Peering into the mist: social learning over an opaque observation network

John Barrdear[1]

[1] Bank of England and CfM

May 2014

Disclaimer: The views expressed are those of the author and do not necessarily reflect the views of the Bank of England, MPC or FPC.

Can we incorporate network learning in a macro model with dispersed information?

Network learning is a natural extension of models of incomplete information and strategic interaction:

- Firms' price-setting.
- Firms' vacancy posting.
- Households with complementarity in consumption.
- Asset pricing with communication between traders.

The problem

Three of the defining features of macroeconomic models ...

- Agents act repeatedly.
- Agents update their beliefs in a Bayesian and model-consistent way.
- Agents act strategically (payoffs are affected by others' actions).

... are precisely those that prevent comprehensive analysis of network learning.

Who's afraid of infinite state vectors?

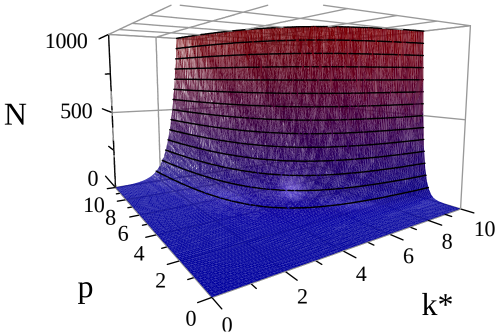

- k^*: The number of higher-order expectations to include
- p: The number of relevant *compound expectations*: linear combinations of individuals' expectations

- The dispersed info and global games literatures set $p = 1$ (the simple average) and place decreasing weight on higher-order beliefs
- For network learning, p is the number of agents
- For macro models, the number of agents is infinite

This paper (in English)

- Bayesian learning about a hidden state
- Agents
  - receive public and private signals
  - observe each others' actions over an exogenous, directed network
- Repeated, simultaneous actions
- Strategic complementarity
- Key assumption: the network is opaque

I solve for the law of motion for the full hierarchy of expectations and show that an arbitrarily accurate finite approximation may be found.

- Herding: aggregate expectations overshoot the truth
- Transitory idiosyncratic shocks have persistent aggregate effects

(A small subset of) previous literature

- Network learning
  - Dropping repeated actions: Banerjee (1992) ... Acemoglu, Dahleh, Lobel and Ozdaglar (2011)
  - Dropping Bayesian updating: DeGroot (1974) ... DeMarzo, Vayanos and Zwiebel (2003); Golub & Jackson (2010)
  - Dropping strategic concerns: Gale & Kariv (2003); Mueller-Frank (2013)

- Global games: Townsend (1983) ... Morris & Shin (2002) ...

- Dispersed information: Woodford (2003); **Nimark (2008, 2011)**; Lorenzoni (2009); Graham (2011)

- Idiosyncratic origins for aggregate volatility: Gabaix (2011); Acemoglu, Carvalho, Ozdaglar & Tahbaz-Saleh (2012)

Outline

The setup

Everything is linear

A continuum of agents, indexed $i \in [0, 1]$

The hidden *underlying state* is AR(1): $\boldsymbol{x}_t = A\boldsymbol{x}_{t-1} + P\boldsymbol{u}_t$

The *full state* includes, at a minimum, the hierarchy of simple-average expectations about the underlying state: $\overline{\boldsymbol{x}}_{t|t}^{(0:\infty)} \in X_t$

Agents' common decision rule: $g_t\left(i\right) = \boldsymbol{\lambda}_1' E_t\left(i\right)\left[X_t\right] + \boldsymbol{\lambda}_2' \boldsymbol{x}_t + \boldsymbol{\lambda}_3' \boldsymbol{v}_t\left(i\right)$

Example (Morris & Shin):

$$\begin{aligned} g_t\left(i\right) &= \left(1-\beta\right) E_t\left(i\right)\left[x_t\right] + \beta E_t\left(i\right)\left[\overline{g}_t\right] \\ &= \left(1-\beta\right)\begin{bmatrix}1 & \beta & \beta^2 & \cdots\end{bmatrix} E_t\left(i\right)\left[\overline{x}_t^{(0:\infty)}\right] \end{aligned}$$

Agents' information

Agents observe public and (conditionally independent) private signals

$$\boldsymbol{s}_t^p(i) = D_1 \boldsymbol{x}_t + D_2 X_{t-1} + R_1 \boldsymbol{v}_t(i) + R_2 \boldsymbol{e}_t + R_3 \boldsymbol{z}_{t-1}$$

$\boldsymbol{v}_t(i)$ are agent i's idiosyncratic shocks, $\boldsymbol{e}_t$ are public noise shocks and $\boldsymbol{z}_t$ are network shocks: weighted sums of idiosyncratic shocks.

