

ECON 2021 - FINANCIAL ECONOMICS I

Lecture 8 – Learning

November 5, 2018

MOTIVATION

- Thus far we have assumed that the only uncertainty agents confront is the future realization of asset returns. More specifically, we've assumed agents know the equations generating these processes. Only the realizations of future exogenous shocks are unknown. For example, they may not know what the future return will be, but they **know** that it is lognormally distributed with a particular mean and variance.
- Although these models work well qualitatively, they suffer from several **quantitative** shortcomings. For example,
 - 1 Prices are more volatile than predicted.
 - 2 Spreads between risky and safe asset returns are larger than predicted.
 - 3 Investors hold fewer risky assets than predicted.
- A natural suspect in these failures is the assumption that investors know the distribution of future returns. The usual defense for this kind of assumption is **history**, i.e., agents have already undergone a process of **adaptive learning**.
- But people don't live forever, they may trust their own experiences more than their ancestors', and the world may be constantly changing. So what if this learning process hasn't finished yet, or what if it is perpetually ongoing?

- We shall see that learning provides a plausible explanation of excess volatility and low market participation. It is somewhat less successful in explaining Sharpe ratios and the Equity Premium Puzzle.
- However, there are many ways to introduce learning into asset pricing models, and today we only study the simplest. We continue to assume agents have identical time-additive CRRA preferences, and although learning is a natural source of heterogeneous beliefs, today we continue to assume a representative agent. We also continue to assume agents 'know the model'. They just don't know one or more of its **parameters**. We study model uncertainty, heterogeneous beliefs, and recursive preferences in future lectures.
- Beliefs about the parameters are assumed to be updated using **Bayes Rule**. Although Bayesians have staked out the high moral ground by calling this 'Rational Learning', there is an active literature exploring the implications of other types of learning. These alternative learning strategies are only 'irrational' relative to a (questionable) set of axioms.
- Conceptually, our job is straightforward. All we need to do is replace objective probabilities with subjective beliefs. In continuous-time, this reduces to a simple Girsanov transformation, which alters a model's drift process. Parameters describing an agent's beliefs now become (hedgeable) state variables. We've already encountered this sort of thing.
- The only new tool we need is to figure out how to keep track of the evolution of beliefs. That is, we need to derive the law of motion for these new state variables.

TWO WORKHORSE FILTERS

- Beliefs are described by a probability distribution. Learning involves updating this distribution in response to new information. In general, probability distributions are infinite-dimensional. Updating infinite-dimensional objects isn't easy.
- 99% of the Bayesian learning literature in macro-finance uses one of two workhorse specifications that make this updating process mathematically tractable:
 - 1 **Kalman Filter:** Gaussian noise corrupted signals of an unobserved continuous-state Gaussian process.
 - 2 **Wonham Filter:** Gaussian noise corrupted signals of an unobserved discrete-state Markov chain with constant transition probabilities.
- Perhaps not surprisingly, the Gaussian distribution figures prominently in both. A Gaussian distribution is fully characterized by only *two* parameters, the mean and the variance.
- Often it is more natural to assume the hidden state can assume a continuum of values (e.g., an unknown mean return). This points to the Kalman filter. Unfortunately, the Kalman filter has the property that uncertainty about parameters monotonically decreases over time. Hence, the Kalman filter is not useful for modeling events that cause uncertainty to increase. In contrast, the Wonham filter easily accommodates this feature, and so has been increasingly used in the macro-finance literature.

A LIMITATION OF CONTINUOUS-TIME

- In continuous-time learning models, at least those based on Brownian motion, we assume agents only have to learn about drift parameters. Volatility parameters are assumed known.
- This is because continuous observations of returns permit agents to obtain noiseless estimates of σ , using even an arbitrarily short sample length. This is due to the highly volatile nature of dB , which is effectively the 'explanatory variable' when estimating σ . (Remember, variation in explanatory variables leads to more precise parameter estimates).
- For example, suppose $dS_t/S_t = \mu dt + \sigma dB_t$. Then note that $\left(\frac{dS_t}{S_t}\right)^2 = \sigma^2 dt + o(dt)$. Hence,

$$\frac{1}{T} \int_0^T \left(\frac{dS_t}{S_t}\right)^2 = \frac{1}{T} \int_0^T \sigma^2 dt = \sigma^2$$

and this is true for any $T \geq dt$!

