



Linear State Space Models, Asset Pricing, and the Kalman Filter

Undergraduate Computational Macro

Jesse Perla

jesse.perla@ubc.ca

University of British Columbia



Table of contents

- Overview
- Linear State Space Models
- Introduction to the Kalman Filter
- Models of Expectations
- Risk-Neutral Asset Pricing



Overview



Motivation and Materials

- In this section we introduce a class of dynamic models that are widely used in economics and finance
- Unlike the previous sections, we will be separating out the equations for the “evolution” of the state and the “observation”
- The main applications will be some simple models of asset pricing, but we will use this machinery in the next section on the permanent income model
- For the asset pricing examples, we will be building off the deterministic versions we discussed **previously**
- Finally, we will introduce the Kalman Filter: a workhorse for estimation, implementing learning in dynamic models, and machine learning

Materials

- Adapted from QuantEcon lectures coauthored with John Stachurski and Thomas J. Sargent
 - **Linear State Space Models**
 - **A First Look at the Kalman Filter**
- The new package, `QuantEcon.jl` is used for some of the code examples for easy simulation

```
1 using Distributions, Plots, LaTeXStrings, LinearAlgebra, Statistics
2 using Plots.PlotMeasures, QuantEcon, StatsPlots
3 default(;legendfontsize=16, linewidth=2, tickfontsize=12,
4         bottom_margin=15mm)
```



Linear State Space Models

State Space Models

- State space models describe state and observation dynamics
 - $\mathbf{x}_t \in \mathbb{R}^n$ denoting the **state**, which may be “latent”
 - $\mathbf{y}_t \in \mathbb{R}^k$ **observables** of that state
 - $\mathbf{w}_{t+1} \in \mathbb{R}^m$ **shocks** which cannot be forecasted
- Where the model includes a pair of equations
 - A law of motion of a state variable \mathbf{x}_t (the “evolution equation”)
 - A law of motion of the observables \mathbf{y}_t given the state \mathbf{x}_t (the “observation equation”)
- A recursive, Markovian model is the goal. Linearity is convenient

Primitives for a LSS

- $A \in \mathbb{R}^{n \times n}$ **transition matrix**
- $C \in \mathbb{R}^{n \times m}$ **volatility matrix**
- $G \in \mathbb{R}^{k \times n}$ **observation matrix** (or output matrix)
- Then the LSS is given by

$$\begin{aligned}x_{t+1} &= Ax_t + Cw_{t+1}, && \text{evolution equation} \\y_t &= Gx_t, && \text{observation equation} \\w_{t+1} &\sim \mathcal{N}(0, I) && \text{shocks}\end{aligned}$$

Initial Conditions

- The initial condition \mathbf{x}_0 could be given, or it could be a distribution
- Given $\boldsymbol{\mu}_0 \in \mathbb{R}^n$ and $\boldsymbol{\Sigma}_0 \in \mathbb{R}^{n \times n}$, a (positive semi-definite) covariance matrix

$$\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$$

- Note that if $\boldsymbol{\Sigma}_0 = [\mathbf{0}]$ then $\mathbf{x}_0 = \boldsymbol{\mu}_0$ deterministically
- Later, when we discuss the Kalman Filter, we will consider this as a “prior” distribution over possible \mathbf{x}_0 states

Example: Difference Equation

- Let $\{y_t\}$ be a deterministic sequence that satisfies

$$y_{t+1} = \phi_0 + \phi_1 y_t + \phi_2 y_{t-1}$$

- Given a y_0, y_{-1}
- Map this into the LSS by choosing a x_t
- “Finding the state is an art”

Example: Difference Equation in LSS


- Fulfill: $y_{t+1} = \phi_0 + \phi_1 y_t + \phi_2 y_{t-1}$
- Define $x_t = [1 \quad y_t \quad y_{t-1}]^\top$, $w_{t+1} \in \mathbb{R}^1$

$$\underbrace{\begin{bmatrix} 1 \\ y_{t+1} \\ y_t \end{bmatrix}}_{\equiv x_{t+1}} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ \phi_0 & \phi_1 & \phi_2 \\ 0 & 1 & 0 \end{bmatrix}}_{\equiv A} \underbrace{\begin{bmatrix} 1 \\ y_t \\ y_{t-1} \end{bmatrix}}_{\equiv x_t} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}}_{\equiv C} \underbrace{[w_{t+1}]}_{\equiv w_{t+1}}$$

$$y_t = \underbrace{[0 \quad 1 \quad 0]}_{\equiv G} \underbrace{\begin{bmatrix} 1 \\ y_t \\ y_{t-1} \end{bmatrix}}_{\equiv x_t}$$

