Lecture 2: Linearized Structural Macro Models

Christian Wolf
MIT

14.462, Fall 2022

Structural business-cycle models

- Structural business-cycle models (RBC, NK, HANK, ...) are mappings from

  structural shocks ε_t

  to

  macroeconomic aggregates y_t

- Today's lecture:

  1. Two ways of constructing a **linear approximation** of this mapping

     $$\left.\begin{array}{l}\text{(a) State-space approach} \\ \text{(b) Sequence-space approach}\end{array}\right\}\text{ yield } \textbf{structural vector moving average} \text{ (SVMA)}$$

  2. Map SVMA model into our **objects of interest** (IRF, FVD, HD)

- Rest of this class: how to use data to learn about coefficients of SVMA model

Running example

- Will illustrate the methods through a **running example**: neoclassical growth model

  1. Model relationships

$$
\begin{aligned}
c_t + i_t + g_t &= y_t \\
y_t &= z_t k_{t-1}^{\alpha} \\
i_t &= k_t - (1-\delta) k_{t-1} \\
q_t c_t^{-\gamma} &= \beta \mathbb{E}_t \left[q_{t+1} c_{t+1}^{-\gamma} \left[\alpha z_{t+1} k_t^{\alpha-1} + (1-\delta) \right] \right]
\end{aligned}
$$

  This is a standard NGM simplified to feature exogenous labor supply, and enriched with gov't spending and a menu of shocks. We will later also consider a richer heterogeneous-agent version.

  2. Exogenous driving forces: shocks to TFP ($\varepsilon_t^z \to z_t$), consumer demand ($\varepsilon_t^q \to q_t$) & gov't spending ($\varepsilon_t^g \to g_t$). Will usually consider simple AR(1)'s for those.

- **Objective**: characterize dynamic behavior of $(y_t, i_t, c_t, \dots)$ to first order

Outline

1. Linear State-Space Methods
   - Model Representation & Solution
   - SVMA Model and Objects of Interest
2. Linear Sequence-Space Methods
   - Model Representation & Solution
   - Heterogeneous-Agent Models & the "Fake-News" Algorithm
3. Summary

Outline

Overview

- There exist many standard treatments for the solution of linear rational expectations systems. The original reference is Blanchard & Kahn (1980).
    - I will follow the `gensys` notation of Sims (1999)
    - My discussion will be brief, as you are already supposed to be familiar with this material from first-year macroeconomics
    - Will consider a rep.-agent model. Similar methods in principle work for het.-agent models [Ahn et al. (2019)], but for those I prefer sequence-space approaches.
- Underlying computational routines: `gensys` or `dynare`
    - We will go through a solution of the running example in `dynare`
    - Exercise: replicate the same numerical solution using `gensys`
- Roadmap: **model solution** $\rightarrow$ **SVMA representation** $\rightarrow$ **objects of interest**

Outline

Linear state-space solution

The solution approach proceeds in two steps:

1. **Linearize** the model's equilibrium conditions to arrive at the form

$$\Gamma_0 \hat{x}_t = \Gamma \hat{x}_{t-1} + \Psi \varepsilon_t + \Pi \eta_t \tag{1}$$

   - $\hat{x}_t$ contains all model variables in (log-)deviation from the deterministic steady state
   - $\varepsilon_t \sim N(0, I)$ are "structural" shocks. Orthogonality by assumption, unit variance is normalization, and normality purely for convenience
   - η_t is a vector of expectational errors satisfying $\mathbb{E}_t [\eta_{t+1}] = 0$, indicating which of the equations in (1) hold only in expectation

2. Solve (1), giving a mapping from **shocks** to **macro variables** in state-space form:

$$\hat{x}_t = G_1 \hat{x}_{t-1} + \Theta \varepsilon_t \tag{2}$$

   We are interested in some particular variables y_t, given as $\hat{y}_t \equiv \tilde{C} \hat{x}_t$

1. Linearization

- Begin with the endogenous model relations. I *log*-linearize:

