Lecture 8: The Origins of Business Cycles

Christian Wolf
MIT

14.461, Fall 2022

Overview

- We'll now turn to our third question: what are the **origins of business cycles**?
- Again proceed in two steps:
  a) Review purely **semi-structural** literature on shock identification
     Can we identify any shocks that are credible main sources of cyclical fluctuations?
  b) How can time-series moments inform **structural business-cycle modeling**?
     (i) Business-cycle anatomy Angeletos et al. (2021)
     (ii) RANK/HANK model estimation using likelihood-based methods Justiniano et al. (2010), Auclert et al. (2021)

Outline

1. Semi-Structural Methods
   - TFP Shocks
   - Investment Technology Shocks

2. Structural Model Estimation
   - Basics of Likelihood-Based Estimation
   - Business-Cycle Anatomy
   - RANK & HANK Model Estimation

Outline

1. Semi-Structural Methods
   - TFP Shocks
   - Investment Technology Shocks

2. Structural Model Estimation
   - Basics of Likelihood-Based Estimation
   - Business-Cycle Anatomy
   - RANK & HANK Model Estimation

Background

- For better or worse, semi-structural analysis of business-cycle origins has been largely focused on "**technology shocks**"
- Background is the history of structural macro modeling, notably Kydland-Prescott (1982)
  - Classical RBC business-cycle analysis is built around the aggregate production function

$$y_t = a_t f(k_t, \ell_t)$$

  - Core finding: in RBC models, shocks to the exogenous process driving a_t can generate what looks like typical aggregate business cycles
    Aside: this celebrated finding somewhat rests on a very high Frisch elasticity that is simply implied by the chosen preference specification. We can chat more if you're interested.
- This created an entire **empirical research agenda**: Can we identify "technology shocks"? If so, do they drive cyclical fluctuations, in our FVD/FVR sense?

Outline

Identification challenge

Q: How should we go about identifying "technology shocks"?

- **Early attempts**
  - The early literature simply used Solow residuals as a direct measure of technology shocks, e.g. Prescott (1986)
  - This of course has a host of problems: assumes specific production function and competition, doesn't allow for variations in capacity utilization, ...
- Literature thus moved ahead quickly. Will review **two more refined approaches**
  1. Galí (1999): VAR in labor productivity and employment, assume only tech. shocks can have long-run effects on productivity. What are the problems with this set-up?
  2. Basu et al. (2006): manually adjust Solow residual for time-varying utilization rates

Note: similar results are also reported in the recent max-share contribution of Francis et al. (2014).

Results

Galí (1999)

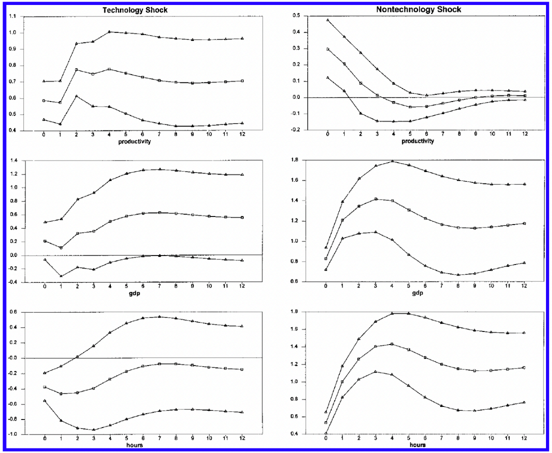

Figure 2. Estimated Impulse Responses from a Bivariate Model: U.S. Data, First-Differenced Hours (Point Estimates and ±2 Standard Error Confidence Intervals)

Results

Basu et al. (2006)

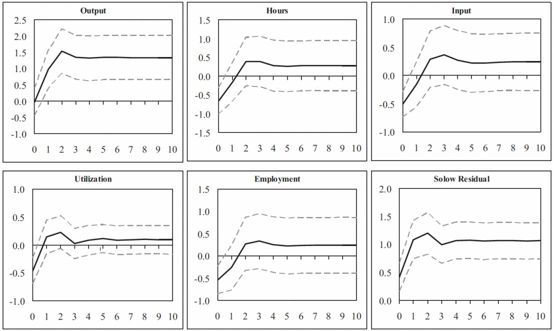

Results

Basu et al. (2006)