Simulation

```
1 phi0, phi1, phi2 = 1.1, 0.8, -0.8
2 A = [1.0 0.0 0
3       phi0 phi1 phi2
4       0.0 1.0 0.0]
5 C = zeros(3, 1)
6 G = [0.0 1.0 0.0]
7 y_0 = 1.0
8 y_m1 = 1.0
9 mu_0 = [1.0, y_0, y_m1]
10 lss = LSS(A, C, G; mu_0)
11 x, y = simulate(lss, 50)
12 plot(y'; xlabel = L"t", label = L"y_t",
13       size=(600, 400))
```



Example: Auto-Regressive Process


- Fulfill: $\mathbf{y}_{t+1} = \phi_1 \mathbf{y}_t + \phi_2 \mathbf{y}_{t-1} + \phi_3 \mathbf{y}_{t-2} + \phi_4 \mathbf{y}_{t-3} + \sigma \mathbf{w}_{t+1}$

$$\underbrace{\begin{bmatrix} \mathbf{y}_{t+1} \\ \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \end{bmatrix}}_{\equiv \mathbf{x}_{t+1}} = \underbrace{\begin{bmatrix} \phi_1 & \phi_2 & \phi_3 & \phi_4 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\equiv A} \underbrace{\begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \mathbf{y}_{t-3} \end{bmatrix}}_{\equiv \mathbf{x}_t} + \underbrace{\begin{bmatrix} \sigma \\ 0 \\ 0 \\ 0 \end{bmatrix}}_{\equiv C} \underbrace{[\mathbf{w}_{t+1}]}_{\equiv \mathbf{w}_{t+1}}$$
$$\mathbf{y}_t = \underbrace{[1 \ 0 \ 0 \ 0]}_{\equiv G} \underbrace{[\mathbf{y}_t \ \mathbf{y}_{t-1} \ \mathbf{y}_{t-2} \ \mathbf{y}_{t-3}]^\top}_{\equiv \mathbf{x}_t}$$



Simulation

```
1 phi1, phi2, phi3, phi4 = 0.5, -0.2, 0, 0.5
2 sigma = 0.2
3 A = [phi1 phi2 phi3 phi4
4       1.0 0.0 0.0 0.0
5       0.0 1.0 0.0 0.0
6       0.0 0.0 1.0 0.0]
7 C = [sigma
8       0.0
9       0.0
10      0.0]
11 G = [1.0 0.0 0.0 0.0]
12 mu_0 = ones(4)
13 lss = LSS(A, C, G; mu_0)
14 x, y = simulate(lss, 200)
15 plot(y'; xlabel = L"t", label = L"y_t",
16       size=(600, 400))
```



Moments and Forecasts

- Given $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, can forecast $\mathbf{x}_{t+1} \sim \mathcal{N}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$

$$\boldsymbol{\mu}_{t+1} = \mathbf{A}\boldsymbol{\mu}_t$$

$$\boldsymbol{\Sigma}_{t+1} = \mathbf{A}\boldsymbol{\Sigma}_t\mathbf{A}^\top + \mathbf{C}\mathbf{C}^\top$$

- And given some $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$

$$\mathbf{y}_{t+1} \sim \mathcal{N}(\mathbf{G}\boldsymbol{\mu}_t, \mathbf{G}\boldsymbol{\Sigma}_t\mathbf{G}^\top)$$

Forecasts and Expected Net Present Values

- Given $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, we can forecast \mathbf{x}_{t+j} and \mathbf{y}_{t+j} for any j

$$\mathbb{E}_t \mathbf{x}_{t+j} = \mathbf{A}^j \boldsymbol{\mu}_t$$

$$\mathbb{E}_t \mathbf{y}_{t+j} = \mathbf{G} \mathbf{A}^j \boldsymbol{\mu}_t$$

- Useful for computing expected net present values of future cash flows

$$\begin{aligned} \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \mathbf{y}_{t+j} &= \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t \mathbf{y}_{t+j} = \sum_{j=0}^{\infty} \beta^j \mathbf{G} \mathbf{A}^j \boldsymbol{\mu}_t \\ &= \mathbf{G} (\mathbf{I} - \beta \mathbf{A})^{-1} \boldsymbol{\mu}_t \end{aligned}$$