1. Output market-clearing
$$\bar{c}\widehat{c}_t + \bar{i}\widehat{i}_t + \bar{g}\widehat{g}_t = \bar{y}\widehat{y}_t$$

2. Production function
$$\widehat{y}_t = \widehat{z}_t + \alpha\widehat{k}_{t-1}$$

3. Investment
$$\widehat{i}_t = \frac{1}{\delta}\left[\widehat{k}_t - (1-\delta)\widehat{k}_{t-1}\right]$$

4. Euler equation
$$\widehat{q}_t - \gamma\widehat{c}_t = \mathbb{E}_t\left[\widehat{q}_{t+1} - \gamma\widehat{c}_{t+1} + [1-\beta(1-\delta)][\widehat{z}_{t+1} + (\alpha-1)\widehat{k}_t]\right]$$

- Finally we have exogenous laws of motion—say AR(1)'s—for $(\hat{z}_t, \hat{q}_t, \hat{g}_t)$

Stacking everything in the required form:

$$\underbrace{\begin{pmatrix} -\bar{c} & -\bar{i} & \bar{y} & 0 & -\bar{g} & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & -\frac{1}{\delta} & 0 & 0 & 0 \\ -\gamma & 0 & 0 & 0 & 0 & [1-\beta(1-\delta)] & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}}_{\Gamma_0} \begin{pmatrix} \hat{c}_t \\ \hat{i}_t \\ \hat{y}_t \\ \hat{k}_t \\ \hat{g}_t \\ \hat{z}_t \\ \hat{q}_t \end{pmatrix}$$

$$= \underbrace{\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{1-\delta}{\delta} & 0 & 0 & 0 \\ -\gamma & 0 & 0 & -[1-\beta(1-\delta)](\alpha-1) & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & \rho_g & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_z & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \rho_q \end{pmatrix}}_{\Gamma_1} \begin{pmatrix} \hat{c}_{t-1} \\ \hat{i}_{t-1} \\ \hat{y}_{t-1} \\ \hat{k}_{t-1} \\ \hat{g}_{t-1} \\ \hat{z}_{t-1} \\ \hat{q}_{t-1} \end{pmatrix} + \underbrace{\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ \sigma_g & 0 & 0 \\ 0 & \sigma_z & 0 \\ 0 & 0 & \sigma_q \end{pmatrix}}_{\Psi} \begin{pmatrix} \varepsilon_t^g \\ \varepsilon_t^z \\ \varepsilon_t^q \end{pmatrix} + \underbrace{\begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}}_{\Pi} \eta_t$$

2. Solution

- **Objective**: find a mapping from ε_t to x_t that is: (i) consistent with (1) and (ii) such that x_t remains bounded with probability 1 [remember the debate about this in 14.451?]
- We will recover such a solution under two addt'l assumptions, to make things easy:
  1. Γ_0 is invertible
  2. The eigenvectors of $A \equiv \Gamma_0^{-1}\Gamma_1$ are linearly independent

  By the first assumption we can re-write the system as

$$\hat{x}_t = A\hat{x}_{t-1} + \Gamma_0^{-1}(\Psi\varepsilon_t + \Pi\eta_t)$$

  Eigen-decompose A as $A = P\Lambda P^{-1}$, where the eigenvalues are ordered from smallest to largest (in absolute value). Defining $\hat{w}_t \equiv P^{-1}\hat{x}_t$, we may write

$$\hat{w}_t = \Lambda\hat{w}_{t-1} + Q(\Psi\varepsilon_t + \Pi\eta_t), \quad Q \equiv P^{-1}\Gamma_0^{-1}$$

2. Solution

- Partition Λ as

$$\Lambda = \begin{pmatrix} \Lambda_s & 0 \\ 0 & \Lambda_e \end{pmatrix}$$

  where Λ_s is diagonal with all elements smaller than 1 in absolute value, and Λ_e is diagonal with all elements greater than 1 in absolute value

- Now write the system with one stable block and one unstable block:

$$\begin{pmatrix} \hat{w}_{1,t} \\ \hat{w}_{2,t} \end{pmatrix} = \begin{pmatrix} \Lambda_s & 0 \\ 0 & \Lambda_e \end{pmatrix} \begin{pmatrix} \hat{w}_{1,t-1} \\ \hat{w}_{2,t-1} \end{pmatrix} + \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} (\Psi \varepsilon_t + \Pi \eta_t)$$