Agents also observe social signals

$$\begin{aligned} \boldsymbol{s}_t^s(i) &= \boldsymbol{g}_{t-1}(\boldsymbol{\delta}_{t-1}(i)) \\ &= \boldsymbol{\lambda}_1' E_{t-1}(\delta_{t-1}(i))[X_{t-1}] + \boldsymbol{\lambda}_2' \boldsymbol{x}_{t-1} + \boldsymbol{\lambda}_3' \boldsymbol{v}_{t-1}(\delta_{t-1}(i)) \end{aligned}$$

$\delta_t(i)$ maps agent i onto their observation target(s), the period-t action of whom will be observed by i (in period $t+1$)

The network is opaque: key assumptions

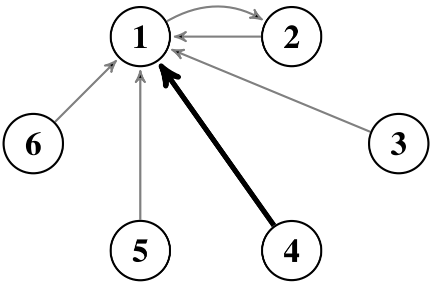

The distribution across observation targets is:

- i.i.d.
- common knowledge
- *asymptotically non-uniform*

Let Φ_n be a discrete distribution with p.m.f. $\phi_n(i)$ and let $\zeta_n \equiv \sum_{i=1}^{n} \phi_n(i)^2$ be its Herfindahl index. Φ_n is asymptotically non-uniform if:

- $\lim_{n\to\infty} \phi_n(i) = 0\ \forall i$; and
- $\lim_{n\to\infty} \zeta_n = \zeta^*$ where $\zeta^* \in (0, 1)$.

Social networks have non-uniform distributions

The degree sequences of most social networks are well approximated with a power law distribution (Jackson, 2008)

$$\phi_n(i) = c_n\, i^{-\gamma} \text{ where } \gamma > 1 \quad \Rightarrow \quad \zeta^* \in (0,1)$$

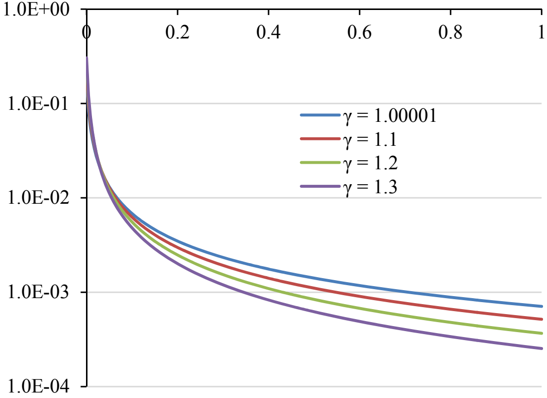

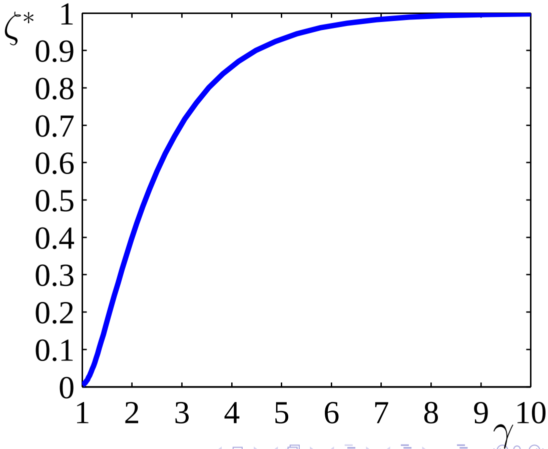

What this buys #1: a transformed problem

Linear + i.i.d. + common knowledge means:

$$
\begin{aligned}
E_t(i)\left[\mathbf{v}_t(\delta_t(i))\right] &= \int \phi(j)\, E_t(i)\left[\mathbf{v}_t(j)\right] dj \\
&= E_t(i)\left[\int \phi(j)\, \mathbf{v}_t(j)\, dj\right] \\
&= E_t(i)\left[\int \mathbf{v}_t(\delta_t(j))\, dj\right] \\
&= E_t(i)\left[\overset{1:\sim}{\mathbf{v}}_t\right]
\end{aligned}
$$