- The fact that agents know σ creates challenges for continuous-time learning models to explain things like Sharpe ratios and risk prices, since these depend on volatilities and covariances.

JENSEN'S INEQUALITY

- We are used to thinking of uncertainty as 'bad', with price depressing effects. However, this isn't necessarily the case.
- When unknown parameters enter valuation formulas nonlinearly, then due to Jensen's inequality, uncertainty about parameter estimates influences prices. If the mapping between prices and parameters is **convex**, then parameter uncertainty will actually **increase** prices.
- Suppose dividends follow the process, $dD/D = gdt + \sigma dB$, and assume agents are risk-neutral. If g is known then the price is given by

$$P = \frac{D}{r - g}$$

- If g is unknown then

$$P = E \left\{ \frac{D}{r - g} \right\} > \frac{D}{r - E[g]}$$

since P is a convex function of g . Some have argued that this explains the dotcom boom of the late 1990s, since this was an era of increased growth uncertainty. Pastor & Veronesi (*JF*, 2003) argue that this explains why Market/Book ratios decline with firm age (because there is more uncertainty about young firms).

- Of course, with risk aversion, uncertainty is always bad, so the net effect will depend on a race between Jensen's inequality and risk aversion. The race is a tie for a log utility investor.

A DISCRETE-TIME REFRESHER

- The continuous-time Kalman filter is a straightforward generalization of the discrete-time Kalman filter, which you (hopefully) already know. (The continuous-time version is often called the ‘Kalman-Bucy filter’).
- Suppose you want to estimate an unknown constant, θ , given a sequence of observed signals, $s_t = \theta + \varepsilon_t$, where ε_t i.i.d. $N(0, \sigma^2)$. Assume your prior is $N(\theta_0, \sigma_0^2)$.
- After observing t signals your posterior is $N(\hat{\theta}_t, \hat{\sigma}_t^2)$,

$$\hat{\theta}_t = \left(\frac{\frac{1}{\sigma_0^2}}{\frac{1}{\sigma_0^2} + \frac{t}{\sigma^2}} \right) \theta_0 + \left(\frac{\frac{t}{\sigma^2}}{\frac{1}{\sigma_0^2} + \frac{t}{\sigma^2}} \right) \bar{s} \quad \frac{1}{\hat{\sigma}_t^2} = \frac{1}{\sigma_0^2} + \frac{t}{\sigma^2}$$

where $\bar{s} = t^{-1} \sum_{i=1}^t s_i$. Hence, your posterior is just a precision-weighted average of your prior and your new information.

- This can be written recursively as follows

$$\begin{aligned} \hat{\theta}_t &= \hat{\theta}_{t-1} + \gamma_t (s_t - \hat{\theta}_{t-1}) \\ \frac{1}{\hat{\sigma}_t^2} &= \frac{1}{\hat{\sigma}_{t-1}^2} + \frac{1}{\sigma^2} \quad \Rightarrow \quad \hat{\sigma}_t^2 = \hat{\sigma}_{t-1}^2 - \frac{(\hat{\sigma}_{t-1}^2)^2}{\sigma^2 + \hat{\sigma}_{t-1}^2} \end{aligned}$$

where $\gamma_t = \frac{\hat{\sigma}_t^2}{\sigma^2}$ is called the Kalman gain. It is just the signal-to-noise ratio.

- The Kalman-Bucy filter generalizes this to the case of diffusion processes. Notice that $\hat{\sigma}_t^2$ declines monotonically.