- The stability requirement will impose restrictions on the explosive block. Write it as

$$\hat{w}_{2,t} = \Lambda_e^{-1} \hat{w}_{2,t+1} - \Lambda_e^{-1} Q_2 (\Psi \varepsilon_{t+1} + \Pi \eta_{t+1})$$

  and so

$$\hat{w}_{2,t} = \lim_{T \to \infty} \Lambda_e^{-T} \hat{w}_{2,t+T} - \sum_{s=1}^{T} \Lambda_e^{-s} Q_2 (\Psi \varepsilon_{t+s} + \Pi \eta_{t+s})$$

2. Solution

- Taking expectations at time t:

$$\hat{w}_{2,t} = \lim_{T\to\infty} \Lambda_e^{-T}\mathbb{E}_t(\hat{w}_{2,t+T}) - \sum_{s=1}^{T} \Lambda_e^{-s} Q_2 \mathbb{E}_t(\Psi\varepsilon_{t+s} + \Pi\eta_{t+s}) = 0$$

- We thus must have $\hat{w}_{2,t} = 0$ and so

$$Q_2(\Psi\varepsilon_t + \Pi\eta_t) = 0$$

  The equilibrium exists and is unique if this equation uniquely defines η_t. Assume that it does, and write the solution as

$$\eta_t = -(Q_2\Pi)^{-1}Q_2\Psi\varepsilon_t$$

- Plugging everything back into the stable block:

$$\hat{w}_{1,t} = \Lambda_s\hat{w}_{1,t-1} + Q_1(\Psi - \Pi(Q_2\Pi)^{-1}Q_2\Psi)\varepsilon_t$$

2. Solution

- Collecting everything, we have

$$\begin{pmatrix} \hat{w}_{1,t} \\ \hat{w}_{2,t} \end{pmatrix} = \begin{pmatrix} \Lambda_s & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \hat{w}_{1,t-1} \\ \hat{w}_{2,t-1} \end{pmatrix} + \begin{pmatrix} Q_1(\Psi - \Pi(Q_2\Pi)^{-1}Q_2\Psi) \\ 0 \end{pmatrix} \varepsilon_t$$

  or

$$\hat{x}_t = \underbrace{P \begin{pmatrix} \Lambda_s & 0 \\ 0 & 0 \end{pmatrix} P^{-1}}_{G_1} \hat{x}_{t-1} + \underbrace{P \begin{pmatrix} Q_1(\Psi - \Pi(Q_2\Pi)^{-1}Q_2\Psi) \\ 0 \end{pmatrix}}_{\Theta} \varepsilon_t$$

- We have thus arrived at the desired model solution—a mapping from exogenous shocks ε_t to the dynamics of the macro variables $\hat{x}_t$, in the form of a VAR(1):

$$\hat{x}_t = G_1 \hat{x}_{t-1} + \Theta \varepsilon_t$$

  - Note: stability implies that all eigenvalues of G_1 are inside the unit circle
    See the posted codes for the running example.

Outline

From state-space model to SVMA

- We write the model solution in canonical **state-space** form:

$$\hat{x}_t = A\hat{x}_{t-1} + B\varepsilon_t, \quad \varepsilon_t \sim N(0, I) \tag{3}$$

$$\hat{y}_t = C\hat{x}_{t-1} + D\varepsilon_t \tag{4}$$

  - From our solution: x_t collects all variables, $A \equiv G_1$, $B \equiv \Theta$, $\hat{y}_t \equiv \tilde{C}\hat{x}_t$ is a selection of variables we are interested in, and $C \equiv \tilde{C}A$ and $D \equiv \tilde{C}B$
  - Note: I prefer the form (3) - (4) because it is general enough to also nest measurement error in (4), if desired (though we won't consider it)

- Substitute recursively, using stability of the system:

$$\hat{y}_t = D\varepsilon_t + CB\varepsilon_{t-1} + CAB\varepsilon_{t-2} + CA^2B\varepsilon_{t-3} + \cdots \equiv \sum_{\ell=0}^{\infty} \Theta_\ell \varepsilon_{t-\ell}$$

- This is a **SVMA(∞) representation**: mapping the history of shocks ε_t to y_t via the Θ's

Objects of interest

Our objects of interest are now given as simple functions of SVMA coefficients:

