m_0 *fixed, and the transversality condition*

$$\lim_{T\to\infty} E\left[\beta^T \frac{\partial F}{\partial x_{+1}}\left(x_{T+1}^{\phi^*}, x_T^{\phi^*}, z_T\right) \cdot x_{T+1}^{\pi-\pi^\infty}\right] = 0, \quad \text{for any } \pi \in \Pi,$$

then π^∞ *is the unique optimal plan in* Π*. Conversely, if* $\pi^* \in \Pi$ *is an optimal plan for which Assumption 10 holds, then there a stationary plan* $\pi^\infty \in \Pi^\infty$ *which is also optimal in* Π*, and which satisfies the above stationary Euler equation and transversality condition.*

Theorem 2 formally establishes the equivalence between the original decision problem with incomplete information described in (24) in which ξ is latent and z observed, and a complete information decision model with exogenous states $(z, \widehat{\xi}, \Sigma, m)$ that evolve as in (28), all of which are observed by the decision maker. In the original problem, the decision maker is considering policies that depend on the entire history of z which makes the problem nonstationary; in its full information equivalent formulation, all the exogenous information needed for decision is summarized in the current values of z and the statistics $(\widehat{\xi}, \Sigma, m)$ that are sufficient for control, which makes the problem stationary.

Put differently, Theorem 2 allow us to turn a nonstationary decision problem, in which an optimal sequence of policies (or plan) $\pi^* = (\pi_0^*, \pi_1^*, \ldots)$ is to be found, into a stationary one, in which we only need to find one optimal policy ϕ^*. Even though the latter problem appears easier to solve, the stationary Euler equation is generally a highly nonlinear function of ϕ^* for which no closed form solutions are available. This raises the question of optimal policy approximation, to which we turn next.

4 Perturbation Approximation

The main idea behind the perturbation approach is to consider approximating the object of interest–here an optimal policy–around a point at which the solution to the original problem–here the decision problem in equation (8)–is known (or at least easier to calculate). In the vast majority of applications of the perturbation method in macroeconomics, the point around which the approximation is sought is that of a *deterministic limit* of the model, see e.g. Judd (1996). While this limit quite naturally obtains in stationary models in which the decision maker perfectly observes a Markovian state, the question of what is a deterministic limit of a nonstationary model that involves learning is less obvious. To setup the perturbation approximation, we therefore start with an asymptotic *small noise* analysis of the conditional Gauss Markov model.

Consider again the conditional Gauss Markov model in (24) with an additional scalar parameter

$\sigma > 0$ which controls the size of the error terms,

$$\begin{aligned}
\xi^\sigma_{t+1} &= A(z^\sigma_t, m^\sigma_t)\xi^\sigma_t + u(z^\sigma_t, m^\sigma_t) + \sigma B(z^\sigma_t, m^\sigma_t)w_{t+1} \\
z^\sigma_t &= C(z^\sigma_{t-1}, m^\sigma_{t-1})\xi^\sigma_t + v(z^\sigma_{t-1}, m^\sigma_{t-1}) + \sigma D(z^\sigma_{t-1}, m^\sigma_{t-1})w_t, \quad t = 0, 1, \ldots \\
m^\sigma_{t+1} &= f(m^\sigma_t, z^\sigma_t)
\end{aligned} \tag{30}$$

where $w_t \sim$ iid $N(0, \mathrm{Id})$, $\xi^\sigma_0 \sim N(\bar{\xi}^\sigma_0, \bar{\Sigma}^\sigma_0)$ and m^σ_0 given. We emphasize the fact that the solution to (30) depends on σ by labeling the resulting processes as $\{\xi^\sigma_t\}^{+\infty}_{t=0}$, $\{z^\sigma_t\}^{+\infty}_{t=0}$, and $\{m^\sigma_t\}^{+\infty}_{t=0}$. Similarly, we let $\widehat{\xi}^\sigma_t \equiv E(\xi^\sigma_t \mid z^{\sigma t})$ and $\Sigma^\sigma_t \equiv E[(\xi^\sigma_t - \widehat{\xi}^\sigma_t)(\xi^\sigma_t - \widehat{\xi}^\sigma_t)' \mid z^{\sigma t}]$, for $t = 0, 1, \ldots$ be the conditional mean and covariance of ξ^σ_t given $z^{\sigma t} \equiv (z^\sigma_0, \ldots, z^\sigma_t)$. The idea in (30) is that a common scale parameter σ controls the speed at which all the stochastic drivers in the model (30) converge to zero. We are interested in the behavior of the full-information equivalent of the conditional Gauss Markov model above as we let $\sigma \to 0$.

For this, we first examine the behavior as $\sigma \to 0$ of the innovations representation of the conditional Gauss Markov model in (30). As argued previously, the innovations representation is key in determining the full-information equivalent model and its state variables. Assuming that the initial distribution of ξ^σ_0 satisfies $\bar{\xi}^\sigma_0 = \bar{\xi}_0$ and $\bar{\Sigma}^\sigma_0 = \sigma^2\bar{\Sigma}_0$ with $(\bar{\xi}_0, \bar{\Sigma}_0)$ independent of σ (Assumption 15), the innovations representation of the conditional Gauss Markov model in (30) is given in the following lemma.

Lemma 4. *Let Assumptions 11 - 15 hold, and consider* $\widehat{\xi}^\sigma_t = E(\xi^\sigma_t \mid z^{t\sigma})$ *and* $\Sigma_t = \frac{1}{\sigma^2}E[(\xi^\sigma_t - \widehat{\xi}^\sigma_t)(\xi^\sigma_t - \widehat{\xi}^\sigma_t)' \mid z^{t\sigma}] = \frac{\Sigma^\sigma_t}{\sigma^2}$. *Then, for* $t = 0, 1, \ldots$

$$\begin{aligned}
\widehat{\xi}^\sigma_{t+1} &= A(z^\sigma_t, m^\sigma_t)\widehat{\xi}^\sigma_t + u(z^\sigma_t, m^\sigma_t) + \sigma K(z^\sigma_t, m^\sigma_t, \Sigma_t)\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)^{1/2}\epsilon_{t+1} \\
z^\sigma_{t+1} &= C(z^\sigma_t, m^\sigma_t)A(z^\sigma_t, m^\sigma_t)\widehat{\xi}^\sigma_t + C(z^\sigma_t, m^\sigma_t)u(z^\sigma_t, m^\sigma_t) + v(z^\sigma_t, m^\sigma_t) + \sigma\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)^{1/2}\epsilon_{t+1} \\
\Sigma_{t+1} &= A(z^\sigma_t, m^\sigma_t)\Sigma_t A(z^\sigma_t, m^\sigma_t)' + B(z^\sigma_t, m^\sigma_t)B(z^\sigma_t, m^\sigma_t)' \\
&\quad - K(z^\sigma_t, m^\sigma_t, \Sigma_t)\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)K(z^\sigma_t, m^\sigma_t, \Sigma_t)' \\
m^\sigma_{t+1} &= f(m^\sigma_t, z^\sigma_t)
\end{aligned} \tag{31}$$

with $\{\epsilon_t\}^{+\infty}_1 \sim$ *iid* $N(0, Id)$, *starting values* $\widehat{\xi}^\sigma_0 = \widehat{\xi}_0$ *and* Σ_0 *calculated as in* (16), *and where* $K(z^\sigma_t, m^\sigma_t, \Sigma_t)$ *and* $\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)$ *are as in* (18)-(19).