THE KALMAN-BUCY FILTER

- Now suppose the underlying hidden state, x_t , is time-varying, and it follows a diffusion process. Also suppose we continuously observe a signal, y_t , which also follows a diffusion process:

$$dx_t = Ax_t dt + CdB_t$$

$$dy_t = Dx_t dt + GdB_t$$

Notice that we are now chasing a moving target. The previous case of learning an unknown constant is a special case, where $A = C = 0$.

- Importantly, we assume (A, C, D, G) are **known** matrices. What is unknown is x_t . (If (A, C, D, G) are also unknown, the problem becomes nonlinear and exact solutions disappear).
- Let $\mathcal{F}_t =$ filtration generated by $\{y_t\}$. Denote $\hat{x}_t = E[x_t | \mathcal{F}_t]$ and $\Sigma_t = E[(x_t - \hat{x}_t)(x_t - \hat{x}_t)' | \mathcal{F}_t]$. Hence, Σ_t captures our uncertainty about x_t . We want to derive equations for \hat{x}_t and Σ_t .
- As always, it is useful to consider a discrete approximation. Note that $E[x_{t+\varepsilon} | x_t] \approx x_t + \varepsilon Ax_t$. Therefore, $E[x_{t+\varepsilon} | \mathcal{F}_t] \approx \hat{x}_t + \varepsilon A\hat{x}_t$. We can then approximate the innovation in x_t as follows

$$x_{t+\varepsilon} - \hat{x}_t - \varepsilon A\hat{x}_t \approx (x_t - \hat{x}_t) + \varepsilon A(x_t - \hat{x}_t) + C(B_{t+\varepsilon} - B_t)$$

- Similarly, we can approximate the evolution of the signal as follows

$$y_{t+\varepsilon} - y_t - \varepsilon D \hat{x}_t \approx \varepsilon D(x_t - \hat{x}_t) + G(B_{t+\varepsilon} - B_t)$$

- Since this is an innovation w.r.t. \mathcal{F}_t , we can change measures and write

$$\varepsilon D(x_t - \hat{x}_t) + G(B_{t+\varepsilon} - B_t) \approx \bar{G}(\bar{B}_{t+\varepsilon} - \bar{B}_t)$$

where \bar{B}_t is a Brownian motion w.r.t. \mathcal{F}_t . In continuous-time, we know we must have $\bar{G}\bar{G}' = GG'$.

- Now, exactly as in the discrete Kalman filter, we want to compute a regression of the hidden state innovation on the signal innovation

$$x_{t+\varepsilon} - \hat{x}_t - \varepsilon A \hat{x}_t = K_t(y_{t+\varepsilon} - y_t - \varepsilon D \hat{x}_t) + \eta_{t+\varepsilon}$$

where the regression coefficient is given by the usual least-squares formula,

$$K_t = E[(x_{t+\varepsilon} - \hat{x}_t - \varepsilon A \hat{x}_t)(y_{t+\varepsilon} - y_t - \varepsilon D \hat{x}_t)'] \cdot (E[(y_{t+\varepsilon} - y_t - \varepsilon D \hat{x}_t)(y_{t+\varepsilon} - y_t - \varepsilon D \hat{x}_t)'])^{-1}$$

- From the above results we can easily evaluate this

$$K_t = [\varepsilon CG' + \varepsilon \Sigma_t D' + \varepsilon^2 A \Sigma_t D'] (\varepsilon GG' + \varepsilon^2 D \Sigma_t D')^{-1} \rightarrow [CG' + \Sigma_t D'] (GG')^{-1}$$

- Letting $\varepsilon \rightarrow 0$ we get the Kalman-Bucy filter

$$d\hat{x}_t = A\hat{x}_t dt + K_t(dy_t - D\hat{x}_t dt) = A\hat{x}_t dt + K_t \bar{G} d\bar{B}_t \quad \text{Kalman-Bucy Filter}$$