1. **Dynamic causal effects**

$$\text{IRF}_{i,j,h} = \Theta_{i,j,h} = \mathbb{E}(\hat{y}_{i,t+h} \mid \varepsilon_{j,t} = 1) - \mathbb{E}(\hat{y}_{i,t+h} \mid \varepsilon_{j,t} = 0), \quad h = 0, 1, 2, \ldots$$

2. **Shock importance** for average cyclical fluctuations

$$\text{FVD}_{i,j,h} \equiv 1 - \frac{\text{Var}(\hat{y}_{i,t+h} \mid \{\varepsilon_{t-\ell}\}_{\ell=0}^{\infty}, \{\varepsilon_{j,t+\ell}\}_{\ell=1}^{h})}{\text{Var}(\hat{y}_{i,t+h} \mid \{\varepsilon_{t-\ell}\}_{\ell=0}^{\infty})} = \frac{\sum_{m=0}^{h-1} \Theta_{i,j,m}^2}{\sum_{j=1}^{n_\varepsilon} \sum_{m=0}^{h-1} \Theta_{i,j,m}^2}$$

3. Contribution of shocks to particular **historical episodes**

Note: here we need the SVMA coefficients (the Θ's) plus the shocks themselves.

$$\text{HD}_{i,j,t} = \mathbb{E}\left(\hat{y}_{i,t} \mid \{\varepsilon_{j,t-\ell}\}_{\ell=0}^{\infty}\right) = \sum_{\ell=0}^{\infty} \Theta_{i,j,\ell}\varepsilon_{j,t-\ell}$$

Outline

Linear sequence-space solution

- Will now present a different solution strategy: the **sequence-space approach**
  Less established, so I will discuss it in more detail, following Auclert et al. (2021).

- Big picture comparison: **stochastic shocks** vs. **MIT shocks**

  1. So far we've considered stochastic shocks of the form

  $$\hat{z}_t = \rho_z \hat{z}_{t-1} + \sigma_z \varepsilon_t^z, \quad \varepsilon_t^z \sim N(0,1)$$

     We reduced the system to a 1st-order stoch. linear difference equation in state variables + shocks ("state-space") and recovered the stable solution through eigendecomposition.

  2. Alternatively we could consider perfect foresight ("MIT") shock paths

  $$\hat{z}_t = \rho_z^t \sigma_z$$

     and solve for linearized transition sequences ("sequence-space") back to steady state.

- **Q**: What's the difference?

Linear sequence-space solution

- Kind of obvious—but still underappreciated—result: they are the **same**!
  - Formal result: see Boppart et al. (2018) or Auclert et al. (2021)
  - Intuition preview: linearity implies certainty equivalence = perfect foresight
    - → We will see this in our running example: sequence-space approach solves the same (linearized) equations as the state-space approach, giving the same SVMA representation
- Though they give the same end result, I think there are at least **two good reasons** to add sequence-space methods to your toolkit:
  1. Sometimes it is more convenient computationally (e.g. HANK models, as you have seen in recitation with Alex)
  2. It is often easier to map model to data (both time-series and cross-sectional)
     Preview: sequence-space will be key for (i) opt. policy and (ii) X-sectional to macro aggregation.

Outline

Sequence-space model

- For sequence-space approaches, we will let boldface denote *sequences*, e.g.

$$\boldsymbol{y} = (y_0, y_1, y_2, \dots)'$$

I.e., $\boldsymbol{y}$ is the perfect-foresight transition path from $t = 0$ to ∞ given exogenous shock paths $\boldsymbol{\varepsilon}$.

- Let x_t denote a model's endogenous variables and ε_t its shocks. A **perfect-foresight equilibrium** given shock paths $\boldsymbol{\varepsilon}$ is then a set of paths $\boldsymbol{x}$ such that

$$F(\boldsymbol{x}, \boldsymbol{\varepsilon}) = \boldsymbol{0} \tag{5}$$

where $F(\bullet)$ embeds the model's equilibrium relations (Euler equation, output market-clearing, ...) In a couple of slides we will make this concrete in our running example ...

- Heuristically, we can to first order write (5) as

$$F_x \hat{\boldsymbol{x}} + F_\varepsilon \boldsymbol{\varepsilon} = 0 \tag{6}$$

(6) implicitly defines a mapping from ε's to x's, just as before.