Lemma 4 has several important implications. Firstly, it shows that the full information equivalent of the conditional Gauss Markov model in (30) has state variables $(\widehat{\xi}^\sigma_t, \Sigma_t, z^\sigma_t, m^\sigma_t)$ where $\widehat{\xi}^\sigma_t = E(\xi^\sigma_t \mid z^{t\sigma})$ and $\Sigma_t = \frac{1}{\sigma^2}E[(\xi^\sigma_t - \widehat{\xi}^\sigma_t)(\xi^\sigma_t - \widehat{\xi}^\sigma_t)' \mid z^{t\sigma}] = \frac{\Sigma^\sigma_t}{\sigma^2}$. Secondly, it shows that the dynamic equations governing the law of motion of Σ_t (and m^σ_t) do not explicitly depend on σ. This is a non-trivial finding which follows from two key properties of the model in (30): (i) the Kalman gain matrix $K(z^\sigma_t, m^\sigma_t, \Sigma_t)$ does not depend on σ directly; and (ii) the covariance matrix of the innovations is of the form $\sigma^2\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)$, where again $\Sigma_e(z^\sigma_t, m^\sigma_t, \Sigma_t)$ only depends on σ indirectly, through its dependence on (z^σ_t, m^σ_t). Thirdly, the dynamic equations governing the laws of motion of $\widehat{\xi}^\sigma_t$ and z^σ_t, depend on σ in a very specific way, which is well suited for examining the behavior of the sequences $\{\widehat{\xi}^\sigma_t\}^{+\infty}_{t=0}$ and $\{z^\sigma_t\}^{+\infty}_{t=0}$ as $\sigma \to 0$ and applying the perturbation method of Fleming (1971).

A Assumptions

This Appendix collects the assumptions on the stochastic decision model defined in Section 3. To clarify their role, we compare and contrast our assumptions with those by Stokey and Lucas (1989), whenever possible. *In the assumptions that follow $x_0 \in X$ as well as all the parameters of the model, β, Γ, F as well as $A, C, u, W, V, \bar{\xi}_0, \bar{\Sigma}_0$ later on, are all given, deterministic, and assumed known to the decision maker.*

We first state the assumptions that ensure that the function $v(\pi)$ in (10) and the optimization problem in (11) are well defined.

Assumption 1. *The set Π is nonempty.*

Assumption 2. *For each $\pi \in \Pi$, the limit*

$$\lim_{T\to\infty} \sum_{t=0}^{T} \beta^t \int_{\mathcal{Y}^t} F(x_{t+1}^\pi, x_t^\pi, z_t)\mu^t(dz^t) \quad \text{exists (it may be } -\infty).$$

Assumption 3. *There is a sequence of $\{m_t\}_{t=0}^{+\infty}$ of nonnegative real numbers such that:*

$$(i) \quad \int_{\mathcal{Z}^t} F(x_{t+1}^\pi, x_t^\pi, z_t)\mu^t(dz^t) \leqslant m_t \quad \text{for every } t = 0, 1, \ldots \text{ and } \pi \in \Pi$$

$$(ii) \quad \sum_{t=0}^{+\infty} \beta^t m_t < +\infty$$

Assumption 4. *There exists $\pi \in \Pi$ such that*

$$E\left[\sum_{t=0}^{+\infty} \beta^t F(x_{t+1}^\pi, x_t^\pi, z)\right] > -\infty$$

Assumption 1 plays the same role as Assumption 9.1 in Stokey and Lucas (1989); see their Lemma 9.1. Assumption 2 is analogous to Assumption 9.2(b) in Stokey and Lucas (1989). Note that this assumption does not preclude F from being unbounded nor from taking both signs. Assumption 3 plays the same role as Assumption 9.3 in Stokey and Lucas (1989): it is used to justify the change in the order of limit and integration (see p. 250 in Stokey and Lucas (1989)). In particular, CRRA utility functions are allowed (including the log utility) so long as Assumptions 2 and 3 hold. Finally, Assumption 4 ensures that $v(\pi)$ in (10) does not take value $-\infty$ for all $\pi \in \Pi$ so that the supremum in (11) is well defined.

Next is the group of assumptions that we impose in order to guarantee that the function $v(\pi)$ has a Gateaux differential.

Assumption 5. *For each possible value $z \in Z$, the function $F(\cdot, \cdot, z)$ is differentiable in the interior of $X \times X$.*

Assumption 6. *Let $\pi \in \Pi$ be an internal plan in the direction $\lambda \in \Lambda$, denote by $\{x_t^\pi\}_{t=1}^{+\infty}$ and $\{x_t^\lambda\}_{t=1}^{+\infty}$ the corresponding endogenous state processes, and let*

$$F_0(\epsilon) = F(x_1^\pi + \epsilon x_1^\lambda, x_0, z_0) \text{ and } F_t(\epsilon) = F(x_{t+1}^\pi + \epsilon x_{t+1}^\lambda, x_t^\pi + \epsilon x_t^\lambda, z_t), \quad t = 1, 2, \ldots.$$

There exists $\bar{\epsilon} > 0$ *such that for all* $|\epsilon| < \bar{\epsilon}$

$$\frac{d}{d\epsilon} E\left[\sum_{t=0}^{+\infty} \beta^t F_t(\epsilon)\right] = E\left[\sum_{t=0}^{+\infty} \beta^t \frac{dF_t}{d\epsilon}(\epsilon)\right].$$

Assumption 5 is analogous to Assumption 9.12 in Stokey and Lucas (1989). Assumption 6 is used to ensure that one can differentiate under the integral in (10).

Last are the convexity and concavity assumptions.

Assumption 7. *The set* X *is a convex subset of* $\mathbb{R}^{n_x}$.

Assumption 8. *For every* $z \in Z$, $(x_1, y_1, x_2, y_2)X \times X \times X \times X$ *and* $0 \leqslant \theta \leqslant 1$ *imply*

$$F(\theta y_1 + (1-\theta)y_2, \theta x_1 + (1-\theta)x_2, z) \geqslant \theta F(y_1, x_1, z) + (1-\theta)F(y_2, x_2, z),$$

and the inequality is strict when $(y_1, x_1) \neq (y_2, x_2)$ *and* $0 < \theta < 1$.

Assumption 9. *For every* $(z, x_1, x_2) \in Z \times X \times X$, $y_1 \in \Gamma(x_1, z)$, $y_2 \in \Gamma(x_2, z)$ *and* $0 \leqslant \theta \leqslant 1$ *imply* $\theta y_1 + (1-\theta)y_2 \in \Gamma(\theta x_1 + (1-\theta)x_2, z)$.

Assumption 7 is the same as Assumption 9.4 in Stokey and Lucas (1989). Assumptions 8 and 9, which are the same as Assumptions 9.10 and 9.11 in Stokey and Lucas (1989), ensure that the set of feasible policies Π is convex and that $v(\cdot)$ is strictly concave on Π.

Assumption 10. *Let* $\pi^* \in \Pi$ *be an optimal plan in* (11). *For any* $t = 0, 1, \ldots$: *(i)* $x_{t+1}^{\pi^*}$ *is an interior point of the feasible set* $\Gamma(x_t^{\pi^*}, z_t)$, *and (ii) there exists* $\epsilon_t > 0$ *such that* $\|x - x_{t+1}^{\pi^*}\| < \epsilon_t$ *implies* $x_{t+2}^{\pi^*} \in \Gamma(x, z_{t+1})$, *for all* $z^{t+1} \in \mathcal{Z}^{t+1}$.