Stationary Distributions

- If they exist, from any gaussian initial condition, the stationary distribution is $x_\infty \sim \mathcal{N}(\mu_\infty, \Sigma_\infty)$
- Must fulfill the fixed points of the previous iteration,

$$\begin{aligned}\mu_\infty &= A\mu_\infty \\ \Sigma_\infty &= A\Sigma_\infty A^\top + CC^\top\end{aligned}$$

- The first is an eigenvalue problem
- The second is a discrete **Lyapunov equation**



Introduction to the Kalman Filter

Noisy Observation Equation

- Given \mathbf{A} and \mathbf{G} matrices, you may be able to recover \mathbf{x}_t from the \mathbf{y}_t
- What if the observations in the LSS are noisy? Then \mathbf{x}_t is truly “latent”

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{C}\mathbf{w}_{t+1}$$

$$\mathbf{y}_t = \mathbf{G}\mathbf{x}_t + \mathbf{H}\mathbf{v}_t$$

$$\mathbf{w}_{t+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

$$\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

→ If $\mathbf{H} = [\mathbf{0}]$, then noiseless observation

Tracking the Distribution of the State

- Can the latent \mathbf{x}_t state be estimated from the noisy $\{\mathbf{y}_0, \dots, \mathbf{y}_t\}$ observations?
- Then we can forecast the state \mathbf{x}_{t+j} and future observables \mathbf{y}_{t+j}
 - However, we also need to “nowcast” the state \mathbf{x}_t since the state is unknown and our observations are noisy
- If we assume that $\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, then we interpret this as a “prior” distribution over the possible states of \mathbf{x}_0
- In that case, we use the \mathbf{y}_1 observation to update our beliefs about the state \mathbf{x}_0 to get a new distribution over the state \mathbf{x}_0
- This is a **Bayesian** approach, $\mathbb{P}(\mathbf{x}_t | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots) \propto \mathbb{P}(\mathbf{y}_t | \mathbf{x}_t) \mathbb{P}(\mathbf{x}_t | \mathbf{y}_{t-1}, \dots)$

Bayesian Approach with Normal Distributions

- In particular, we want to take our “prior” $\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and \mathbf{y}_1 to build new beliefs about $\mathbf{x}_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$
 - This is more complicated than a normal Bayesian update because the \mathbf{x}_t is moving with the evolution equation
- The key here, as with our **derivation with the AR(1)** is that a **linear combination of Gaussians is Gaussian**
- Because of this, it is sufficient to write a recurrence for $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$
 - Given $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ and \mathbf{y}_{t+1} what is $\mathbf{x}_{t+1} \sim \mathcal{N}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$?

Kalman Filter

- The Kalman Filter is the recursive, Bayesian updating of the distribution of the state \mathbf{x}_{t+1} given the observations \mathbf{y}_{t+1} and a prior $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$
- See [here](#) and other places for the derivation

$$\begin{aligned} \mathbf{K}_t &\equiv \mathbf{A}\boldsymbol{\Sigma}_t\mathbf{G}^\top (\mathbf{G}\boldsymbol{\Sigma}_t\mathbf{G}^\top + \mathbf{H}\mathbf{H}^\top)^{-1} \\ \boldsymbol{\mu}_{t+1} &= \mathbf{A}\boldsymbol{\mu}_t + \mathbf{K}_t(\mathbf{y}_t - \mathbf{G}\boldsymbol{\mu}_t) \\ \boldsymbol{\Sigma}_{t+1} &= \mathbf{A}\boldsymbol{\Sigma}_t\mathbf{A}^\top - \mathbf{K}_t\mathbf{G}\boldsymbol{\Sigma}_t\mathbf{A}^\top + \mathbf{C}\mathbf{C}^\top \end{aligned}$$

- \mathbf{K}_t is the “Kalman Gain” and $\mathbf{y}_t - \mathbf{G}\boldsymbol{\mu}_t$ is called the “innovation”
- The last equation is called a matrix Ricatti equation

The **Kalman Smoother** is when you go back at time t and update all of your previous distributions $\{\mathbf{x}_0, \dots, \mathbf{x}_{t-1}\}$ given **all** of the observations $\{\mathbf{y}_0, \dots, \mathbf{y}_t\}$.