$$E_t(i)\left[\mathbf{v}_t(\delta_t(\delta_t(i)))\right] = E_t(i)\left[\overset{2:\sim}{\mathbf{v}}_t\right]$$

$$E_t(i)\left[\mathbf{v}_t(\delta_t(\delta_t(\delta_t(i))))\right] = E_t(i)\left[\overset{3:\sim}{\mathbf{v}}_t\right]$$

$$\vdots$$

What this buys #1: a transformed problem

Linear + i.i.d. + common knowledge means:

$$\begin{aligned} E_t(i)[\mathbf{v}_t(\delta_t(i))] &= \int \phi(j)\, E_t(i)[\mathbf{v}_t(j)]\, dj \\ &= E_t(i)\left[\int \phi(j)\, \mathbf{v}_t(j)\, dj\right] \\ &= E_t(i)\left[\int \mathbf{v}_t(\delta_t(j))\, dj\right] \\ &= E_t(i)\left[\overset{1:\sim}{\mathbf{v}}_t\right] \end{aligned}$$

$$\begin{aligned} E_t(i)[\mathbf{v}_t(\delta_t(\delta_t(i)))] &= E_t(i)\left[\overset{2:\sim}{\mathbf{v}}_t\right] \\ E_t(i)[\mathbf{v}_t(\delta_t(\delta_t(\delta_t(i))))] &= E_t(i)\left[\overset{3:\sim}{\mathbf{v}}_t\right] \\ &\vdots \end{aligned}$$

What this buys #2: we break the law of large numbers

An asymptotically non-uniform distribution means:

$$
\begin{aligned}
Var\left[\overset{1:\sim}{\mathbf{v}}_t\right] &= Var\left[\int \phi(j)\, \mathbf{v}_t(j)\, dj\right] \\
&= \int Var\left[\phi(j)\, \mathbf{v}_t(j)\right] dj \\
&= \int \phi(j)^2\, \Sigma_{vv} dj \\
&= \zeta^* \Sigma_{vv} \neq 0
\end{aligned}
$$

Define network shocks: $\mathbf{z}_t \equiv \begin{bmatrix} \overset{1:\sim}{\mathbf{v}}_t \\ \overset{2:\sim}{\mathbf{v}}_t \\ \vdots \end{bmatrix}$

$$Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] = \left(1 - (1 - \zeta^*)^p\right) \Sigma_{vv}$$

$$Cov\left[\overset{p:\sim}{\mathbf{v}}_t, \overset{r:\sim}{\mathbf{v}}_t\right] = Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] \quad \forall p < r$$

What this buys #2: we break the law of large numbers

An asymptotically non-uniform distribution means:

$$
\begin{aligned}
Var\left[\overset{1:\sim}{\mathbf{v}}_t\right] &= Var\left[\int \phi(j)\,\mathbf{v}_t(j)\,dj\right] \\
&= \int Var\left[\phi(j)\,\mathbf{v}_t(j)\right] dj \\
&= \int \phi(j)^2\,\Sigma_{vv} dj \\
&= \zeta^* \Sigma_{vv} \neq 0
\end{aligned}
$$

Define network shocks: $\mathbf{z}_t \equiv \begin{bmatrix} \overset{1:\sim}{\mathbf{v}}_t \\ \overset{2:\sim}{\mathbf{v}}_t \\ \vdots \end{bmatrix}$

$$
\begin{aligned}
Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] &= \left(1 - (1 - \zeta^*)^p\right) \Sigma_{vv} \\
Cov\left[\overset{p:\sim}{\mathbf{v}}_t, \overset{r:\sim}{\mathbf{v}}_t\right] &= Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] \quad \forall p < r
\end{aligned}
$$

What this buys #2: we break the law of large numbers

An asymptotically non-uniform distribution means:

$$Var\left[\overset{1:\sim}{\mathbf{v}}_t\right] = Var\left[\int \phi(j)\,\mathbf{v}_t(j)\,dj\right] = \int Var\left[\phi(j)\,\mathbf{v}_t(j)\right] dj = \int \phi(j)^2\,\Sigma_{vv} dj = \zeta^* \Sigma_{vv} \neq 0$$

Define network shocks: $\mathbf{z}_t \equiv \begin{bmatrix} \overset{1:\sim}{\mathbf{v}}_t \\ \overset{2:\sim}{\mathbf{v}}_t \\ \vdots \end{bmatrix}$

$$Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] = \left(1 - (1-\zeta^*)^p\right)\Sigma_{vv}$$

$$Cov\left[\overset{p:\sim}{\mathbf{v}}_t, \overset{r:\sim}{\mathbf{v}}_t\right] = Var\left[\overset{p:\sim}{\mathbf{v}}_t\right] \quad \forall p < r$$

