Lecture 9: Cross-Sectional Analysis – Interpretation

Christian Wolf

MIT

14.461, Fall 2022

Overview

- There's been a recent push towards **micro data for macro**. Examples:
  - Consumer spending responses to stimulus payments (Great Recession, Covid)
    Johnson-Parker-Souleles (2006), Fagereng-Holm-Natvik (2020), ...
  - Labor supply responses to UI benefits
    Ganong-Greig-Liebeskind-Noel-Sullivan-Vavra (2021)
  - Firm responses to investment stimulus/easier financing conditions
    Zwick-Mahon (2017), Ottonello-Winberry (2020)
- We'll try to connect this empirical work to our **SVMA framework**

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

1. What does micro data identify? **A**: key inputs to the Θ's, not the Θ's themselves [today]
2. How can we go from micro data to the Θ's? [next time]

Outline

1. Background: Two Models in Sequence-Space
   - Intertemporal Keynesian Cross
   - Neoclassical Growth Model
2. Cross-Individual Analysis
   - Household MPCs
   - Firm MPIs
3. Cross-Regional Analysis

Outline

1. Background: Two Models in Sequence-Space
   - Intertemporal Keynesian Cross
   - Neoclassical Growth Model

2. Cross-Individual Analysis
   - Household MPCs
   - Firm MPIs

3. Cross-Regional Analysis

Micro data & macro models

- Will see: **sequence-space methods** allow us to formalize the connection between

  the econometric estimands of cross-sectional methods

  and

  our macroeconomic causal effects of interest (the Θ's)

- Preview of **big-picture intuition**
  - Can represent linearized macro models as 1. (dynamic) supply & demand curves (e.g., inv. demand, labor supply, ...) + 2. policy rules + 3. shocks shifting supply/demand/policy
  - The Θ's come from solving a linear system in 1., 2., and 3. We'll show that that (idealized) cross-sectional experiments identify the slopes of 1.

- We'll illustrate this through **two examples**: (i) an NK model & (ii) a neoclassical model

Outline

1. Background: Two Models in Sequence-Space
    - Intertemporal Keynesian Cross
    - Neoclassical Growth Model
2. Cross-Individual Analysis
    - Household MPCs
    - Firm MPIs
3. Cross-Regional Analysis

The Intertemporal Keynesian Cross

- The first model is a version of the **intertemporal Keynesian cross** [Auclert et al. (2018)]
  Overview: standard NK model + HA block. Look at linearized transition paths = "sequence-space".

- **Model details**

  1. Households

     → A unit continuum of households solves a standard consumption-savings problem:

     $$\max_{\{c_{it}\}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_{it}) \right]$$

     subject to

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

     where e_{it} denotes individual i's (stochastic) productivity, with $\int e_{it} di = 1$

     → Aggregating across households, we get an aggregate consumption function:

     $$\boldsymbol{c} = \mathcal{C}(\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}), \quad \Rightarrow \quad \widehat{\boldsymbol{c}} = \mathcal{C}_y \widehat{\boldsymbol{y}} + \mathcal{C}_r \widehat{\boldsymbol{r}} + \mathcal{C}_\tau \widehat{\boldsymbol{\tau}}$$

The Intertemporal Keynesian Cross

- **Model details**

2. Firms & unions

   → Household labor supply is intermediated by sticky-prices and unions

   → Optimal price-setting gives rise to an aggregate NKPC: [I am skipping many steps here.]

$$\widehat{\pi}_t = \kappa \widehat{y}_t + \beta \widehat{\pi}_{t+1}$$

3. Government

   → The government pays out uniform transfers τ_t (or imposes taxes, if negative), consumes (g_t), and issues debt (b_t). Its budget constraint is

$$b_t = \frac{1 + i_{t-1}}{1 + \pi_t} b_{t-1} + g_t + \tau_t$$

   → Assume the monetary authority fixes the real rate of interest (i.e., $1 + r_t = \frac{1+i_{t-1}}{1+\pi_t} = 1 + \bar{r}$). A fiscal policy is then simply paths $(\boldsymbol{g}, \boldsymbol{\tau})$ such that