Assumption 10 is used to ensure that an optimal plan π^* is interior in any direction $\lambda_\tau = (0, \ldots, 0, \nu_\tau, 0, \ldots)$ with $\tau = 0, 1, \ldots$, and ν_τ that is $\mathcal{B}_{\mathcal{Z}^t}$-measurable.

Assumption 11. $\{w_t\}_{t=1}^{+\infty}$ *is independent and normally distributed* $N(0, Id)$; ξ_0 *is normally distributed* $N(\bar{\xi}_0, \bar{\Sigma}_0)$ *with* $\bar{\Sigma}_0$ *positive semidefinite;* $\{w_t\}_{t=1}^{+\infty}$ *and* ξ_0 *are independent.*

Assumption 12. *For* $t = 0, 1, \ldots$, *the entries of* $A(z^t), B(z^t), C(z^t), D(z^t), u(z^t), v(z^t)$ *are: (i)* $\mathcal{B}_{\mathcal{Z}}^t$*-measurable; (ii) finite with probability one.*

Assumption 13. *For* $t = 0, 1, \ldots$, *the matrices* $\Sigma_e(z^t, \Sigma_t)$ *are nonsingular with probability one.*

Assumption 14. *There exists a random variable* $m \in \mathbb{R}^{n_m}$ *such that for* $t = 1, 2, \ldots$, *(i)* m_t *is* $\mathcal{B}_{\mathcal{Z}}^{t-1}$ *measurable, (ii)* m_t *is recursive* $m_{t+1} = f(m_t, z_t)$, *for some function* $f : \mathbb{R}^{n_m} \times \mathbb{R}^{n_z} \to \mathbb{R}^{n_m}$, *and (iii) such that* A, B, C, D, u, v *depend on* z^t *only through* (z_t, m_t), *i.e.* $A(z^t) = A(z_t, m_t)$ *with analogous restrictions on the other system matrices.*

Assumption 15. *The initial distribution of* ξ_0^σ *in* (30) *satisfies* $\bar{\xi}_0^\sigma = \bar{\xi}_0$ *and* $\bar{\Sigma}_0^\sigma = \sigma^2 \bar{\Sigma}_0$ *with* $(\bar{\xi}_0, \bar{\Sigma}_0)$ *independent of* σ.

B Proofs

This Appendix contains the proofs of the results stated in the main text.

Proof of Lemma 1. Under Assumptions 1-4 the function $v(\cdot)$ in (10) is well defined. Let π and λ be as in Assumption 6. For any $|\epsilon| < \min\{\epsilon_\lambda, \bar{\epsilon}\}$ we have

$$\frac{dv}{d\epsilon}(\pi + \epsilon\lambda) = \frac{d}{d\epsilon} E\left[F(x_1^\pi + \epsilon x_1^\lambda, x_0, z_0) + \sum_{t=1}^{+\infty} \beta^t F(x_{t+1}^\pi + \epsilon x_{t+1}^\lambda, x_t^\pi + \epsilon x_t^\lambda, z_t)\right]$$
$$= E\left[\frac{d}{d\epsilon} F(x_1^\pi + \epsilon x_1^\lambda, x_0, z_0) + \sum_{t=1}^{+\infty} \beta^t \frac{d}{d\epsilon} F(x_{t+1}^\pi + \epsilon x_{t+1}^\lambda, x_t^\pi + \epsilon x_t^\lambda, z_t)\right]$$
$$= E\left[\frac{\partial F}{\partial x_{+1}}(x_1^\pi + \epsilon x_1^\lambda, x_0, z_0) \cdot x_1^\lambda + \sum_{t=1}^{+\infty} \beta^t \left(\frac{\partial F}{\partial x_{+1}}(x_{t+1}^\pi + \epsilon x_{t+1}^\lambda, x_t^\pi + \epsilon x_t^\lambda, z_t) \cdot x_{t+1}^\lambda + \frac{\partial F}{\partial x}(x_{t+1}^\pi + \epsilon x_{t+1}^\lambda, x_t^\pi + \epsilon x_t^\lambda, z_t) \cdot x_t^\lambda\right)\right]$$

where the second equality uses Assumption 6, and where $\partial/\partial x_{+1}$ (resp. $\partial/\partial x$) denotes the gradient with respect to the first variable x_{+1} (resp. second variable x) in $F(x_{+1}, x, z)$, and $\cdot$ denotes the scalar product. Taking $\epsilon = 0$ then yields

$$\delta v(\pi, \lambda) = E\left[\frac{\partial F}{\partial x_{+1}}(x_1^\pi, x_0, z_0) \cdot x_1^\lambda + \sum_{t=1}^{+\infty} \beta^t \left(\frac{\partial F}{\partial x_{+1}}(x_{t+1}^\pi, x_t^\pi, z_t) \cdot x_{t+1}^\lambda + \frac{\partial F}{\partial x}(x_{t+1}^\pi, x_t^\pi, z_t) \cdot x_t^\lambda\right)\right].$$

□

Proof of Theorem 1. The proof proceeds in three steps: first, we establish that the feasible set Π is convex and that $v(\cdot)$ is strictly concave on Π. Second, we show that the Euler equation is necessary for optimality. Third, we show that it is also sufficient. We note that the convexity of Π and concavity of $v(\cdot)$ only play a role in the last step, i.e. the Euler equation remains necessary for optimality even if the convexity requirements are not met.

Step 1: Concavity of $v(\cdot)$. We first establish that under Assumption 8, the set Π is convex. For this, consider two policies π_1 and π_2 in Π, and take any $0 \leqslant \theta \leqslant 1$. The policy $\pi_\theta = \theta\pi_1 + (1-\theta)\pi_2$ generates the following sequence $\{x_t^{\pi_\theta}\}_{t=1}^{+\infty}$:

$$x_1^{\pi_\theta} = \theta\pi_0^1(z_0) + (1-\theta)\pi_0^2(z_0) = \theta x_1^{\pi_1} + (1-\theta)x_1^{\pi_2} \in \Gamma(x_0, z_0)$$
$$x_{t+1}^{\pi_\theta} = \theta\pi_t^1(z^t) + (1-\theta)\pi_t^2(z^t) \in \Gamma\left(\theta\pi_{t-1}^1(z^{t-1} + (1-\theta)\pi_{t-1}^2(z^{t-1}), z_t)\right)$$

where the inclusions follow by Assumption 7 (convexity of X) and Assumption 9. Next, we establish that $v(\pi)$ in (10) is strictly concave on Π. For this, consider $\pi_1, \pi_2 \in \Pi$ and $0 \leqslant \theta \leqslant 1$ and let

$\pi_\theta = \theta\pi_1 + (1-\theta)\pi_2$; then

$$\begin{aligned} v\left(\theta\pi_1 + (1-\theta)\pi_2\right) &= E\left[\sum_{t=0}^{+\infty} \beta^t F(x_{t+1}^{\pi_\theta}, x_t^{\pi_\theta}, z_t)\right] \\ &= E\left[\sum_{t=0}^{+\infty} \beta^t F(\theta\pi_{1t}(z^t) + (1-\theta)\pi_{2t}(z^t), \theta\pi_{1,t-1}(z^{t-1}) + (1-\theta)\pi_{2,t-1}(z^{t-1}), z_t)\right] \\ &\geqslant E\left[\sum_{t=0}^{+\infty} \beta^t \left(\theta F(\pi_{1t}(z^t), \pi_{1,t-1}(z^{t-1}), z_t) + (1-\theta)F(\pi_{2t}(z^t), \pi_{2,t-1}(z^{t-1}), z_t)\right)\right] \\ &= \theta v(\pi_1) + (1-\theta)v(\pi_2) \end{aligned}$$

where the inequality uses Assumption 8, and is strict when $\pi_1 \neq \pi_2$ and $0 < \theta < 1$.