Interpreting the Gain

- Consider the simple case where $\mathbf{x}_t \in \mathbb{R}$, $\mathbf{y}_t \in \mathbb{R}$, $\mathbf{A} = \mathbf{1}$, $\mathbf{G} = \mathbf{1}$ and $\mathbf{C}, \mathbf{H} \in \mathbb{R}$

$$\begin{aligned} \mathbf{K}_t &= \Sigma_t / (\Sigma_t + \mathbf{H}^2) \\ \mu_{t+1} &= \mu_t + \underbrace{\mathbf{K}_t(\mathbf{y}_t - \mu_t)}_{\text{innovation}} = (1 - \mathbf{K}_t)\mu_t + \mathbf{K}_t\mathbf{y}_t \\ \Sigma_{t+1} &= (1 - \mathbf{K}_t)\Sigma_t + \mathbf{C}^2 \end{aligned}$$

- The μ_{t+1} equation is a weighted average of the forecast (i.e. μ_t since $\mathbf{A} = \mathbf{1}$) and the observation \mathbf{y}_t
- \mathbf{K}_t says how much to update the forecast of the mean. Small “gain” means less weight on new observations

Forecasting and Nowcasting

- Future states are forecasted by the Kalman Filter itself
- The state is a hidden Markov variable, but we can forecast the current state and future observations
- Current state is constructed to be $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$
- Given a \mathbf{x}_t distribution, can get the \mathbf{y}_t distribution as

$$\mathbf{y}_t \sim \mathcal{N}(G\boldsymbol{\mu}_t, G\boldsymbol{\Sigma}_tG^\top + \mathbf{H}\mathbf{H}^\top)$$

- Useful for forecasting (i.e., what would the observation distribution be for a future distribution)
- Also useful for estimation and likelihoods in structural models

Different Canonical Forms

- When looking at software packages, you may need to map to different version
- For example, another common one is

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{w}_{t+1}$$

$$\mathbf{y}_t = \mathbf{G}\mathbf{x}_t + \mathbf{v}_t$$

$$\mathbf{w}_{t+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$$

$$\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$$

- Which maps to ours if $\mathbf{Q} = \mathbf{C}\mathbf{C}^\top$ and $\mathbf{R} = \mathbf{H}\mathbf{H}^\top$
- Can go other direction with a Cholesky decomposition
- Others may have an additional “control” term in the \mathbf{x}_{t+1} equation

Example Implementation

- We will use the `QuantEcon.jl` package for the Kalman Filter, uses the \mathbf{Q} and \mathbf{R} form
- Consider a univariate function

$$\mathbf{x}_{t+1} = \mathbf{x}_t$$

$$\mathbf{y}_t = \mathbf{x}_t + \mathbf{v}_t$$

$$\mathbf{v}_t \sim \mathcal{N}(0, 1.0^2)$$

→ We will assume that $\mathbf{x}_0 \sim \mathcal{N}(8.0, 1.0^2)$

→ We will assume that the true $\mathbf{x}_0 = \mathbf{10.0}$ and hence $\mathbf{x}_t = \mathbf{10.0}$ for all t


Simulation

```
1 A, G, Q, R = 1.0, 1.0, 0.0, 1.0
2 x_hat_0, Sigma_0 = 8.0, 1.0
3 x_true = 10.0
4 # initialize Kalman filter
5 kalman = Kalman(A, G, Q, R)
6 set_state!(kalman, x_hat_0, Sigma_0)
7 plt = plot(;title="First 5 Densities")
8 for i in 1:5
9     # record the current predicted mean and variance, and plot their densities
10    m, v = kalman.cur_x_hat, kalman.cur_sigma
11    plot!(Normal(m, sqrt(v)); label = "t = $i")
12    # Generate signal and update
13    y = x_true + sqrt(R) * randn() # i.e. x_t + v_t
14    update!(kalman, y)
15 end
16 plt
```