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}} \right)^t \widehat{g}_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}} \right)^t \widehat{\tau}_t$$

The Intertemporal Keynesian Cross

- **Equilibrium characterization** is now straightforward:
  - Given a fiscal policy $(\boldsymbol{g}, \boldsymbol{\tau})$, an equilibrium is a path of aggregate output $\boldsymbol{y}$ s.t.

$$\mathcal{C}(\boldsymbol{y}, \bar{\boldsymbol{r}}, \boldsymbol{\tau}) + \boldsymbol{g} = \boldsymbol{y}$$

    Approximating this to first order:

$$\mathcal{C}_y \widehat{\boldsymbol{y}} + \mathcal{C}_\tau \widehat{\boldsymbol{\tau}} + \widehat{\boldsymbol{g}} = \widehat{\boldsymbol{y}}$$

  - Why is this enough? output market and so asset market clear, hh's behave optimally, π_t can be recovered residually from the NKPC (and i_t then set to keep r fixed)

Note: it depends a bit on the details of the asset structure whether $r_0 = \bar{r}$ is actually implementable. I will ignore those details because they are orthogonal to our focus here.

The Intertemporal Keynesian Cross

- Now let's solve from here for a **VMA representation** for fiscal shocks
  - Consider some fiscal rule $\widehat{\boldsymbol{\tau}} = \widehat{\boldsymbol{\tau}}(\boldsymbol{\varepsilon}_g, \boldsymbol{\varepsilon}_\tau)$ and $\widehat{\boldsymbol{g}} = \boldsymbol{\varepsilon}_g$, i.e., two fiscal shocks. Then:

$$\widehat{\boldsymbol{y}} = (I - \mathcal{C}_y)^{-1} \times (\mathcal{C}_\tau \widehat{\boldsymbol{\tau}}(\boldsymbol{\varepsilon}_g, \boldsymbol{\varepsilon}_\tau) + \boldsymbol{\varepsilon}_g) \equiv \Theta_{y,g} \times \boldsymbol{\varepsilon}_g + \Theta_{y,\tau} \times \boldsymbol{\varepsilon}_\tau$$

  Note: $(I - \mathcal{C}_y)^{-1}$ is heuristic—this map is actually not invertible. See Auclert et al. for details.

- What have we learned from this **equilibrium representation**?
  - Mapping from fiscal spending shocks to output depends only on 1. $\mathcal{C}_y$ and $\mathcal{C}_\tau$ & 2. the fiscal rule $\widehat{\boldsymbol{\tau}}(\bullet)$—exactly my claim from before
  - Note that $\mathcal{C}_y$ and $\mathcal{C}_\tau$ are "slopes" of an aggregate consumption function—i.e., how additional income/taxes/transfers map into consumer spending

$\Rightarrow$ Where we're going: cross-sectional experiments give entries of those $\mathcal{C}_\bullet$ matrices

Outline

1. Background: Two Models in Sequence-Space
    - Intertemporal Keynesian Cross
    - Neoclassical Growth Model
2. Cross-Individual Analysis
    - Household MPCs
    - Firm MPIs
3. Cross-Regional Analysis

A simple neoclassical model

- The second model is a version of the standard **neoclassical growth model**

  Overview: households supply labor and save in capital, firms rent capital and hire labor. Will again look at linearized perfect-foresight transition paths, now after a TFP shock.