The main result

The full hierarchy of expectations is defined recursively and follows an ARMA(1,1) process:

$$X_t \equiv \begin{bmatrix} \boldsymbol{x}_t \\ \overline{E}_t [X_t] \\ \overset{1:\sim}{E}_t [X_t] \\ \overset{2:\sim}{E}_t [X_t] \\ \vdots \end{bmatrix} = F X_{t-1} + G_1 \boldsymbol{u}_t + G_2 \boldsymbol{z}_t + G_3 \boldsymbol{e}_t + G_4 \boldsymbol{z}_{t-1}$$

An arbitrarily accurate approximation is obtained by defining cut-offs:

- k^*: Number of higher *orders* to include (how deep into the recursion)
- p^*: Number of higher *weights* to include (how deep into the network)

More detail

Outline

A simplified example using Morris & Shin preferences

$$g_t(i) = (1-\beta)\,E_t(i)[x_t] + \beta E_t(i)[\bar{g}_t]$$

Uni-variate state:	$x_t = \rho x_{t-1} + u_t$	$u_t \sim N(0, \sigma_u^2)$
Private signal:	$s_t^p(i) = x_t + v_t(i)$	$v_t(i) \sim N(0, \sigma_v^2)$

Result:

$$F = \begin{pmatrix} \rho & 0 & 0 & 0 & \cdots \\ B & C & D & 0 & \\ B & 0 & C & D & \\ B & 0 & 0 & C & \ddots \\ \vdots & & & & \ddots \end{pmatrix}$$

$$x_t = \rho\, x_{t-1} + u_t$$

$$\bar{E}_t[X_t] = B\,x_{t-1} + C\,\bar{E}_{t-1}[X_{t-1}] + D\,E^{1:\sim}_{t-1}[X_{t-1}] + Hu_t$$

$$E^{1:\sim}_t[X_t] = B\,x_{t-1} + C\,E^{1:\sim}_{t-1}[X_{t-1}] + D\,E^{2:\sim}_{t-1}[X_{t-1}] + Hu_t + Q\,v^{1:\sim}_t$$

$$E^{2:\sim}_t[X_t] = B\,x_{t-1} + C\,E^{2:\sim}_{t-1}[X_{t-1}] + D\,E^{3:\sim}_{t-1}[X_{t-1}] + Hu_t + Q\,v^{2:\sim}_t$$

$$\vdots$$

A simplified example using Morris & Shin preferences

$$g_t(i) = (1-\beta)\, E_t(i)[x_t] + \beta E_t(i)[\overline{g}_t]$$

Uni-variate state:	$x_t = \rho x_{t-1} + u_t$	$u_t \sim N(0, \sigma_u^2)$
Private signal:	$s_t^p(i) = x_t + v_t(i)$	$v_t(i) \sim N(0, \sigma_v^2)$

Result:

$$F = \begin{pmatrix} \rho & 0 & 0 & 0 & \cdots \\ B & C & D & 0 & \\ B & 0 & C & D & \\ B & 0 & 0 & C & \ddots \\ \vdots & & & & \ddots \end{pmatrix}$$

$$x_t = \rho\, x_{t-1} + u_t$$

$$\overline{E}_t[X_t] = B\, x_{t-1} + C\, \overline{E}_{t-1}[X_{t-1}] + D\, \overset{1:\sim}{E}_{t-1}[X_{t-1}] + H u_t$$

$$\overset{1:\sim}{E}_t[X_t] = B\, x_{t-1} + C\, \overset{1:\sim}{E}_{t-1}[X_{t-1}] + D\, \overset{2:\sim}{E}_{t-1}[X_{t-1}] + H u_t + Q\, \overset{1:\sim}{v}_t$$

$$\overset{2:\sim}{E}_t[X_t] = B\, x_{t-1} + C\, \overset{2:\sim}{E}_{t-1}[X_{t-1}] + D\, \overset{3:\sim}{E}_{t-1}[X_{t-1}] + H u_t + Q\, \overset{2:\sim}{v}_t$$

$$\vdots$$

A “true” aggregate shock #1

The hierarchy of simple-average expectations $(\overline{x}_{t|t}^{(0:\infty)})$ following a one standard deviation shock to the underlying state