Simulation

First 5 Densities



Applications of the Kalman Filter

- The LSS with noisy observation is an example of a “hidden Markov model”
 - i.e., Observe only a noisy version of a Markovian state
- The Kalman Filter is used in many applications
 - Estimating and forecasting the state of the economy given noisy data
 - Estimating the state of a latent variable or forming a likelihood in a structural model
 - Machine learning and reinforcement learning (e.g., estimating the position of a car or pedestrian given noisy sensor data)
 - **Apollo 11** used a Kalman Filter to estimate the position of the spacecraft



Models of Expectations

Forecasts and Expectations

- The emphasis on stochastic processes serves a dual role
 1. Model of the economy which you can conduct quantitative experiments as an “econometrician”
 2. Model of the formation of expectations and the pricing of assets for an “agent” inside of the model
- This wasn’t required when we have exogenously given, ad-hoc decisions like the savings rate previously
 - But if we want to model agent decisions, they need to form expectations about the future
 - Without that model of decisions, can we conduct policy counterfactuals?

Alternative Approaches

- The baseline approach for most of macroeconomics in the 1960s+ is called “rational expectations”
 - Use the mathematical expectation, and assume agent’s have a well-specified model of the economy
- Using that as a baseline, there are many models of **bounded rationality** which deviate from this in various ways
 - e.g., what it agents don’t fully know the evolution of the economy but have priors (use Kalman Filter?)
 - what is agent’s only learn from their own past observations? The oldest versions of this are called “adaptive expectations” and it is related to modern methods in machine learning

Using the Mathematical Expectation

- A good starting point for a model of expectations is to assume that agents use the mathematical expectation
- This requires that they have an internal model of a stochastic process for the data-generating process - conditional on their choices
 - Then they can use probabilities to calculate expected values
- One benefit is that we can use the mathematical expectation and its properties (e.g., linearity)

$$\mathbb{E}(aX + bY|Z) = a\mathbb{E}(X|Z) + b\mathbb{E}(Y|Z)$$

- Requires a model of joint distribution of the data-generating process, and theory of which values to condition on

Information Sets

- Think of the values we can condition on in our expectations as being the “information set” of the agent
 - For stochastic processes that unfold over time, a good default is to think of all information up to time t being available
 - Call this the “Information Set”. Shorthand denote with subscript

$$\mathbb{E}[X_{t+1} | \underbrace{X_t, X_{t-1}, \dots}_{\text{Information Set}}] = \mathbb{E}_t[X_{t+1}]$$

- If Markov, information set is summarized by the current state

$$\mathbb{E}[X_{t+1} | X_t, X_{t-1}, \dots] = \mathbb{E}[X_{t+1} | X_t]$$

Law of Iterated Expectations

- Frequently you will find yourself taking expectations of future expectations
- A useful property of mathematical expectations is the “Law of Iterated Expectations”

$$\mathbb{E}_t[\mathbb{E}_{t+1}[\mathbf{X}_{t+2}]] = \mathbb{E}_t[\mathbf{X}_{t+2}]$$
$$\mathbb{E}[\mathbb{E}[\mathbf{X}_{t+2} | \mathbf{X}_{t+1}, \mathbf{X}_t, \dots] | \mathbf{X}_t, \mathbf{X}_{t-1}, \dots] = \mathbb{E}[\mathbf{X}_{t+2} | \mathbf{X}_t, \mathbf{X}_{t-1}, \dots]$$

Models of Learning a Hidden State

- Specifying the information set is a key part of economic models
- If a state is hidden, then the agent must expectations from observables
- Models of learning a hidden value or latent state are often built around some form of state-space model
- For example, with a LSS model $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{C}\mathbf{w}_{t+1}$ and $\mathbf{y}_t = \mathbf{G}\mathbf{x}_t + \mathbf{H}\mathbf{v}_t$

$$\mathbb{E}(\mathbf{x}_{t+1} | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots)$$

- In that case, could use a posterior probabilities from Kalman Filter
- Not all models of learning are Bayesian, but economists often case about

$$\mathbb{E}[(\mathbf{x}_{t+1} - \mathbb{E}[\mathbf{x}_{t+1} | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots])^2 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots]$$