Step 2: Necessity. Let $\pi^* \in \Pi$ be an optimal plan in (11). Pick an integer $\tau \geqslant 0$ and pick a measurable function $\nu_\tau : \mathcal{Z}^\tau \to \mathbb{R}^{n_x}$. Define a plan $\lambda^\tau \in \Lambda$ by letting $\lambda_t^\tau = 0$ if $t \neq \tau$ and $\lambda_\tau^\tau = \nu_\tau$. By Assumption 10, there exists $\epsilon_\tau > 0$ such that for every $|\epsilon| < \epsilon_\tau$ we have

$$(i) \quad x_{\tau+1}^{\pi^*} + \epsilon\nu_\tau \in \Gamma(x_\tau^{\pi^*}, z_\tau), \quad \text{and} \quad (ii) \quad x_{\tau+2}^{\pi^*} \in \Gamma(x_{\tau+1}^{\pi^*} + \epsilon\nu_\tau, z_{\tau+1})$$

for all $z^{\tau+1} \in \mathcal{Z}^{\tau+1}$, $\tau = 0, 1, \ldots$. That is, the plan π^* is internal in the direction λ^τ. Moreover, π^* and λ^τ satisfy Assumption 6. Then, by Lemma 1, the Gateaux differential $\delta v(\pi^*, \lambda^\tau)$ exists and equals

$$\delta v(\pi^*, \lambda^\tau) = E\left[\left(\frac{\partial F}{\partial x_{+1}}(x_{\tau+1}^{\pi^*}, x_\tau^{\pi^*}, z_\tau) + \beta\frac{\partial F}{\partial x}(x_{\tau+2}^{\pi^*}, x_{\tau+1}^{\pi^*}, z_{\tau+1})\right) \cdot \nu_\tau\right] = 0, \quad \tau = 0, 1, \ldots \tag{32}$$

where the second equality follows by Theorem 2.1 in Fleming and Rishel (1975). Since (32) holds for any $\nu_\tau : \mathcal{Z}^\tau \to \mathbb{R}^{n_x}$ measurable, it follows by the definition of conditional expectations (see, e.g. Section 34 in Billingsley (1995)) that

$$E\left[\frac{\partial F}{\partial x_{+1}}(x_{\tau+1}^{\pi^*}, x_\tau^{\pi^*}, z_\tau) + \beta\frac{\partial F}{\partial x}(x_{\tau+2}^{\pi^*}, x_{\tau+1}^{\pi^*}, z_{\tau+1})\middle| z^\tau\right] = 0 \quad a.s. \quad \tau = 0, 1, \ldots.$$

Since the first term inside the expectation is $\mathcal{B}_\mathcal{Z}^\tau$ measurable, the above equality can also be written as

$$\frac{\partial F}{\partial x_{+1}}(x_{\tau+1}^{\pi^*}, x_\tau^{\pi^*}, z_\tau) + \beta E\left[\frac{\partial F}{\partial x}(x_{\tau+2}^{\pi^*}, x_{\tau+1}^{\pi^*}, z_{\tau+1})\middle| z^\tau\right] = 0 \quad a.s. \quad \tau = 0, 1, \ldots. \tag{33}$$

Now take any other plan $\pi \in \Pi$ and consider the direction $\varphi = \pi - \pi^* \in \Lambda$. Since the set Π is convex, we have for every $0 \leqslant \varepsilon \leqslant 1$, $\pi^* + \varepsilon\varphi \in \Pi$, that is, π^* is an internal point of Π in the direction φ. Again, using Theorem 2.1 in Fleming and Rishel (1975), we necessarily have $\delta v(\pi^*, \varphi) = 0$. Combining the latter with (12) and (33) yields

$$E\left[\beta^T \frac{\partial F}{\partial x_{+1}}(x_{T+1}^{\pi^*}, x_T^{\pi^*}, z_T) \cdot x_{T+1}^{\varphi} + \sum_{t=T+1}^{+\infty} \beta^t \left(\frac{\partial F}{\partial x_{+1}}(x_{t+1}^{\pi}, x_t^{\pi}, z_t) \cdot x_{t+1}^{\varphi} + \frac{\partial F}{\partial x}(x_{t+1}^{\pi}, x_t^{\pi}, z_t) \cdot x_t^{\varphi}\right)\right] = 0, \quad \forall T = 0, 1, \ldots$$

Taking the limit as $T \to \infty$ then gives

$$\lim_{T\to\infty} E\left[\beta^T \frac{\partial F}{\partial x_{+1}}\left(x_{T+1}^{\pi^*}, x_T^{\pi^*}, z_T\right) \cdot x_{T+1}^{\varphi}\right] = 0. \tag{34}$$

STEP 3: SUFFICIENCY. Take $\pi, \pi^* \in \Pi$ ($\pi \neq \pi^*$) and $0 < \theta < 1$; it follows from strict concavity of $v(\cdot)$ (established in Step 1) that

$$v\left(\theta\pi + (1-\theta)\pi^*\right) > \theta v(\pi) + (1-\theta)v(\pi^*),$$

which can be re-arranged as

$$\frac{1}{\theta}\left[v\left(\pi^* + \theta(\pi - \pi^*)\right) - v(\pi^*)\right] > v(\pi) - v(\pi^*).$$

Taking the limit as $\theta \to 0$ then gives

$$\delta v(\pi^*, \pi - \pi^*) > v(\pi) - v(\pi^*)$$

Thus, if $\delta v(\pi^*, \pi - \pi^*) = 0$ for all $\pi \in \Pi$, then $v(\cdot)$ has a unique maximum at π^*. Let then π^* be a policy that satisfies the Euler equation (33), and consider any $\pi \in \Pi$, $\pi \neq \pi^*$. Define the direction $\varphi \equiv \pi - \pi^*$. Since for any $t = 0, 1, \ldots$, $\varphi_t = \pi_t^* - \pi_t$ is $\mathcal{B}_{Z^t}$-measurable, we have from (33) and (34) that

$$\begin{aligned}
\delta v(\pi^*, \varphi) &= E\left[\tfrac{\partial F}{\partial x_{+1}}(x_1^{\pi}, x_0, z_0) \cdot x_1^{\lambda} + \sum_{t=1}^{+\infty} \beta^t \left(\tfrac{\partial F}{\partial x_{+1}}(x_{t+1}^{\pi}, x_t^{\pi}, z_t) \cdot x_{t+1}^{\varphi} + \tfrac{\partial F}{\partial x}(x_{t+1}^{\pi}, x_t^{\pi}, z_t) \cdot x_t^{\varphi}\right)\right] \\
&= \lim_{T\to\infty} E\left[\sum_{t=0}^{T-1} \beta^t \left(\tfrac{\partial F}{\partial x_{+1}}(x_{t+1}^{\pi^*}, x_t^{\pi^*}, z_t) + \beta \tfrac{\partial F}{\partial x}(x_{t+2}^{\pi^*}, x_{t+1}^{\pi^*}, z_{t+1})\right) \cdot x_{t+1}^{\varphi} + \beta^T \tfrac{\partial F}{\partial x_{+1}}(x_{T+1}^{\pi^*}, x_T^{\pi^*}, z_T) \cdot x_{T+1}^{\varphi}\right] \\
&= \lim_{T\to\infty} E\left[\beta^T \frac{\partial F}{\partial x_{+1}}\left(x_{T+1}^{\pi^*}, x_T^{\pi^*}, z_T\right) \cdot x_{T+1}^{\varphi}\right] \\
&= 0
\end{aligned}$$

Thus, $v(\cdot)$ has a strict maximum on Π at π^*.