- To compute the evolution of the conditional variance, Σ_t , notice that

$$x_{t+\varepsilon} - \hat{x}_{t+\varepsilon} = (I + \varepsilon A)(x_t - \hat{x}_t) + C(B_{t+\varepsilon} - B_t) - K_t(y_{t+\varepsilon} - y_t - \varepsilon D\hat{x}_t)$$

Using the fact $y_{t+\varepsilon} - y_t - \varepsilon D\hat{x}_t = G(B_{t+\varepsilon} - B_t) + \varepsilon D(x_t - \hat{x}_t)$ we can write this as

$$x_{t+\varepsilon} - \hat{x}_{t+\varepsilon} = (I + \varepsilon A - \varepsilon K_t D)(x_t - \hat{x}_t) + (C - K_t G)(B_{t+\varepsilon} - B_t)$$

- Now, if we square both sides, take expectations, and then drop $O(\varepsilon^2)$ terms we get

$$\Sigma_{t+1} = \Sigma_t + \varepsilon(A - K_t D)\Sigma_t + \varepsilon\Sigma_t(A - K_t D)' + (B - K_t G)(B - K_t G)'$$

Finally, letting $\varepsilon \rightarrow 0$, and rearranging gives

$$\frac{d\Sigma_t}{dt} = A\Sigma_t + \Sigma_t A' + CC' - K_t GG' K_t$$

- Notice that Σ_t evolves **deterministically** (ie, independently from the signals), and that it monotonically decreases. Uncertainty never increases. Also note that as long as $CC' \neq 0$ (ie, the hidden state is always moving), the steady state value of $\Sigma_t > 0$.

APPLICATION 1: PORTFOLIO CHOICE

- As a simple (univariate) example, let's revisit the Merton portfolio choice problem, but relax the assumption that the mean return is known. Suppose the risky asset follows the process

$$\frac{dS_t}{S_t} = \mu dt + \sigma dB_t$$

but now the investor doesn't know μ (but he does know it is constant). He must use observations of S_t to learn about it.

- Letting (m_t, V_t) be his conditional mean and variance of μ , the Kalman filter gives us

$$\begin{aligned} dm_t &= \frac{V_t}{\sigma^2} \left(\frac{dS_t}{S_t} - m_t dt \right) \\ dV_t &= -\frac{V_t^2}{\sigma^2} dt \end{aligned}$$

- By defining the innovation,

$$d\hat{B}_t = \frac{1}{\sigma} \left(\frac{dS_t}{S_t} - m_t dt \right) = dB_t + \left(\frac{\mu - m_t}{\sigma} \right) dt$$

we can write the perceived asset price process in terms of the investor's observed filtration,

$$\frac{dS_t}{S_t} = m_t dt + \sigma d\hat{B}_t$$

- The problem is now the same as before, except now the time-varying state variables (m_t, V_t) enter the value function. Since V_t is deterministic, it does not create a hedging demand, and we can capture its influence simply by including t as a state variable.
- One can readily verify that the value function is homogeneous in W as before, and has the functional form

$$J(W, m, t) = \frac{1}{1 - \gamma} W^{1-\gamma} H(m, t)$$

- After optimizing out C and canceling out W we get the following PDE for the unknown function $H(m, t)$

$$0 = \max_{\alpha} \left\{ H_t + [(1 - \gamma)(r + \alpha(m - r)) - \frac{1}{2}\gamma(1 - \gamma)\alpha^2\sigma^2]H + (1 - \gamma)\alpha V_t H_m + \frac{1}{2} \frac{V_t^2}{\sigma^2} H_{mm} \right\}$$

s.t. the boundary condition $H(m, T) = 1$. Note: if $\gamma = 1$ then $H(m, t) = 1 \forall t$.