Sequence-space model

- In fact we often times can arrive at a lower-dimensional representation of equilibria that does not involve all variables x: given $\boldsymbol{\varepsilon}$, may be able to characterize equilibria via

$$H(\boldsymbol{u}, \boldsymbol{\varepsilon}) = \mathbf{0} \tag{7}$$

where $n_u < n_x$, with x then given residually via $\boldsymbol{x} = M(\boldsymbol{u}, \boldsymbol{\varepsilon})$. To first order we can write this as

$$H_u \hat{\boldsymbol{u}} + H_\varepsilon \boldsymbol{\varepsilon} = \mathbf{0} \quad \Leftrightarrow \quad \hat{\boldsymbol{u}} = -H_u^{-1} H_\varepsilon \boldsymbol{\varepsilon}$$

- We thus recover

$$\hat{\boldsymbol{x}} = M_u \left[-H_u^{-1} H_\varepsilon \boldsymbol{\varepsilon}\right] + M_\varepsilon \boldsymbol{\varepsilon}$$

- In particular, if we are again just interested in some var's $\boldsymbol{y} = N(\boldsymbol{u}, \boldsymbol{\varepsilon})$, then

$$\hat{\boldsymbol{y}} = N_u \left[-H_u^{-1} H_\varepsilon \boldsymbol{\varepsilon}\right] + N_\varepsilon \boldsymbol{\varepsilon}$$

Let us illustrate using the simple example ...

Running example

- **Claim**: We can write the system as one set of equations (market-clearing) in one unknown u_t (which will be capital k_t). How?
    - We are given $(\boldsymbol{\varepsilon}_g, \boldsymbol{\varepsilon}_z, \boldsymbol{\varepsilon}_q)$ and so $(\boldsymbol{g}, \boldsymbol{z}, \boldsymbol{q})$. We want to write the eq'm as one equation in the unknown $\boldsymbol{k}$ given $\boldsymbol{\varepsilon}$. So suppose we knew $\boldsymbol{k}$ and then go through eq'm conditions.
    - $\boldsymbol{k}$ and $\boldsymbol{z}$ give $\boldsymbol{y}$ from the production function. We also get $\boldsymbol{i}$ from $\boldsymbol{k}$ via the def'n of i.
    - The Euler equation maps $(\boldsymbol{q}, \boldsymbol{z}, \boldsymbol{k})$ into $\boldsymbol{c}$.
    - If indeed $\boldsymbol{c} + \boldsymbol{i} + \boldsymbol{g} = \boldsymbol{y}$ then we have verified all eq'm relations

  $\Rightarrow$ This function from $\boldsymbol{k}$ and shocks $\boldsymbol{\varepsilon}$ to market-clearing is our $H(\bullet)$ function. We get the other variables in x (output, investment, consumption, ...) from k and ε– the auxiliary function $M(\bullet)$.

- Differentiating this function gives H_u. Evaluating it for $\boldsymbol{k} = \bar{\boldsymbol{k}}$ and a given shock path $\boldsymbol{\varepsilon}$ gives $H_\varepsilon \boldsymbol{\varepsilon}$. Thus we can solve for the equilibrium $\boldsymbol{k}$, and from there get $\boldsymbol{y}$, $\boldsymbol{i}$ and $\boldsymbol{c}$
- This is best seen by walking through the code ...

From sequence-space to SVMA

- Note that the model solution is immediately in the form of **shock IRFs**:

$$\Theta_{\bullet,j,\bullet} \equiv N_u \left[-H_u^{-1} H_{\varepsilon_j} e_1\right] + N_{\varepsilon_j} e_1$$

  where the ε_j subscript denotes differentiation w.r.t. shock j and $e_1 = (1, 0, 0, \dots)'$

- The perfect foresight solution has thus directly given us **SVMA coefficients**:

$$y_t = \sum_{\ell=0}^{\infty} \Theta_\ell \varepsilon_{t-\ell}$$

  The Θ's will be the same as before, as we are solving the exact same equations:

  → Recall: IRFs in stochastic model = shock $(1, 0, 0, \dots)'$ + linearized optimality conditions that hold at $t = 0$ and (in expectation) at $t = 1, 2, \dots$ + return to steady state (stability)

  → But the sequence-space approach imposes the exact same linear relations at $t = 0, 1, \dots$ + return to steady state, so you get the same numbers!