□

Proof of Lemma 2. To prove the lemma, we use the results from Ch 13 of Liptser and Shiryayev (1978). Note that by combining the transition and observation equations (13)-(14), we can write

$$z_{t+1} = \underbrace{C(z^t)A(z^t)}_{A_1(t,z)}\xi_t + C(z^t)u(z^t) + v(z^t) + \underbrace{\left[C(z^t)B(z^t) + D(z^t)\right]}_{B_1(t,z)} w_{t+1},$$

which is the form of the observation equation (13.47) in Liptser and Shiryayev (1978); our transition equation (13) is of the form (13.46) used in Liptser and Shiryayev (1978). We now follow the reasoning in Theorem 13.5 in (Liptser and Shiryayev, 1978). In order to establish the conditional Gaussianity of ξ_t given z^t, we do not use their Theorem 13.3 which requires the stringent hypotheses of finite second

moments of the entries of the random system matrices B, D, u, v and boundedness of the entries of A, C (those are conditions (1) and (2) on p. 62 of Liptser and Shiryayev (1978)) – rather, we use our Assumption 12 that those entries are finite with probability one.[16] Then, by the main Theorem of Chen, Kumar, and van Schuppen (1989), the conditional Gaussianity follows. The reasoning in Theorem 13.5 in Liptser and Shiryayev (1978) applies provided their condition

$$B_1(t,z)B_1(t,z)' + A_1(t,z)\Sigma_t A_1(t,z)' \quad \text{is nonsingular with probability one, } t = 0, 1, \ldots$$

holds. Our Assumption 13 ensures that for $t = 0, 1, \ldots$, the matrix $B_1(t,z)B_1(t,z)'$ is positive definite with probability one, thus the same holds for $B_1(t,z)B_1(t,z)' + A_1(t,z)\Sigma_t A_1(t,z)'$ and the above condition is satisfied. □

Proof of Lemma 3. To prove the lemma, we follow the same steps as Rhenius (1974) in his proof of Theorem 6. First, we need to establish that under our Assumptions 1-4 for any $\overline{\pi} \in \overline{\Pi}$, $v(\overline{\pi})$ in (10) and $v(\overline{\pi}^*)$ in (11) are well defined. For this, note that if Π is nonempty (Assumption 1), then $\overline{\Pi}$ is nonempty. Next, note that for any plan $\overline{\pi} \in \overline{\Pi}$, the policy $\overline{\pi}_t$ is $\mathcal{B}_{\mathcal{Z}^t}$-measurable so we can find $\pi_t : \mathcal{Z}^t \to X$ such that $\overline{\pi}_t(\overline{I}_t) = \pi_t(z^t)$; thus, if $\pi = (\pi_0, \pi_1, \ldots) \in \Pi$ satisfies Assumptions 2-4, so does $\overline{\pi}$. Thus $v(\overline{\pi})$ is well defined and moreover $v(\overline{\pi}) = v(\pi)$. □

Proof of Theorem 2. The proof of Theorem 2 consists of two parts: part 1 establishes the sufficiency of the stationary Euler equation and transversality condition, while part 2 shows necessity. Throughout, we let Assumptions 1-11 hold, so in particular Theorem 1 and Lemma 3 apply.

PART 1: SUFFICIENCY. Consider a stationary plan $\pi^\infty = (\phi^*, \phi^*, \ldots)$ in Π^∞ with $\phi^* : \mathbb{R}^{n_z} \times \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\xi(n_\xi+1)/2} \times X \to X$ and the corresponding endogenous variable process $\{x_t^{\phi^*}\}_{t=1}^{+\infty}$ that satisfy the stationary Euler equation

$$\frac{\partial F}{\partial x_{+1}}\left(x_{t+1}^{\phi^*}, x_t^{\phi^*}, z_t\right) + \beta E\left[\frac{\partial F}{\partial x}\left(x_{t+2}^{\phi^*}, x_{t+1}^{\phi^*}, z_{t+1}\right)\middle| z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right] = 0 \quad a.s. \quad t = 0, 1, \ldots$$

with $\{\widehat{\xi}_t\}_{t=0}^{+\infty}$ and $\{\Sigma_t\}_{t=0}^{+\infty}$ as given in (16)-(20). Now, the expectation term above can be written as:

$$\begin{aligned}
&E\left[\frac{\partial F}{\partial x}\left(x_{t+2}^{\phi^*}, x_{t+1}^{\phi^*}, z_{t+1}\right)\middle| z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right] \\
&= E\left[\frac{\partial F}{\partial x}\left(\phi^*\left(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right)\right), \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right), z_{t+1}\right) \;\middle|\; z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right] \\
&= E\left[\frac{\partial F}{\partial x}\left(\phi^*\left(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right)\right), \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right), z_{t+1}\right) \;\middle|\; \overline{I}_t\right] \\
&= E\left[\frac{\partial F}{\partial x}\left(\phi^*\left(z_{t+1}, \widehat{\xi}_{t+1}, \Sigma_{t+1}, \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right)\right), \phi^*\left(z_t, \widehat{\xi}_t, \Sigma_t, x_t^{\phi^*}\right), z_{t+1}\right) \;\middle|\; z^t\right] \qquad (35)
\end{aligned}$$

where the second equality uses the Markov property (27), and the third equality uses the fact that $\overline{I}_t$ is $\mathcal{B}_{\mathcal{Z}}^t$-measurable; thus π^∞ satisfies the Euler equation in Theorem 1. If in addition π^∞ satisfies the

[16] Note that conditions (3) and (4) on p. 62 of Liptser and Shiryayev (1978) are satisfied here.

transversality condition

$$\lim_{T\to\infty} E\left[\beta^T \frac{\partial F}{\partial x_{+1}}\left(x_{T+1}^{\phi^*}, x_T^{\phi^*}, z_T\right)\cdot x_{T+1}^{\pi-\pi^\infty}\right] = 0, \quad \text{for any } \pi\in\Pi,$$

then by Theorem 1 (sufficiency), π^∞ is the unique optimal plan in Π.

PART 2: NECESSITY. Let $\pi^*\in\Pi$ be an optimal plan for which Assumption 10 holds. By Lemma 3, π^* is also optimal in $\overline{\Pi}$. Now, we can put our setup in the framework considered by Blackwell (1965) by letting $s_t \equiv (x_t, z_t, \widehat{\xi}_t, \Sigma_t)$ be the state at time $t = 0, 1, \ldots$ taking values in $S = X\times\mathbb{R}^{n_z}\times\mathbb{R}^{n_x i}\times\mathbb{R}^{n_\xi(n_\xi+1)/2}$, $a_t = x_{t+1}$ be the action at time t taking values in $A = X$, and letting $r(s_t, a_t, s_{t+1}) = F(x_t, z_t, x_{t+1})$ be the reward at time t. By Theorem 6(c) in Blackwell (1965), there exists an optimal plan π^∞ which is stationary. Since π^* satisfies Assumption 10, so does π^∞, so by Theorem 1 π^∞ satisfies the Euler equation and transversality condition. Using the equality in (35) then shows that π^∞ satisfies the stationary Euler equation given in Theorem 2.