- The FOC for the risky portfolio share is

$$\alpha = \frac{m_t - r}{\gamma\sigma^2} + \frac{V_t}{\gamma\sigma^2} \frac{H_m}{H}$$

- Observe that learning produces two changes in the portfolio:
 - 1 The myopic/tangency portfolio is now time-varying. The agent engages in a momentum/trend-chasing strategy, in which higher returns increase share demand because they increase the estimate of the mean return, m_t .
 - 2 The 2nd component is a hedge against **future** changes in expected returns. It reflects **estimation risk**. The sign of the hedging component is ambiguous. Because $J_m = \frac{1}{1-\gamma} W^{1-\gamma} H_m > 0$, we know $\text{sign}(H_m) = \text{sign}(1 - \gamma)$. If the agent is relatively risk averse ($\gamma > 1$), he will reduce his holdings the risky asset. In this case, changes in expected returns are negatively correlated with the marginal value of wealth. Low returns not only make you poorer and more pessimistic about the future, but they also occur when wealth is especially valuable to you.
- Note that the strength of the hedging component depends on the agent's **confidence**, as measured by V_t . If the agent lacks confidence about the mean return and so V_t is large, then demand for the risky asset can be greatly discouraged if the agent is relatively risk averse ($\gamma > 1$). However, also note that if μ is constant, $V_t \rightarrow 0$, and so the hedging component ultimately disappears. (From the law of large numbers, we also know $m_t \rightarrow \mu$, so asymptotically the portfolio converges to the Merton portfolio.)
- Results like these can perhaps explain why investors seem to invest too little in risky assets. In fact, the entire 'financial literacy' literature can be interpreted as an extension of this basic insight. Lack of knowledge discourages risk taking. (Of course, this is a partial equilibrium model. *Somebody* has to hold the risky asset, so to make this work in general equilibrium requires the introduction of some heterogeneity).

APPLICATION 2: EQUILIBRIUM ASSET PRICES

- Let's now reverse what is exogenous and what is endogenous. Following Lucas/Breeden, assume consumption/dividends following the exogenous process

$$\frac{dD_t}{D_t} = \mu dt + \sigma dB_t$$

However, unlike Lucas/Breeden, let's suppose agents don't know μ , and they must learn about it by observing the realization of dividends. They do this using the Kalman filter, exactly as before. Our task is to compute equilibrium, market-clearing, asset prices.

- Suppose markets are complete and agents have identical CRRA preferences. The unique SDF is then given by $M_t = e^{-\delta t} D_t^{-\gamma}$, and equilibrium prices are given by

$$S_t = \frac{1}{M_t} E \left[\int_t^T M_s D_s ds | \mathcal{F}_t \right] = D_t^\gamma E \left[\int_t^T e^{-\delta(s-t)} D_s^{1-\gamma} ds | \mathcal{F}_t \right]$$

- With respect to the agent's own information filtration, \mathcal{F}_t , dividends are given by

$$\log(D_s) = \log(D_t) + \left(m_t - \frac{1}{2}\sigma^2 \right) (s - t) + \sigma(\hat{B}_s - \hat{B}_t)$$

where as before we've replaced the unknown drift, μ , with its current estimate, m_t , and changed probability measures from dB to $d\hat{B}$.

- From this it follows,

$$E_t \log(D_s) = \log(D_t) + \left(m_t - \frac{1}{2}\sigma^2\right)(s - t)$$

$$\text{var}_t[\log(D_s)] = \sigma^2(s - t) + V_t(s - t)^2$$

Note that we have used the fact that m_t is a martingale under Bayesian updating, so that $E_t m_s = m_t$.

- Using standard properties of the lognormal distribution we then have

$$E_t D_s^{1-\gamma} = \exp \left\{ (1-\gamma) \left[\log(D_t) + \left(m_t - \frac{1}{2}\sigma^2\right)(s - t) \right] + \frac{1}{2}(1-\gamma)^2 \left[\sigma^2(s - t) + V_t(s - t)^2 \right] \right\}$$

- Defining $\psi_t \equiv (1 - \gamma) \left(m_t - \frac{1}{2}\gamma\sigma^2\right) - \delta$, we can then write the equilibrium price as

$$S_t = D_t F(m_t, t) = D_t \int_t^T \exp \left[\psi_t(s - t) + \frac{1}{2}(1 - \gamma)^2 V_t(s - t)^2 \right] ds$$