Forecast Errors

- With either a learning model or one with noiseless observations, we can define the one-period ahead forecast error as

$$FE_{t,t+1} \equiv \mathbf{x}_{t+1} - \mathbb{E}_t[\mathbf{x}_{t+1}]$$

- With our LSS, the information sets are $\mathbf{x}_{t+1}, \mathbf{w}_{t+1}, \mathbf{x}_t, \mathbf{w}_t, \dots$ vs. $\mathbf{x}_t, \mathbf{w}_t, \dots$

$$FE_{t,t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{C}\mathbf{w}_{t+1} - \mathbb{E}_t[\mathbf{A}\mathbf{x}_t + \mathbf{C}\mathbf{w}_{t+1}] = \mathbf{C}\mathbf{w}_{t+1}$$

- The forecast error is a random variable, but it is uncorrelated with the information set

Systematic Bias in Forecasts

- How far off do you expect your forecasts to be?

$$\mathbb{E}_t[FE_{t,t+1}] = \mathbb{E}_t[x_{t+1} - \mathbb{E}_t[x_{t+1}]] = 0$$

- That comes from the linearity of expectations
- A hallmark of rational expectations is that agents don't systematically over or underestimate the future
- If they did, why not just manually adjust fudge expectations?

Variance of Forecast Errors

- While there is no systematic bias, that doesn't mean the forecasts are correct
- In some cases the agent may care deeply about how precise they are
- For a LSS we can calculate the variance (using the mean zero result)

$$\begin{aligned}\mathbb{V}_t(\mathbf{F}\mathbf{E}_{t+1}) &\equiv \mathbb{E}_t[\mathbf{F}\mathbf{E}_{t,t+1}\mathbf{F}\mathbf{E}_{t,t+1}^\top] \\ &= \mathbb{E}_t[(\mathbf{C}\mathbf{w}_{t+1})(\mathbf{C}\mathbf{w}_{t+1})^\top] = \mathbf{C}\mathbf{C}^\top\end{aligned}$$

- With the observation equation (and possible measurement error)

$$\mathbb{V}_t(\mathbf{F}\mathbf{E}_{t+1}) = \mathbf{G}\mathbf{C}\mathbf{C}^\top\mathbf{G}^\top + \mathbf{H}\mathbf{H}^\top$$

Forecasting Error with Learning Models

- If \mathbf{x}_t is not in the information set, then (conditional on \mathbf{x}_t) the expected forecast error may not be zero.

→ With a LSS and a Kalman Filter and $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$,

$$\begin{aligned} FE_{t,t+1} &= \mathbf{y}_{t+1} - \mathbb{E}[\mathbf{y}_{t+1} | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots] \\ &= \mathbf{G}(\mathbf{A}\mathbf{x}_t + \mathbf{C}\mathbf{w}_{t+1}) - \mathbf{G}(\mathbf{A}\boldsymbol{\mu}_t) = \mathbf{G}\mathbf{A}(\mathbf{x}_t - \boldsymbol{\mu}_t) + \mathbf{C}\mathbf{w}_{t+1} \end{aligned}$$

- In that setup, the agent has an unbiased estimate if they use their \mathbf{x}_t estimate since $\mathbb{E}_t(\mathbf{x}_t) = \boldsymbol{\mu}_t$

$$\mathbb{E}_t[FE_{t,t+1}] = \mathbb{E}_t[\mathbf{G}\mathbf{A}(\mathbf{x}_t - \boldsymbol{\mu}_t) + \mathbf{C}\mathbf{w}_{t+1}] = \mathbf{0}$$

Martingales

- An important type of stochastic process are Martingales where

$$\mathbb{E}_t[\mathbf{X}_{t+1}] = \mathbf{X}_t$$

- i.e., in expectation, the future value is the current value
- Inductively you can see that $\mathbb{E}_t[\mathbf{X}_{t+j}] = \mathbf{X}_t$ for all j
- The canonical Martingale is a random walk $\mathbf{X}_{t+1} = \mathbf{X}_t + \mathbf{w}_{t+1}$ where $\mathbf{w}_{t+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$
- Forecasting: the best guess for the future is the current value
- Martingales have many applications in asset pricing and finance, models of learning, and in consumption and savings models



Risk-Neutral Asset Pricing

Risk-Neutrality

- Economically what does linearity in payoffs $u(c) = c$ mean?
 - It means that the agent is “risk-neutral” and only cares about the expected value of the payoff
 - This is a strong assumption, but it may be accurate in many cases (e.g., institutional investors)
- Risk averse agents have concave utility functions, and risk-loving agents have convex utility functions
- Linearity is especially useful for stochastic processes because we can use it with the mathematical expectation