Outline

Motivation

- In our RA running example, it was easy to arrive at the eq'm characterization

$$H(\boldsymbol{u}, \boldsymbol{\varepsilon}) = \mathbf{0}$$

  and then derive the required objects $\{\mathcal{H}_u, \mathcal{H}_\varepsilon \boldsymbol{\varepsilon}\}$

- Sequence-space methods however really shine in **heterogeneous-agent models** [notably the recently popular heterogenous-household NK models]
  - In principle both the state-space and sequence-space methods reviewed here can be used to solve het.-agent models. Just the dimensionality of the system increases.
  - Sequence-space methods are however particularly appealing because the solution only require inverses of $(n_u \cdot T) \times (n_u \cdot T)$ matrices rather than Schur decompositions of $n_x \times n_x$ matrices, where n_x can end up very large [see Ahn et al. (2019) for an example]

- Will illustrate using a sequence-space algorithm for an **HA extension** of our example

HA running example

- **Consumption-savings decisions** are now made by a continuum of households i
    - Households save in an asset (capital) with return r_t, have idiosyncratic earnings risk e_{it}, and pay taxes τ_t
    - The household consumption-savings problem is to

$$\max_{a_{it}, c_{it}} \quad \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t q_t \frac{c_{it}^{1-\gamma} - 1}{1-\gamma} \right]$$

      subject to the budget constraint

$$c_{it} + a_{it} = e_{it} w_t + (1 + r_t) a_{it-1} - \tau_t, \quad a_{it} \geq \underline{a}$$

      where e_{it} is household i's (stochastic) productivity and $\int_0^1 e_{it} = 1$
    - The solution to this problem gives (complicated) consumption and savings decisions rules (previously summarized by a simple Euler equation relation)

HA running example

- **Rest of the model**
  - Optimal firm behavior gives the factor prices as
$$r_t = \alpha z_t k_{t-1}^{\alpha-1} + (1-\delta), \quad w_t = (1-\alpha) y_t$$
  - The government finances its expenditure with taxes today, so $g_t = \tau_t$
  - The asset market clears with
$$k_t = \int_0^1 a_{it} di$$
    This implies output market-clearing:
$$\int_0^1 c_{it} di + i_t + g_t = y_t$$
    Exercise: solve the RBC model version under this decentralization & verify that the solution is the same as the planning equilibrium considered so far.
- Same **objective** as before: characterize dynamic behavior of $(y_t, i_t, c_t, \ldots)$ to first order

HA running example: equilibrium representation

- **Claim**: We can again write the eq'm system as one set of equations (market-clearing) in one unknown (capital $u_t = k_t$):
    - As before we arrive at $\boldsymbol{y}$ and $\boldsymbol{i}$. We also get $(\boldsymbol{r}, \boldsymbol{w})$ from firm behavior and $\boldsymbol{\tau}$ from the government budget.
    - For each household i we get consumption $\boldsymbol{c}_i$ as a function of real rates, wages, taxes, and the demand shock. Summing over households, we get an aggregate consumption function:

$$\boldsymbol{c} = \mathcal{C}(\boldsymbol{r}, \boldsymbol{w}, \boldsymbol{\tau}, \boldsymbol{q})$$

and so, to first order,

$$\hat{\boldsymbol{c}} = \mathcal{C}_r \hat{\boldsymbol{r}} + \mathcal{C}_w \hat{\boldsymbol{w}} + \mathcal{C}_\tau \hat{\boldsymbol{\tau}} + \mathcal{C}_q \hat{\boldsymbol{q}} \tag{8}$$

Thus the only change is that (8) replaces the Euler equation

- The remaining computational challenge is how to get the $\mathcal{C}_\bullet$'s ...