- **Model details**

  1. Households

     → Households face sequences of wages $\boldsymbol{w}$, real interest rates $\boldsymbol{r}$ and firm dividends $\boldsymbol{d}$, and decide how much to consume, save, and work

     → Rather than specifying the decision problem in detail [see e.g. Lecture Note 2], I will just summarize its solution through the implied optimal decision rules:

$$\begin{aligned}
\boldsymbol{c} &= \boldsymbol{c}(\boldsymbol{w}, \boldsymbol{r}, \boldsymbol{d}) \\
\boldsymbol{\ell}^h &= \boldsymbol{\ell}^h(\boldsymbol{w}, \boldsymbol{r}, \boldsymbol{d}) \\
\boldsymbol{k}^h &= \boldsymbol{k}^h(\boldsymbol{w}, \boldsymbol{r}, \boldsymbol{d})
\end{aligned}$$

A simple neoclassical model

- **Model details**

2. Firms

→ Firms also face wages $\boldsymbol{w}$ and real interest rates $\boldsymbol{r}$. They decide how much output to produce, capital to rent, labor to hire, and dividends to pay out, subject to productivity shocks $\boldsymbol{z}$.

→ I again summarize the solution through the implied optimal decision rules:

$$\begin{aligned} \boldsymbol{y} &= \boldsymbol{y}(\boldsymbol{\ell}^f, \boldsymbol{k}^f; \boldsymbol{z}) \\ \boldsymbol{\ell}^f &= \boldsymbol{\ell}^f(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}) \\ \boldsymbol{k}^f &= \boldsymbol{k}^f(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}) \\ \boldsymbol{d} &= \boldsymbol{d}(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}) \end{aligned}$$

→ Will also be useful to define investment $\boldsymbol{i} = \boldsymbol{k} - (1-\delta)\boldsymbol{k}_{-1}$ and its corresponding decision rule $\boldsymbol{i} = \boldsymbol{i}(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z})$

A simple neoclassical model

- **Equilibrium characterization** is again straightforward:
  - Given exogenous TFP $\boldsymbol{z}$, An equilibrium is now a pair $(\boldsymbol{w}, \boldsymbol{r})$ such that the output and labor markets clear:

$$
\begin{aligned}
\boldsymbol{c}(\boldsymbol{w}, \boldsymbol{r}, \boldsymbol{d}(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z})) + \boldsymbol{i}(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}) &= \boldsymbol{y}(\boldsymbol{\ell}^f(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}), \boldsymbol{k}^f(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z}); \boldsymbol{z}) \\
\boldsymbol{\ell}^h(\boldsymbol{w}, \boldsymbol{r}, \boldsymbol{d}(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z})) &= \boldsymbol{\ell}^f(\boldsymbol{w}, \boldsymbol{r}; \boldsymbol{z})
\end{aligned}
$$

  - As before, we can take a first-order expansion of this relation to arrive at a big matrix system of the form:

$$
A \times \begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{r} \end{pmatrix} = B\boldsymbol{z}
$$

    The solution gives mappings from $\boldsymbol{z}$ to $(\boldsymbol{w}, \boldsymbol{r})$ and so all other eq'm aggregates (our Θ's)
  - A and B again collect the slopes of various supply and demand functions & $\boldsymbol{z}$ is the shock

$\Rightarrow$ Will again see: cross-sectional experiments give entries of A & B, e.g. $\mathcal{I}_r \equiv \frac{\partial \boldsymbol{i}(\bullet)}{\partial \boldsymbol{r}}$

Outline

1. Background: Two Models in Sequence-Space
    - Intertemporal Keynesian Cross
    - Neoclassical Growth Model

2. Cross-Individual Analysis
    - Household MPCs
    - Firm MPIs

3. Cross-Regional Analysis

The estimand of cross-individual regressions

- **Main result**: cross-sectional micro regressions recover entries of PE elasticity matrices
  - More precisely: cross-*individual* (e.g., across-household, across-firm) regressions give such slopes—need the estimates to be free of any (local or aggregate) GE interactions
  - The same is thus *not* true for cross-regional regressions. Later.
- I will illustrate through two examples, linked to the two models we saw:
  1. **Lottery wins** identify parts of $\mathcal{C}_\tau$
     Fagereng-Holm-Natvik (2020)
  2. **Investment tax write-offs** identify parts of $\mathcal{I}_r$
     Zwick-Mahon (2017)