□

Proof of Lemma 4. Using exactly the same reasoning as in Sections 3.3 and 3.4 but with the system matrices (B, D) being replaced with $(\sigma B, \sigma D)$, and assuming that the starting values $(\bar{\xi}_0^\sigma, \bar{\Sigma}_0^\sigma)$ are such that $\bar{\xi}_0^\sigma = \bar{\xi}_0$ and $\bar{\Sigma}_0^\sigma = \sigma^2\bar{\Sigma}_0$, gives the initial values

$$\widehat{\xi}_0^\sigma = \widehat{\xi}_0 \quad \text{and} \quad \Sigma_0^\sigma = \sigma^2\Sigma_0, \tag{36}$$

where $(\widehat{\xi}_0, \Sigma_0)$ are as given in equation (16). Note that neither ξ_0 nor Σ_0 depend on σ. It then follows from (19) and (36) that

$$\begin{aligned}\Sigma_e^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma) \equiv\ & C(z_0^\sigma, m_0^\sigma)A(z_0^\sigma, m_0^\sigma)\Sigma_0^\sigma A(z_0^\sigma, m_0^\sigma)'C(z_0^\sigma, m_0^\sigma)' \\ & + [C(z_0^\sigma, m_0^\sigma)\sigma B(z_0^\sigma, m_0^\sigma) + \sigma D(z_0^\sigma, m_0^\sigma)][C(z_0^\sigma, m_0^\sigma)\sigma B(z_0^\sigma, m_0^\sigma) + \sigma D(z_0^\sigma, m_0^\sigma)]' \\ =\ & \sigma^2\Sigma_e(z_0^\sigma, m_0^\sigma, \Sigma_0),\end{aligned} \tag{37}$$

with $\Sigma_e(z_0^\sigma, m_0^\sigma, \Sigma_0)$ being calculated exactly as in (19). Again, note that none of the matrices in (19) depend on σ directly, i.e. $\Sigma_e(z_0^\sigma, m_0^\sigma, \Sigma_0)$ depends on σ only indirectly through z_0^σ and m_0^σ. Moving on to (18), we have

$$\begin{aligned}K^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma) \equiv\ & \big[A(z_0^\sigma, m_0^\sigma)\Sigma_0^\sigma A(z_0^\sigma, m_0^\sigma)'C(z_0^\sigma, m_0^\sigma)' + \sigma^2 B(z_0^\sigma, m_0^\sigma)B(z_0^\sigma, m_0^\sigma)'C(z_0^\sigma, m_0^\sigma)' \\ & + \sigma^2 D(z_0^\sigma, m_0^\sigma)D(z_0^\sigma, m_0^\sigma)'\big]\Sigma_e^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma)^+ \\ =\ & K(z_0^\sigma, m_0^\sigma, \Sigma_0)\end{aligned} \tag{38}$$

where the second equality uses (36) and (37), and where $K(z_0^\sigma, m_0^\sigma, \Sigma_0)$ is being calculated exactly as in (18), i.e. its only dependence on σ is indirectly through z_0^σ and m_0^σ. Finally, to calculate Σ_1^σ, we have

from (20), (36), (37) and (38)

$$\begin{aligned}\Sigma_1^\sigma &\equiv A(z_0^\sigma, m_0^\sigma)\Sigma_0^\sigma A(z_0^\sigma, m_0^\sigma)' + \sigma^2 B(z_0^\sigma, m_0^\sigma)B(z_0^\sigma, m_0^\sigma)' - K^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma)\Sigma_e^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma)K^\sigma(z_0^\sigma, m_0^\sigma, \Sigma_0^\sigma)' \\ &= \sigma^2\big[A(z_0^\sigma, m_0^\sigma)\Sigma_0 A(z_0^\sigma, m_0^\sigma)' + B(z_0^\sigma, m_0^\sigma)B(z_0^\sigma, m_0^\sigma)' - K(z_0^\sigma, m_0^\sigma, \Sigma_0)\Sigma_e(z_0^\sigma, m_0^\sigma, \Sigma_0)K(z_0^\sigma, m_0^\sigma, \Sigma_0)'\big] \\ &= \sigma^2\Sigma_1 \qquad (39)\end{aligned}$$

i.e. Σ_1 is calculated exactly as in (20) and depends on σ only indirectly through z_0^σ and m_0^σ.

Note that Σ_0 does not depend on σ, while Σ_1 depends on σ only indirectly through z_0^σ and m_0^σ. Iterating through equations (37)-(39) for $t = 1, \ldots$ shows that for all $t = 0, 1, \ldots$

$$\Sigma_e^\sigma(z_t^\sigma, m_t^\sigma, \Sigma_t^\sigma) = \sigma^2\Sigma_e(z_t^\sigma, m_t^\sigma, \Sigma_t), \quad K^\sigma(z_t^\sigma, m_t^\sigma, \Sigma_t^\sigma) = K(z_t^\sigma, m_t^\sigma, \Sigma_t), \quad \Sigma_t^\sigma = \sigma^2\Sigma_t \qquad (40)$$

with $\Sigma_e(z_t^\sigma, m_t^\sigma, \Sigma_t)$, $K(z_t^\sigma, m_t^\sigma, \Sigma_t)$ and Σ_t calculated exactly as in (18)-(20), i.e. through recursions that depend on σ only indirectly through z_t^σ and m_t^σ. Combining (40) with (28) then yields the result of the lemma. □

C Matrix Cholesky Square Roots

Let Σ be an $n \times n$ positive semidefinite matrix. Then, there exists at least one, and in general many, $n \times n$ matrices R such that $RR' = \Sigma$. When Σ is positive definite, then a *unique* Cholesky lower triangular square root $\Sigma^{1/2}$ can be defined such that

$$\Sigma^{1/2}(\Sigma^{1/2})' = \Sigma,$$

and the diagonal entries of $\Sigma^{1/2}$ are all positive. The elements $\Sigma_{ij}^{1/2}$ of the Cholesky square root matrix can be computed sequentially, row by row, from the following system:

$$\begin{pmatrix} \Sigma_{11}^{1/2} & 0 & \ldots & 0 \\ \Sigma_{21}^{1/2} & \Sigma_{22}^{1/2} & & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{n1}^{1/2} & \Sigma_{n2}^{1/2} & \ldots & \Sigma_{nn}^{1/2} \end{pmatrix} \begin{pmatrix} \Sigma_{11}^{1/2} & \Sigma_{21}^{1/2} & \ldots & \Sigma_{n1}^{1/2} \\ 0 & \Sigma_{22}^{1/2} & & \Sigma_{n2}^{1/2} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \Sigma_{nn}^{1/2} \end{pmatrix} = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} & \ldots & \Sigma_{1n} \\ \Sigma_{12} & \Sigma_{22} & & \Sigma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{1n} & \Sigma_{2n} & \ldots & \Sigma_{nn} \end{pmatrix}$$

which yields the recursion: for $i = 1, \ldots, n$, compute

$$\Sigma_{ij}^{1/2} = \begin{cases} \frac{\Sigma_{ij} - \sum_{k=1}^{j-1}\Sigma_{ik}^{1/2}\Sigma_{jk}^{1/2}}{\Sigma_{jj}^{1/2}} & , \quad j = 1, 2, \ldots, i-1 \\ \sqrt{\Sigma_{ii} - \sum_{k=1}^{i-1}(\Sigma_{ik}^{1/2})^2} & , \quad j = i \\ 0 & , \quad j > i \end{cases}$$

D Details of the motivating example

In this section, we give the detailed derivations for the motivating example in Section 2.