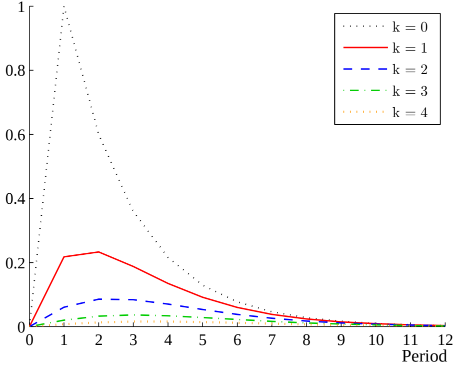

(a) Without network learning ($q = 0$)

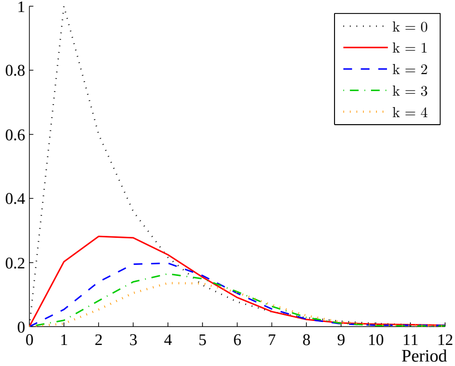

(b) With network learning ($q = 2$)

A “true” aggregate shock #2

Varying the number of other agents observed (q)

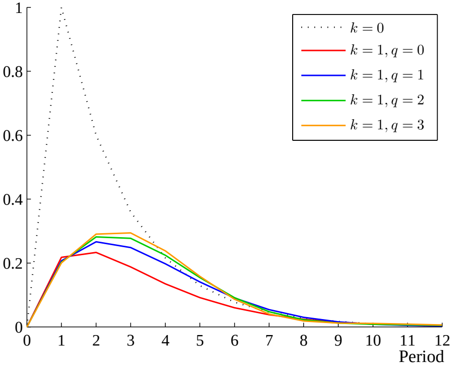

(c) Simple-average expectations

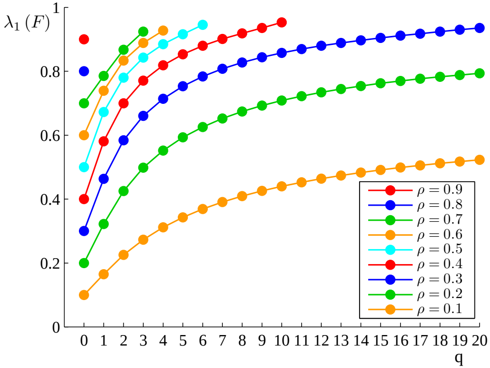

(d) Largest absolute eigenvalues of F

A network shock #1

The hierarchy of simple-average expectations ($\overline{x}_{t|t}^{(0:\infty)}$) following a one standard deviation network shock

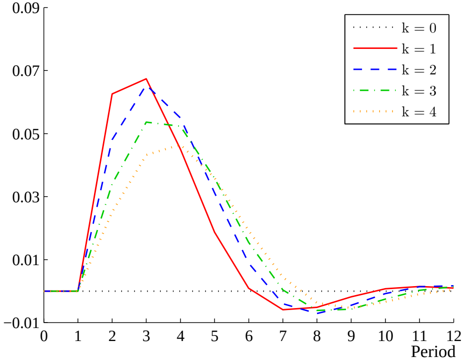

Implemented as a one standard deviation shock to $\overset{1:\sim}{\boldsymbol{v}}_t$ and the corresponding conditional expected value for higher-weighted averages with agents each observing two competitors ($q = 2$).

A network shock #2

Recall that $Var\left[\overset{1:\sim}{\boldsymbol{v}}_t\right] = \zeta^* \sigma_v^2$

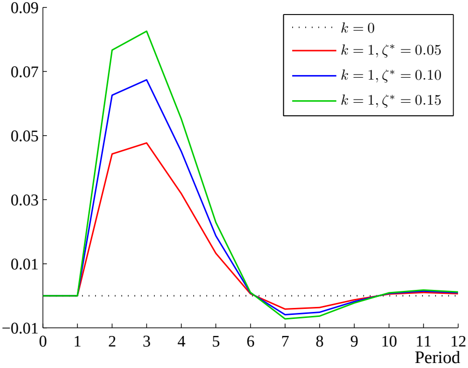

(e) Varying the degree of network irregularity (ζ^*)