Outline

Fagereng et al. (2020)

- Econometric set-up of **Fagereng et al. (2020)**
  Applies similarly to Johnson-Parker-Souleles (2008), Parker-Johnson-Souleles-McClelland (2013)
  - The authors observe comprehensive data on (random) lottery wins, consumption, and wealth holdings of Norwegian households
  - Using their data, they estimate micro regressions of the following form:

$$c_{it} = \alpha_i + \delta_t + \sum_{t=0}^{H} \gamma_k \text{lottery}_{i,t-k} + \theta X_{it} + \varepsilon_{it}$$

- Interpretation: this teaches us about entries of the ***first column* of $\mathcal{C}_\tau$ and $\mathcal{C}_y$**
  - $\mathcal{C}_\tau$: for uniform \$ transfers τ, the γ_k's give the first column of the response matrix
  - $\mathcal{C}_y$: for Keynesian GE effects, MPCs need to be weighted by GE incidence
    Auclert et al. (2018) consider the unit-elasticity case, resulting in MPCs weighted by pre-tax income. See Patterson (2021) for more general estimates.

Implication: testing consumption models

Figure 2: iMPCs in the Norwegian data and several models.

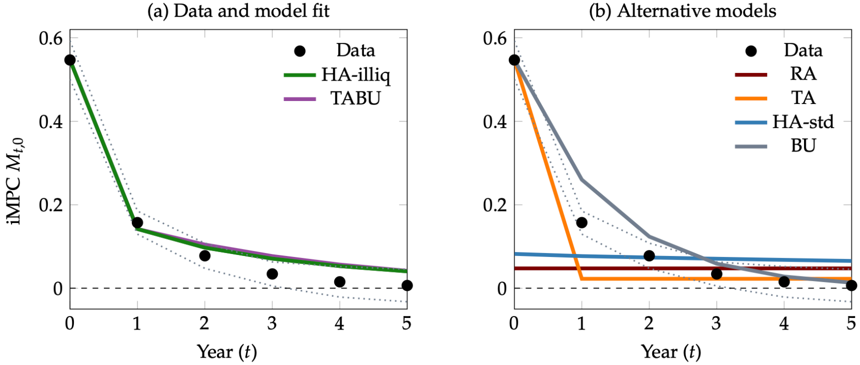

Auclert et al. (2018): can reject permanent-income and spender-saver models. But note: much less evidence on higher-order columns of $\mathcal{C}_\tau/\mathcal{C}_y$

Outline

Zwick & Mahon (2017)

- Econometric set-up of **Zwick-Mahon (2017)**
  - Background: "bonus depreciation" is a policy that allows firms to write off long-lived capital goods at a faster rate, thus resulting in greater tax savings today
    Can show: without financial frictions this is equivalent to cut in cost of capital (Winberry, 2018).
  - The authors exploit differences across firms in their exposure to this policy (due to differences in how long-lived their capital goods are). They run

$$\log(i_{jt}) = \alpha_j + \delta_t + \beta_{ZM} \times b_{jt} + \text{controls} + \text{error}$$

    where b_{jt} is the firm-specific exposure variable
- Under some assumptions this **teaches us about $\mathcal{I}_r$**
  - Let β denote the average contemporaneous interest rate elasticity of investment, i.e. entry (1,1) of the PE elasticity matrix $\mathcal{I}_r$.
  - Koby-Wolf (2020) give (strong) conditions under which the estimand β_{ZM} satisfies $\beta_{ZM} = \beta$

Implication: testing investment models

As it turns out, different models of investment *really* differ on $\mathcal{I}_r(1,1)$:

INVESTMENT ELASTICITIES: KHAN & THOMAS (2008) VS. WINBERRY (2018)