- Applying Ito's lemma to this we can derive a diffusion process for S ,
 $\frac{dS}{S} = \mu_{s,t}dt + \sigma_{s,t}d\hat{B}_t$, where

$$\begin{aligned}\mu_{s,t} &= m_t - \psi_t - \frac{1}{F(m_t, t)} + (1 - \gamma)V_t \frac{F_m(m_t, t)}{F(m_t, t)} \\ \sigma_{s,t} &= \sigma \left[1 + \frac{V_t}{\sigma^2} \frac{F_m(m_t, t)}{F(m_t, t)} \right]\end{aligned}$$

where

$$F_m(m_t, t) = (1 - \gamma) \int_t^T (s - t) \exp \left[\psi_t(s - t) + \frac{1}{2}(1 - \gamma)^2 V_t(s - t)^2 \right] ds$$

so that $\text{sign}(F_m) = \text{sign}(1 - \gamma)$.

Comments

- Note that when agents learn, the Price/Dividend ratio [given by the function $F(m_t, t)$] fluctuates over time, even if the true (unobserved) μ is constant. Also, because m_t mean-reverts, the P/D ratio mean reverts. These are both properties of the data that the benchmark Lucas/Breeden model cannot explain, so they represent notable successes of the learning model.

2 However, note that the *sign* of the effect of m_t on the P/D ratio is ambiguous. If $\gamma < 1$, then we get the intuitive result that optimism about dividend growth increases the P/D ratio. On the other hand, if $\gamma > 1$, then optimism about dividend growth **lowers** the P/D ratio. This ambiguity reflects the dual role of dividends in general equilibrium. Higher dividends increase the cash flows from holding the asset, so by itself, this raises the value of the asset (the 'numerator effect'). However, higher dividends raises the equilibrium discount rate (remember $M_t \sim D_t^{-\gamma}$) and so makes those increased cash flows less valuable (the 'denominator effect'). If $\gamma > 1$ the discount rate effect dominates. This is sometimes regarded as a problem, since $\gamma > 1$ is viewed as the empirically more relevant case.

3 A similar ambiguity appears in the model's predicted equity premium. The riskless rate takes the usual form

$$r_t = \delta + \gamma m_t - \frac{1}{2} \gamma (1 + \gamma) \sigma^2$$

with m_t in place of the unobserved μ . Note that it is now time-varying. Given this, we can derive the following expression for the equity premium

$$\mu_{s,t} - r_t = \gamma \sigma^2 \left[1 + \frac{V_t}{\sigma^2} \frac{F_m(m_t, t)}{F(m_t, t)} \right]$$

Since $\text{sign}(F_m) = \text{sign}(1 - \gamma)$, if $\gamma > 1$ then learning **reduces** the equity premium. Likewise, since $\sigma_{s,t} = \sigma \left[1 + \frac{V_t}{\sigma^2} \frac{F_m(m_t, t)}{F(m_t, t)} \right]$, return volatility is **lower** with learning when $\gamma > 1$.

- 4 Intuitively, when $\gamma > 1$ the denominator effect is countervailing the numerator effect in response to changes in m_t . Because stock prices are less volatile, they earn a lower risk premium.
- 5 Interestingly, the declines in $\mu_{s,t}$ and $\sigma_{s,t}$ exactly offset each other, so that the Sharpe ratio, or 'price of risk', remains unchanged in the presence of learning

$$\begin{aligned}\mu_{s,t} - r_t &= \gamma\sigma^2 \left[1 + \frac{V_t}{\sigma^2} \frac{F_m(m_t, t)}{F(m_t, t)} \right] \\ &= \gamma\sigma\sigma_{s,t} \quad \Rightarrow \quad \frac{\mu_{s,t} - r_t}{\sigma_{s,t}} = \gamma\sigma\end{aligned}$$

Remember, in continuous-time, learning only affects the conditional mean of fundamentals, not the volatility. In this representative agent economy, learning does not affect the quantity of risk the agent must bear, so the price of risk doesn't change.