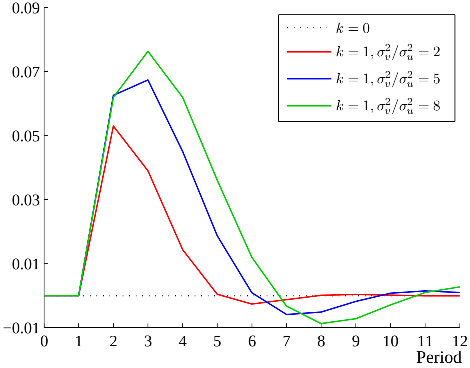

(f) Varying the relative innovation variance (σ_v^2/σ_u^2)

Adding a (lagged) public signal

$$\text{Scenario 1: } s_t^{pub} = \mathbf{1}'\overline{x}_{t-1|t-1}^{(0:\infty)} + e_t$$

$$\text{Scenario 2: } s_t^{pub} = \mathbf{1}'X_{t-1} + e_t$$

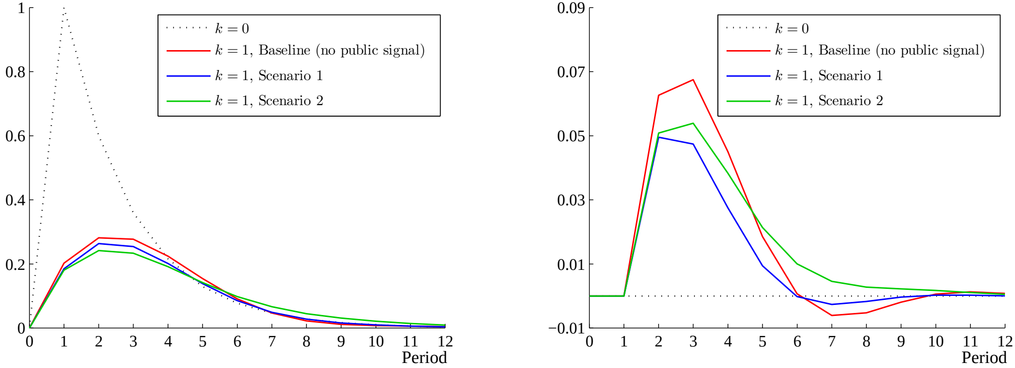

(g) A shock to the underlying state

(h) A network shock

Outline

Conclusions

- Network opacity lets us to combine (a) repeated actions; (b) rational expectations; and (c) strategic complementarity

- Underlying state follows AR(1) $\Rightarrow$ Full hierarchy follows ARMA(1,1) with $\lambda_1(F) > \lambda_1(A)$

- Herding: network learning causes aggregate beliefs to overshoot the truth following a shock to the underlying state

- Transitory idiosyncratic shocks have aggregate effects (b/c of asymptotic non-uniformity) that are persistent (b/c of recursive learning + herding)

- The model is readily nested into wider GE models of the economy

Extra slides

More detail: the Kalman filter

$$E_t(i)[X_t] = E_{t-1}(i)[X_t] + K_t \underbrace{\left(\boldsymbol{s}_t(i) - E_{t-1}(i)[\boldsymbol{s}_t(i)]\right)}_{\boldsymbol{s}^{err}_{t|t-1}(i)}$$

$$K_t = Cov(X_t, \boldsymbol{s}^{err}_{t|t-1}(i)) \left[Var\left(\boldsymbol{s}^{err}_{t|t-1}(i) \right) \right]^{-1}$$

$$\begin{aligned} \boldsymbol{s}^{err}_{t|t-1}(i) &= M_1^* X^{err}_{t-1|t-1}(i) \\ &+ N_1^* \boldsymbol{u}_t + N_2^* \boldsymbol{v}_t(i) + N_3^* \boldsymbol{e}_t \\ &= M_1 X^{err}_{t-1|t-1}(i) + M_2 X^{err}_{t-1|t-1}(\delta_{t-1}(i)) + M_3 X_{t-1} \\ &+ N_1 \boldsymbol{u}_t + N_2 \boldsymbol{v}_t(i) + N_3 \boldsymbol{e}_t + N_4 \boldsymbol{v}_{t-1}(\delta_{t-1}(i)) + N_5 \boldsymbol{z}_{t-1} \end{aligned}$$

$$V_{t|t} = Var\left[X^{err}_{t|t}(i)\right]$$

$$W_{t|t} = Cov\left[X^{err}_{t|t}(i), X^{err}_{t|t}(j)\right]$$

▸ Back