Risk-Neutral Asset Pricing

- If an agent has no risk aversion, their preferences can be rationalized by something proportional to a linear utility
- For our LSS model we can just use the EPDV to calculate a price

$$\begin{aligned} p(\mathbf{x}_t) &= \mathbb{E} \left[\sum_{j=0}^{\infty} \beta^j \mathbf{y}_{t+j} \mid \mathbf{x}_t \right] \\ &= \mathbf{G}(\mathbf{I} - \beta \mathbf{A})^{-1} \mathbf{x}_t \end{aligned}$$

- The interpretation is that the agent is using their internal model to forecast the evolution of the state and the observable payoff
- Also has a second interpretation called “certainty equivalent” where the agent is indifferent to the risks and volatility in any decisions (i.e., no \mathbf{C})

Risk-Neutral Asset Pricing with Hidden States

- If the state is hidden, then the agent must use their internal model to forecast the future
- With a Kalman Filter, this becomes

$$\begin{aligned} p(\mu_t, \Sigma_t) &= \mathbb{E} \left[\sum_{j=0}^{\infty} \beta^j y_{t+j} \mid \mu_t, \Sigma_t \right] \\ &= G(I - \beta A)^{-1} \mu_t \end{aligned}$$

Example: AR(1) Process


- Consider the AR(1) process $y_{t+1} = a + \rho y_t + \sigma w_{t+1}$ with $w_{t+1} \sim \mathcal{N}(0, 1)$ and $y_0 = 0.5$
- Let $x_t \equiv [y_t \quad 1]^\top$ the LSS is

$$x_{t+1} = \underbrace{\begin{bmatrix} \rho & a \\ 0 & 1 \end{bmatrix}}_{\equiv A} x_t + \underbrace{\begin{bmatrix} \sigma \\ 0 \end{bmatrix}}_{\equiv C} w_{t+1}$$
$$y_t = \underbrace{[1 \quad 0]}_{\equiv G} x_t$$

- The price of the asset is then $p(x_t) = G(I - \beta A)^{-1} x_t$

Simulation

```
1 rho = 0.9
2 a = 1.0
3 sigma = 0.2
4 beta = 0.8
5 A = [rho a
6      0 1]
7 C = [sigma; 0]
8 G = [1 0]
9 x_0 = [0.5, 1.0]
10 lss = LSS(A, C, G; mu_0 = x_0)
11 x, y = simulate(lss, 30)
12 H = G * inv(I - beta * A)
13 p = H * x
14 plot(p'; label = L"p_t",
15       xlabel = L"t", size=(600, 400))
16 plot!(y', label = L"y_t")
```



Example: Wages and Productivity

- Wages $y_t = \theta z_t + (1 - \theta)q_t$. Human capital: $\mathbb{E}[\sum_{j=0}^{\infty} \beta^j y_{t+j} | z_t, q_t]$
 - Workers productivity follow $z_{t+1} = z_t + \alpha + \sigma w_{t+1}$ given z_0
 - Firm productivity follows $q_{t+1} = q_t + \gamma$ given q_0
- Guess a state of $x_t \equiv [z_t \quad q_t \quad 1]^\top$

$$x_{t+1} = \underbrace{\begin{bmatrix} 1 & 0 & \alpha \\ 0 & 1 & \gamma \\ 0 & 0 & 1 \end{bmatrix}}_{\equiv A} x_t + \underbrace{\begin{bmatrix} \sigma \\ 0 \\ 0 \end{bmatrix}}_{\equiv C} w_{t+1}$$
$$y_t = \underbrace{[\theta \quad 1 - \theta \quad 0]}_{\equiv G} x_t$$

Simulation

```
1 alpha, gamma = 0.1, 0.1
2 sigma = 0.2
3 theta = 0.5
4 beta = 0.8
5 A = [1 0 alpha
6      0 1 gamma
7      0 0 1]
8 C = [sigma; 0; 0]
9 G = [theta 1-theta 0]
10 mu_0 = [1.0, 1.0, 1.0]
11 lss = LSS(A, C, G; mu_0)
12 x, y = simulate(lss, 30)
13 H = G * inv(I - beta * A)
14 p = H * x
15 plot(p'; label = L"p_t",
16      xlabel = L"t", size=(600, 400))
```