- 6 Finally, note that when μ is constant, the agent eventually learns it (due to the law of large numbers). All of these effects eventually dissipate as $(m_t, V_t) \rightarrow (\mu, 0)$. Fortunately, this is easily remedied. All we need to do is assume the true μ is time-varying. For example, if we assume $d\mu_t = \sigma_\mu dB_t^\mu$, then the only change we need to make in the Kalman filter is in the update equation for V_t

$$dV_t = \left(\sigma_\mu^2 - \frac{V_t^2}{\sigma^2} \right) dt$$

and so now in the steady state $V_t \rightarrow \sigma\sigma_\mu$. The previous formulas apply if we simply modify our expression for $\text{var}_t[\log(D_s)]$

$$\text{var}_t[\log(D_s)] = \sigma^2(s-t) + V_t(s-t)^2 + \frac{1}{3}\sigma_\mu^2(s-t)^3$$

THE WONHAM FILTER

- A significant drawback of the Kalman filter is that it does not accommodate 'black swans', i.e., events that increase uncertainty. In the Kalman filter, uncertainty monotonically decreases.
- In response, researchers in macro-finance have increasingly used an alternative framework, based on noisy observations of an underlying discrete-state Markov Chain. These are sometimes called 'regime-switching' models. In continuous-time, this produces the Wonham filter.
- The Wonham filter is actually easier to work with, since we don't need to keep track of a separate conditional variance estimate. With discrete hidden states, current state probabilities capture both the conditional mean and variance.
- As with the Kalman filter we suppose we observe noisy signals of an underlying hidden state

$$dy_t = x_t dt + \sigma dB_t$$

The only difference is that now x_t can only assume one of N values. It is convenient to define the coordinate vector z_t , whose i th element equals 1 if state- i occurs at time- t and equals 0 otherwise. We can then represent x_t as $\mu \cdot z_t$, where μ is a $1 \times N$ row vector representing the values that x_t takes in each state. Note that $\hat{z}_t = E_t(z_t)$ gives the vector of conditional state probabilities.

- In continuous-time we write the state transition matrix as $\exp(\varepsilon A)$. The ‘intensity matrix’ A contains the state transition **rates**. Its diagonal elements are negative, and its rows sum to 0.
- If we approximate $\exp(\varepsilon A) \approx 1 + \varepsilon A$, and then proceed exactly as before when deriving the Kalman filter, we obtain the following diffusion process for the conditional state probabilities

$$d\hat{z}_t = A' \hat{z}_t dt + K(\hat{z})(dy_t - \mu \cdot \hat{z}_t dt)$$

where, as before, the matrix $K(\hat{z})$ is the coefficient matrix from the (local) regression of the hidden state onto the signal. It is given by

$$K(\hat{z}) = \frac{1}{\sigma^2} [\text{diag}(\hat{z}_t) - \hat{z}_t \hat{z}_t'] \mu'$$

- As in the Kalman filter, we can define the change of measure

$$d\hat{B}_t = \mu \cdot (z_t - \hat{z}_t) + dB_t$$

where $d\hat{B}_t$ is a Brownian motion with respect to the agent’s filtration. We can then write the signal process in terms of observables

$$dy_t = \mu \cdot \hat{z}_t dt + \sigma d\hat{B}_t$$

APPLICATION 1: EQUILIBRIUM ASSET PRICES

- As an application, let's revisit the Lucas/Breeden model with an unobserved dividend growth rate.

$$\frac{dD_t}{D_t} = \mu dt + \sigma dB_t$$

However, now suppose there are only two growth states, $\mu_1 > \mu_0$, so that state 1 is the 'good' state. Assume the transition rate from state 1 is λ_1 , and the transition rate from state 0 is λ_0 .