Step 1: conditional Gauss Markov framework Let

$$w_t = \begin{pmatrix} \epsilon_{at} \\ \epsilon_{\theta t} \end{pmatrix}, \quad w_t \sim N(0, \mathrm{Id}_2)$$

be all the disturbances in the system (notice the stacking of technology shocks on top and parameter shocks at bottom). Then, the state transition equation is

$$\underbrace{\begin{pmatrix} a_{t+1} \\ \theta_{t+1} \end{pmatrix}}_{\xi_{t+1}} = \underbrace{\begin{pmatrix} 0 & a_t \\ 0 & 1 \end{pmatrix}}_{A(z_t, m_t)} \begin{pmatrix} a_t \\ \theta_t \end{pmatrix} + \underbrace{\begin{pmatrix} \exp[\frac{1}{2}(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t)] & 0 \\ 0 & \sigma_\theta \end{pmatrix}}_{B(z_t, m_t)} w_{t+1} \tag{41}$$

and the observation equation is

$$\underbrace{a_t}_{z_t} = \underbrace{\begin{pmatrix} 1 & 0 \end{pmatrix}}_{C(z_{t-1}, m_{t-1})} \begin{pmatrix} a_t \\ \theta_t \end{pmatrix} + \underbrace{\begin{pmatrix} 0 & 0 \end{pmatrix}}_{D(z_{t-1}, m_{t-1})} w_t \tag{42}$$

where we have defined

$$m_t = \sigma_{at}^2, \quad \text{with} \quad m_0 = \sigma_{a0}^2 \text{ given}$$

Note that for all $t = 0, 1, \ldots$, we have $C(z_t, m_t) = C$ and $D(z_t, m_t) = D$.

Step 2: Kalman Filtering equations The conditional distribution of ξ_t given z^t is Gaussian with mean $\widehat{\xi}_t$ and covariance Σ_t. Let us start with prior mean and covariance,

$$\bar{\xi}_0 = \begin{pmatrix} \bar{a} \\ \bar{\theta} \end{pmatrix} \quad \text{and} \quad \bar{\Sigma}_0 = \begin{pmatrix} \bar{\Sigma}_a & 0 \\ 0 & \bar{P} \end{pmatrix}$$

At time $t = 0$, after observing $z_0 = a_0$, we get from (42),

$$\begin{aligned} \widehat{\xi}_0 &= \bar{\xi}_0 + \bar{\Sigma}_0 C'(C\bar{\Sigma}_0 C' + DD')^{-1}(z_0 - C\bar{\xi}_0) = \begin{pmatrix} a_0 \\ \bar{\theta} \end{pmatrix} \\ \Sigma_0 &= \bar{\Sigma}_0 - \bar{\Sigma}_0 C'(C\bar{\Sigma}_0 C' + DD')^{-1} C\bar{\Sigma}_0 = \begin{pmatrix} 0 & 0 \\ 0 & \bar{P} \end{pmatrix} \end{aligned} \tag{43}$$

In what follows, let

$$\widehat{\theta}_0 = \bar{\theta} \quad \text{and} \quad P_0 = \bar{P}$$

so that

$$\widehat{\xi}_0 = \begin{pmatrix} a_0 \\ \widehat{\theta}_0 \end{pmatrix} \quad \text{and} \quad \Sigma_0 = \begin{pmatrix} 0 & 0 \\ 0 & P_0 \end{pmatrix}$$

Now time moves on to $t = 1$ and we have

$$\begin{aligned}\Sigma_e(z_0, m_0, \Sigma_0) &= CA(z_0, m_0)\Sigma_0 A(z_0, m_0)'C' + [CB(z_0, m_0) + D][CB(z_0, m_0) + D]' \\ &= a_0^2 P_0 + \sigma_{a1}^2 \end{aligned} \tag{44}$$

so

$$\begin{aligned}K(z_0, m_0, \Sigma_0) &= \left[A(z_0, m_0)\Sigma_0 A(z_0, m_0)'C' + B(z_0, m_0)B(z_0, m_0)'C' + B(z_0, m_0)D'\right]\Sigma_e(z_0, m_0, \Sigma_0)^{-1} \\ &= \begin{pmatrix} 1 \\ \frac{a_0 P_0}{a_0^2 P_0 + \sigma_{a1}^2} \end{pmatrix} \end{aligned} \tag{45}$$

and

$$\begin{aligned}\Sigma_1 &= A(z_0, m_0)\Sigma_0 A(z_0, m_0)' + B(z_0, m_0)B(z_0, m_0)' - K(z_0, m_0, \Sigma_0)\Sigma_e(z_0, m_0, \Sigma_0)K(z_0, m_0, \Sigma_0)' \\ &= \begin{pmatrix} 0 & 0 \\ 0 & \sigma_\theta^2 + \frac{\sigma_{a1}^2}{a_0^2 P_0 + \sigma_{a1}^2} P_0 \end{pmatrix} \equiv \begin{pmatrix} 0 & 0 \\ 0 & P_1 \end{pmatrix} \end{aligned} \tag{46}$$

Finally, we have

$$\begin{aligned}\widehat{\xi}_1 &= A(z_0, m_0)\widehat{\xi}_0 + K(z_0, m_0, \Sigma_0)\Sigma_e(z_0, m_0, \Sigma_0)^{1/2}\epsilon_1 \\ z_1 &= CA(z_0, m_0)\widehat{\xi}_0 + \Sigma_e(z_0, m_0, \Sigma_0)^{1/2}\epsilon_1,\end{aligned}$$

where $\epsilon_1 \sim N(0, 1)$ so using (44) and (45), gives

$$\begin{aligned}\widehat{\xi}_1 &= \begin{pmatrix} a_0\widehat{\theta}_0 \\ \widehat{\theta}_0 \end{pmatrix} + \begin{pmatrix} \sqrt{a_0^2 P_0 + \sigma_{a1}^2} \\ \frac{a_0 P_0}{\sqrt{a_0^2 P_0 + \sigma_{a1}^2}} \end{pmatrix} \epsilon_1 \\ z_1 &= a_0\widehat{\theta}_0 + \sqrt{a_0^2 P_0 + \sigma_{a1}^2}\epsilon_1\end{aligned}$$

Since $\xi_1 = (a_1, \theta_1)$ and $z_1 = a_1$ it follows from the equations above that the first component $\widehat{\xi}_{11}$ of $\widehat{\xi}_1$ equals z_1 which is as we would expect.