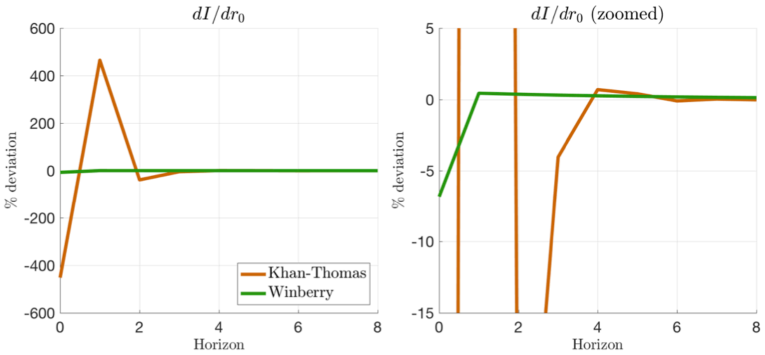

Micro evidence selects small elasticities. Can show that the level of the elasticity matters for various GE IRFs (i.e., our Θ's). [see Koby-Wolf (2020) for details, not our focus here]

Summary

- In neither experiment did the **micro data** answer the **macro-relevant question**
  - What's the response of aggregate consumption to **stimulus checks**?
  - What's the response of aggregate investment to **bonus depreciation**?

  We can't tell from the micro experiment alone ...

- Rather, we got useful *inputs*—not $\Theta_{c,\tau}$ or $\Theta_{i,b}$, but $\mathcal{C}_\tau$ or $\mathcal{I}_r$

  Though note: we don't even get the full inputs (which are infinite-dimensional), but some entries

  - My "application" slides reflect this: I used evidence on $\mathcal{C}_\tau$ or $\mathcal{I}_r$ to test models, but *not* to make statements about the aggregate effects of stimulus checks/bonus depreciation

Q for next time: how can we get from PE elasticities to GE counterfactuals?

Outline

1. Background: Two Models in Sequence-Space
   - Intertemporal Keynesian Cross
   - Neoclassical Growth Model

2. Cross-Individual Analysis
   - Household MPCs
   - Firm MPIs

3. Cross-Regional Analysis

Cross-regional analysis

- So far we focussed on cross-individual micro analysis. Natural **Q**: what about **cross-regional** variation?
- I'll try to make two main points here:
  1. Cross-regional variation identifies neither aggregate effects nor slopes of PE functions, but something in between: it gives local causal effects that contain local GE forces.
  2. Local causal effects may be largely driven by the strength of regional spillover effects, which in turn can be completely irrelevant for aggregate causal effects.
- My entire analysis will focus on a simple but (I think) instructive example.
  For very useful complementary perspectives see Nakamura-Steinsson (2014) and Chodorow-Reich (2019).

A multi-region Keynesian Cross

- I consider a multi-region version of the **IKC model** studied before
  - There is now a unit continuum of regions i. Each region buys its own good and the goods of all other regions, with home bias λ
  - To make things particularly simple I will assume that all prices are fully rigid, so output is demand-determined everywhere
- **Q**: What are the effects of aggregate and local fiscal spending shocks?

1. Aggregate fiscal multiplier

   → Since all regions are symmetric the analysis is exactly as in the baseline case, requiring only the aggregate Keynesian cross

   → In particular, aggregate output in response to a fiscal expansion $\widehat{\boldsymbol{g}}$ with financing $\widehat{\boldsymbol{\tau}}$ satisfies

$$\widehat{\boldsymbol{y}} = (I - \mathcal{C}_y)^{-1} \times (\widehat{\boldsymbol{g}} + \mathcal{C}_\tau \widehat{\boldsymbol{\tau}})$$

A multi-region Keynesian Cross

- **Q**: What are the effects of aggregate and local fiscal spending shocks?

  2. Regional fiscal multiplier

     → Next I consider a spending expansion in region i, financed with taxes imposed on every region

     → Since region i is infinitesimal, its relative output $\widehat{\boldsymbol{y}}_i$ solves the regional Keynesian cross

$$\widehat{\boldsymbol{y}}_i = \lambda \mathcal{C}_y \widehat{\boldsymbol{y}}_i + \widehat{\boldsymbol{g}}_i$$

     → Solving, we find

$$\widehat{\boldsymbol{y}}_i = (I - \lambda \mathcal{C}_y)^{-1} \times \widehat{\boldsymbol{g}}_i$$

     This establishes the first point: the regional multiplier does not yield PE slopes, but is shaped by local GE. Aside: if $\lambda \in (0,1)$ then this notation is not heuristic—$I - \lambda \mathcal{C}_y$ is invertible.