- Let z be the indicator for state 1, so that \hat{z}_t is the time- t conditional probability of μ_1 , the high-growth state. Applying the Wonham filter we get the following diffusion for \hat{z}_t

$$d\hat{z}_t = -(\lambda_0 + \lambda_1) \left[\hat{z}_t - \frac{\lambda_0}{\lambda_0 + \lambda_1} \right] dt + \frac{\mu_1 - \mu_0}{\sigma} \hat{z}_t (1 - \hat{z}_t) d\hat{B}_t$$

Note that $\lambda_0/(\lambda_0 + \lambda_1)$ is the unconditional probability of being in state 1. Hence, the drift term produces mean reversion to this long-run level. The bigger the transition rates, the stronger the mean reversion.

- The innovation $d\hat{B}_t$ is given by

$$d\hat{B}_t = \frac{1}{\sigma} \left(\frac{dD_t}{D_t} - [\mu_1 \hat{z}_t + \mu_0 (1 - \hat{z}_t)] \right)$$

so that whenever dividends grow faster than expected, you increase the probability of state 1.

- The really interesting and novel aspect of this belief revision process is that the **responsiveness** to new information depends on the agent's current state of **confidence**.
- If the agent lacks confidence about the current state, $\hat{z}_t \approx \frac{1}{2}$, and the term $\hat{z}_t(1 - \hat{z}_t)$ is relatively large. Hence, beliefs will change quickly. Conversely, if the agent is currently confident about the state, $\hat{z}_t(1 - \hat{z}_t)$ will be small, and it will take many 'surprises' before he significantly changes his beliefs.
- Since beliefs are constantly fluctuating in response to ongoing fluctuations in μ , we obtain an endogenous source of **stochastic volatility**. Learning causes volatility in the *mean* to produce volatility in the *variance*.
- As always, the stock price is just the expected value of future dividends, scaled by the SDF process

$$S_t = D_t E \left[\int_t^\infty e^{-\delta(s-t)} \left(\frac{D_s}{D_t} \right)^{1-\gamma} ds \mid \mathcal{F}_t \right]$$

Given the time-homogeneous 2-state Markov process for dividend growth, we only need to evaluate this conditional expectation for **two** values, depending on whether you're currently in state 1 or in state 0.

- Let $\pi_i = E \left[\int_t^\infty e^{-\delta(s-t)} \left(\frac{D_s}{D_t} \right)^{1-\gamma} ds \mid z_t = i \right]$. Note that it is a constant. To compute these constants we must solve two ODEs. (See if you can derive and solve them). However, we don't need to solve them to know that $\pi_1 > \pi_0$.

- Hence, we obtain the following simple expression for the Price/Dividend ratio

$$\frac{S_t}{D_t} = (1 - \hat{z}_t)\pi_0 + \hat{z}_t\pi_1$$

Note that it is time-varying, so again, learning can help us understand fluctuations in Price/Dividend ratios.

- Using our previous update equation for the state probability, $d\hat{z}_t$, we get

$$d\left(\frac{S_t}{D_t}\right) = (\pi_1 - \pi_0) \left\{ -(\lambda_0 + \lambda_1) \left[\hat{z}_t - \frac{\lambda_0}{\lambda_0 + \lambda_1} \right] dt + \frac{\mu_1 - \mu_0}{\sigma} \hat{z}_t(1 - \hat{z}_t) d\hat{B}_t \right\}$$

- Finally, since the volatility of dS/S equals the volatility of dD/D plus D/S times the volatility of $d(S/P)$, we get

$$\sigma_{s,t} = \sigma + \frac{(\pi_1 - \pi_0)(\mu_1 - \mu_0)}{\sigma} \cdot \frac{\hat{z}_t(1 - \hat{z}_t)}{(1 - \hat{z}_t)\pi_0 + \hat{z}_t\pi_1} > \sigma$$

- Hence, this simple model can account for both **excess** and **stochastic** volatility. Since $\pi_1 > \pi_0$, note that the denominator is smaller when the conditional probability of being in the low growth state is relatively high. This is consistent with the observation that volatility is higher during downturns (sometimes called the 'leverage effect').