Naive approach

- Need to know $\mathcal{C}_r(t,s) \equiv \frac{\partial c_t}{\partial r_s}$ for $t, s \in \{0, 1, \dots, T-1\}$ (and similarly for w, τ, q)
- Simple idea: compute $\mathcal{C}_r$ **column by column**
  - Set $\boldsymbol{r} = \bar{\boldsymbol{r}} + \varepsilon \times e_s$, where e_s is a vector of 0's and 1 at entry s, and ε is a small number
  - Iterate backward from T to get consumption and asset policies $c_t^s(e,a)$ and $a_t^s(e,a)$ everywhere on the productivity/asset state space $\mathcal{E} \times \mathcal{A}$ given prices $(\boldsymbol{r}, \bar{\boldsymbol{w}}, \bar{\boldsymbol{\tau}}, \bar{\boldsymbol{q}})$ [e.g. using endogenous gridpoint method (Carroll, 2005)]
  - Iterate forward using policies to get the distribution $D_t^s(e,a)$ over the state space
  - Compute the implied path of aggregate consumption, with

$$c_t^s = \int c_t^s(e,a) dD_t^s(e,a)$$

  - Approximate the sth column as $(c_t^s - \bar{c})/\varepsilon$
- Note that this is **costly**: needs to be done T times for each input to the $\mathcal{C}$-function

Fake-news algorithm

- The **"fake-news" algorithm**—the key contribution of Auclert et al. (2021)—exploits the fact that there is a lot of redundancy in the naive approach
  Note: my discussion here will follow their paper quite closely.

- Main appeal: it is T-**times** faster
  - Relies on a single backward and a single forward iteration, rather than T of those
  - Loosely speaking: exploits certain symmetries in agent decisions around the steady state

  This can make a difference for large HANK-type models ...

Fake-news algorithm

- Key to the algorithm is the so-called **"fake-news" matrix** $\mathcal{F}_r$:

$$\mathcal{C}_r = \begin{pmatrix} \mathcal{C}_r(0,0) & \mathcal{C}_r(0,1) & \mathcal{C}_r(0,2) & \dots \\ \mathcal{C}_r(1,0) & \mathcal{C}_r(1,1) & \mathcal{C}_r(1,2) & \dots \\ \mathcal{C}_r(2,0) & \mathcal{C}_r(1,2) & \mathcal{C}_r(2,2) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \quad \mathcal{F}_r = \begin{pmatrix} \mathcal{C}_r(0,0) & \mathcal{C}_r(0,1) & \mathcal{C}_r(0,2) & \dots \\ \mathcal{C}_r(1,0) & \mathcal{C}_r(1,1) - \mathcal{C}_r(0,0) & \mathcal{C}_r(1,2) - \mathcal{C}_r(0,1) & \dots \\ \mathcal{C}_r(2,0) & \mathcal{C}_r(1,2) - \mathcal{C}_r(1,0) & \mathcal{C}_r(2,2) - \mathcal{C}_r(1,1) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

  - First column: responses to surprise shock today
  - All further columns: shock at $s > 0$ is announced at $t = 0$, but then reversed at $t = 1$ (it was "fake news")

- The algorithm will yield $\mathcal{F}_r$ very quickly. From there we'll be able to get $\mathcal{C}_r$

The backward iteration

- Recall that $c_t^s(e,a)$ and $a_t^s(e,a)$ give time-t decision rules for a perturbation at time s
- **Claim**: we can get these decision rules from a single backward intuition
    - Key insight: only the time $s-t$ until the perturbation matters:

$$c_t^s(e,a) = \begin{cases} \bar{c}(e,a) & \text{if } s < t \\ c_{T-1-(s-t)}^{(T-1)}(e,a) & \text{if } s \geq t \end{cases}$$

      where bar denotes steady state. Why? remember first-year dynamic programming!

    - Implies: a single backward iteration from $s = T-1$ is enough to get all of the $c(\bullet)$ and $a(\bullet)$ decision rules on the state space
- We can use these decision rules to get two objects:
    1. $c_0^s = \int c_0^s(e,a) dD_0^s(e,a) = \int c_0^s(e,a) d\bar{D}(e,a)$, i.e. the first row of $\mathcal{F}_r$
    2. The time-1 distributions implied by the various time-0 decisions: $dD_1^s(e,a)$

Taking stock

- So far we already have the first row:

$$\mathcal{F}_r = \begin{pmatrix} \checkmark & \checkmark & \checkmark & \dots \\ ? & ? & ? & \dots \\ ? & ? & ? & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

- The key insight of the fake-news algorithm is that the other rows are easy:
  - Obvious: we can complete the first column by just iterating the distribution forward from $dD_1^0(e, a)$, using steady-state decision rules
  - But crucially: because the initial news are reversed at $t = 1$, we can similarly get all other columns by iterating forward the respective $dD_1^s(e, a)$ using steady-state decision rules
- The forward iteration step formalizes this ...