     → In particular, relative to aggregate multipliers, the regional multiplier is independent of financing effects ($\boldsymbol{\tau}$), but on the other hand shaped by the strength of spillovers (λ).

From regional to aggregate multipliers

What's the relationship between regional and aggregate multipliers?

- Main insight: they can sometimes be **completely orthogonal objects**

From regional to aggregate multipliers

What's the relationship between regional and aggregate multipliers?

- Main insight: they can sometimes be **completely orthogonal objects**
- Let's begin with the **regional multiplier**
  - Let $\boldsymbol{r} \equiv (1, \frac{1}{1+\bar{r}}, \dots)$, so that $\boldsymbol{r}'\boldsymbol{x}$ gives the present value of some vector $\boldsymbol{x}$, and $\boldsymbol{r}'\widehat{\boldsymbol{y}}/\boldsymbol{r}'\widehat{\boldsymbol{g}}$ gives cumulative fiscal multipliers
  - Since necessarily $\boldsymbol{r}'\mathcal{C}_y\widehat{\boldsymbol{y}}_i = \boldsymbol{r}'\widehat{\boldsymbol{y}}_i$ (by the household budget constraint), it follows that

$$\boldsymbol{r}'\widehat{\boldsymbol{y}}_i/\boldsymbol{r}'\widehat{\boldsymbol{g}}_i = \frac{1}{1-\lambda}$$

    That is, the cumulative regional multiplier depends *only* on spillovers λ

From regional to aggregate multipliers

What's the relationship between regional and aggregate multipliers?

- Main insight: they can sometimes be **completely orthogonal objects**
- Now let's consider the **aggregate multiplier**
    - Consider first the case of tax-financed spending increase, $\widehat{\boldsymbol{g}} = \widehat{\boldsymbol{\tau}}$. Here for $\mathcal{C}_y = \mathcal{C}_\tau$ the cumulative multiplier is equal to 1, so below the regional multiplier.
    - Now return to the general deficit-financed case (i.e., $\widehat{\boldsymbol{g}} \neq \widehat{\boldsymbol{\tau}}$). Here, depending on $\mathcal{C}_y$ & $\mathcal{C}_\tau$, cumulative multipliers can be $\gg 1$ [see next slide for an example]

From regional to aggregate multipliers: HANK model example

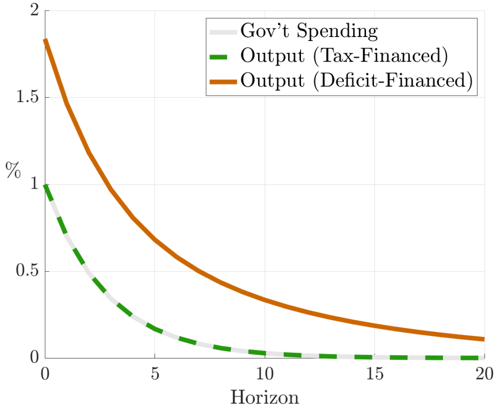

Takeaways: cumulative regional multipliers depend only on λ, while aggregate (cumulative) multipliers depend on $\mathcal{C}_y$, $\mathcal{C}_\tau$ and tax financing. Could have regional $>$ agg. or vice-versa.

Summary

- **Main takeaways**
    - Formalized via sequence-space methods: micro experiments give inputs to our Θ's—slopes of supply & demand functions—not the Θ's themselves
    - Additional challenge for regional variation: also capture regional GE (which may reflect very different mechanisms from aggregate GE)
- **Next time**: how can we go from these slopes to our Θ's?