Step 3: Full information equivalent model In what follows let $\widehat{\theta}_1$ denote the second component of $\widehat{\xi}_1$, and more generally, for any $t = 0, 1, \ldots$, let $\widehat{\theta}_t$ denote the second component of $\widehat{\xi}_t$. In addition, let

P_t denote the bottom diagonal element of Σ_t. Iterating in t then gives, for every $t = 0, 1, \ldots$,

$$\begin{aligned} a_{t+1} &= a_t \widehat{\theta}_t + \sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}\epsilon_{t+1} \\ \widehat{\theta}_{t+1} &= \widehat{\theta}_t + \frac{a_t P_t}{\sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}}\epsilon_{t+1} \\ P_{t+1} &= \sigma_\theta^2 + \frac{\sigma_{a,t+1}^2}{a_t^2 P_t + \sigma_{a,t+1}^2} P_t \\ \sigma_{a,t+1}^2 &= \exp(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t) \end{aligned} \tag{47}$$

where $\{\epsilon_t\}_{t=1}^{+\infty} \sim$ iid $N(0,1)$, and where $\widehat{\theta}_0 = \bar{\theta}$, $P_0 = \bar{P}$, σ_{a0}^2 given.

The exogenous states in the full information equivalent of the model (1)-(2) are then

$$(a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$$

and for every $t = 0, 1, \ldots$, the optimal consumption $C_t = C(K_{t-1}, a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$ is obtained by solving the stationary Euler equation

$$E\left[\beta\left(\alpha \exp(a_t) K_{t-1}^{\alpha-1} + (1-\delta)\right)\left(\frac{C_{t+1}}{C_t}\right)^{-\nu} \;\middle|\; K_{t-1}, a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2\right] = 1 \quad a.s. \quad t = 0, 1, \ldots$$

in which the exogenous states $(a_{t+1}, \widehat{\theta}_{t+1}, P_{t+1}, \sigma_{a,t+1}^2)$ follow (47) and the endogenous state $K_t = (1-\delta)K_{t-1} + \exp(a_t)K_{t-1}^\alpha - C(K_{t-1}, a_t, \widehat{\theta}_t, P_t, \sigma_{at}^2)$, with K_{-1} given.

Step 4: Small noise limit Now consider the case in which the variances of the shocks ϵ_{at} and $\epsilon_{\theta t}$ are small as in (5), i.e.

$$\{\epsilon_{at}\}_{t=1}^{+\infty} \sim \text{ iid } N(0, \sigma^2) \quad \text{and} \quad \{\epsilon_{\theta t}\}_{t=1}^{+\infty} \sim \text{ iid } N(0, \sigma^2), \quad \text{with} \quad \sigma \to 0.$$

In this case, the matrices $B(z_t, m_t)$ and D in equations (41) and (42) need to be replaced with $\sigma B(z_t, m_t)$ and σD. Assuming the prior is modified as $\theta_0 \sim N(\bar{\theta}, \sigma^2 \bar{P})$, and following the same calculations as in Steps 1 and 2 above, it is straighforward to see that this induces the following transformations on the system matrices in (44), (45) and (46):

$$K(z_t, m_t, \Sigma_t) \mapsto K(z_t, m_t, \sigma^{-2}\Sigma_t), \quad \Sigma_e(z_t, m_t, \Sigma_t) \mapsto \sigma^2 \Sigma_e(z_t, m_t, \sigma^{-2}\Sigma_t), \quad \Sigma_t \mapsto \sigma^2 \Sigma_t$$

Defining

$$\widehat{\theta}_t = E(\theta_t \mid a_t, \ldots, a_0) \quad \text{and} \quad P_t = \frac{1}{\sigma^2} E[(\theta_t - \widehat{\theta}_t)^2 \mid a_t, \ldots, a_0]$$

then gives

$$
\begin{aligned}
a_{t+1} &= a_t \widehat{\theta}_t + \sigma \sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}\, \epsilon_{t+1} \\
\widehat{\theta}_{t+1} &= \widehat{\theta}_t + \sigma \frac{a_t P_t}{\sqrt{a_t^2 P_t + \sigma_{a,t+1}^2}} \epsilon_{t+1} \\
P_{t+1} &= \sigma_\theta^2 + \frac{\sigma_{a,t+1}^2}{a_t^2 P_t + \sigma_{a,t+1}^2} P_t \\
\sigma_{a,t+1}^2 &= \exp(\omega_a + \lambda_a \ln \sigma_{at}^2 + \alpha_a a_t)
\end{aligned}
$$

with $\{\epsilon_t\}_{t=1}^{+\infty} \sim$ iid $N(0,1)$, which yields the result in the main text.

References

Acemoglu, D. (2009): *Introduction to Modern Economic Growth*. Princeton University Press.

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109(4), 1263â89.

Anderson, B. D., and J. B. Moore (1979): *Optimal Filtering*. Prentice-Hall, Inc.

Auerbach, A. J., and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4(2), 1â27.

Bahadur, R. R. (1954): "Sufficiency and Statistical Decision Functions," *The Annals of Mathematical Statistics*, 25(3), 423–462.

Bertsekas, D. P., and S. E. Shreve (1978): *Stochastic Optimal Control: The Discrete-Time Case*. Academic Press.

Billingsley, P. (1995): *Probability and Measure*. John Wiley & Sons New York.

Blackwell, D. (1965): "Discounted Dynamic Programming," *The Annals of Mathematical Statistics*, 36(1), pp. 226–235.

Chen, H.-F., P. R. Kumar, and J. H. van Schuppen (1989): "On Kalman Filter for Conditionally Gaussian Systems with Random Matrices," *Systems & Control Letters*, 5(13), 397–404.

Cogley, T., and T. J. Sargent (2005): "Drifts and volatilities: monetary policies and outcomes in the post WWII US," *Review of Economic Dynamics*, 8(2), 262–302, Monetary Policy and Learning.

Fleming, W. (1971): "Stochastic Control for Small Noise Intensities," *SIAM Journal on Control*, 9(3), 473–517.

Fleming, W., and R. Rishel (1975): *Deterministic and Stochastic Optimal Control*. Springer Berlin.

Gonzalez-Sanches, D., and O. Hernandez-Lerma (2013): *Discrete-Time Stochastic Control and Dynamic Potential Games*. Springer Berlin.

Harvey, A. (1989): *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.

Jin, H.-H., and K. L. Judd (2002): "Perturbation Methods For General Dynamic Stochastic Models," Working paper, Stanford University.

Judd, K. L. (1996): *Oxford Handbook of Innovation*vol. I, chap. Approximation, Perturbation, and Projection Methods in Economic Analysis. Handbook of Computational Economics.

Liptser, R., and A. Shiryayev (1978): *Statistics of Random Processes II: Applications*. Springer-Verlag.

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 72(3), 821–852.

Rhenius, D. (1974): "Incomplete Information in Markovian Decision Models," *The Annals of Statistics*, 2(6), pp. 1327–1334.

Santos, M. S. (2000): "Accuracy of Numerical Solutions Using the Euler Equation Residuals," *Econometrica*, 68(6), 1377–1402.

Sokey, N., and R. E. Lucas (1989): *Recursive methods in economic dynamics*. Harvard University Press.

Stokey, N. L., and R. E. Lucas (1989): *Recursive Methods in Economics Dynamics*. Harvard University Press.

Striebel, C. (1965): "Sufficient statistics in the optimum control of stochastic systems," *Journal of Mathematical Analysis and Applications*, 12(3), 576 – 592.

Yushkevich, A. A. (1976): "Reduction of a Controlled Markov Model with Incomplete Data to a Problem with Complete Information in the Case of Borel State and Control Space," *Theory of Probability & Its Applications*, 21(1), 153–158.