The forward iteration

- Let $\bar{\Lambda}$ denote the **transition matrix** across the state space $\mathcal{E} \times \mathcal{A}$ using steady-state decision rules. Forward iteration gives

$$d\tilde{D}_t^s(e,a) = \bar{\Lambda} \times d\tilde{D}_{t-1}^s(e,a)$$

as the "fake-news" distributions. We can start this using $d\tilde{D}_1^s(e,a) = dD_1^s(e,a)$.

- Using this, we can get

$$\tilde{c}_t^s \equiv \int \bar{c}(e,a) d\tilde{D}_t^s(e,a) = \bar{c}'(\bar{\Lambda}')^{t-1} dD_1^s(e,a)$$

Finite differences thus give us all the entries of $\mathcal{F}_r$ via $(\tilde{c}_t^s - \bar{c})/\varepsilon$, and so $\mathcal{C}_r$.

  - Note that this requires just one distributional forward iteration using steady-state consumption decision rules

Closed-form example

- To provide a different perspective on the intuition let's consider a "one-agent" example, allowing us to ignore the distributional component and thus giving closed-form solutions
- Environment: perpetual-youth OLG [Blanchard (1985)]
    - In that case can show that the first column of the MPC matrix $\mathcal{C}_\tau$ satisfies

$$\mathcal{C}_\tau^{(}\bullet, 1) = \underbrace{\left(1 - \frac{\theta}{1+\bar{r}}\right)}_{\text{MPC}} \times \underbrace{\left\{1, \theta, \theta^2, \dots\right\}'}_{\text{spending decay}}$$

while the first row is given as

$$\mathcal{C}_\tau^{(}1, \bullet) = \underbrace{\left(1 - \frac{\theta}{1+\bar{r}}\right)}_{\text{MPC}} \times \underbrace{\left\{1, \frac{\theta}{1+\bar{r}}, \left(\frac{\theta}{1+\bar{r}}\right)^2, \dots\right\}}_{\text{anticipation effects}}$$

    - We can use the logic of the fake news algorithm to complete the matrix

Closed-form example

- What can we say about the second column?
  - Fake-news logic: must be equal to (i) news shock that is reversed at date 1 + (ii) a new surprise shock at date 1
  - But we can write this as

$$\mathcal{C}_\tau(\bullet, 2) = \underbrace{\mathcal{C}_\tau(1,2) \times \begin{pmatrix} 1 \\ -\mathcal{C}_\tau(\bullet, 1)(1+\bar{r}) \end{pmatrix}}_{\text{response to fake news}} + \underbrace{\begin{pmatrix} 0 \\ \mathcal{C}_\tau(\bullet, 1) \end{pmatrix}}_{\text{reverse the fake news}}$$

Simpler than HANK: no distributional considerations needed, just one agent whose asset holdings have declined by $-\mathcal{C}_\tau(1,2) \times (1+\bar{r})$.

- All other columns then follow recursively:

$$\mathcal{C}_\tau(\bullet, h) = \mathcal{C}_\tau(1,h) \times \begin{pmatrix} 1 \\ -\mathcal{C}_\tau(\bullet, 1)(1+\bar{r}) \end{pmatrix} + \begin{pmatrix} 0 \\ \mathcal{C}_\tau(\bullet, h-1) \end{pmatrix}, \quad h = 3, 4, \ldots$$

Outline

Summary

- In line with the **impulse-propagation framework**, modern structural business-cycle models are mappings from exogenous shocks ε_t to macro outcomes y_t

- Today we saw two **numerical techniques** to linearly approximate this mapping:

  1. State-space methods (= first-order perturbation)
  2. Sequence-space methods (= linearized perfect foresight)

  Both yield the same **SVMA(∞) representation** of the mapping $\varepsilon_t \to y_t$

- Rest of the class: how to **use data** to learn about the Θ's of the VMA mapping