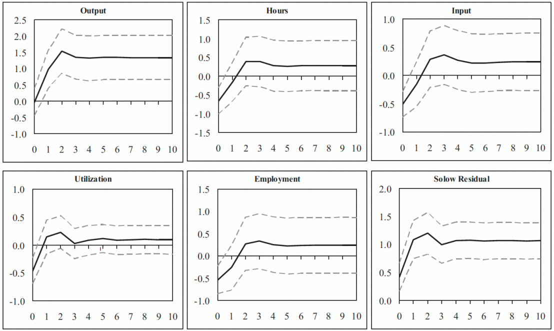

Interpretation: consistent with NK, not RBC. Implies little role for technology shocks as a source of cyclical fluctuations (since IRFs don't look like business cycles).

Technology news

- If not contemporaneous technology shocks, then maybe **news shocks**?
  Basic idea goes back to Beveridge (1909) and Pigou (1927).
    - RBC model challenge: positive news about the future make households wealthier, leading to a decline in labor supply, leading to a recession today
    - But can make it work with a couple of twists, chiefly requiring a weak wealth effect in labor supply See Jaimovich and Rebelo (2009) for the details.
- What can we see in the **data**?
    1. Beaudry-Portier (2006): bivariate VAR in stock prices and TFP, find that the same shock drives both (i) short-run stock price changes $\perp$ TFP and (ii) long-run TFP
       Important concern is non-invertibility. See Forni et al. (2014) for a discussion.
    2. Barsky et al. (2014): maximize TFP contribution at medium horizons and include larger information set $\rightarrow$ alleviates non-invertibility concerns with Beaudry-Portier

Results

Beaudry-Portier (2006)

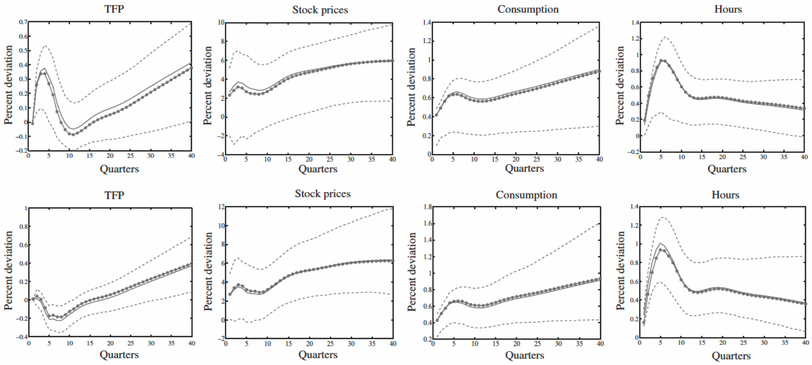

Figure 9. Impulse Responses to ε_2 and $\tilde{\varepsilon}_1$ in the (*TFP*, *SP*, *C*, *H*) VECM, without (upper panels) or with (lower panels) Adjusting TFP for Capacity Utilization

Results

Barsky et al. (2014)

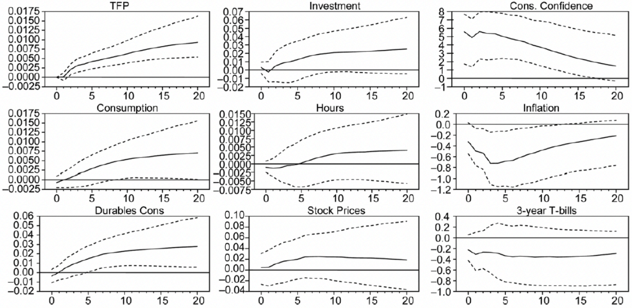

Fig. 3. Impulse responses to orthogonalized news shocks (9-variable VAR, hybrid specification).

Results

Barsky et al. (2014)

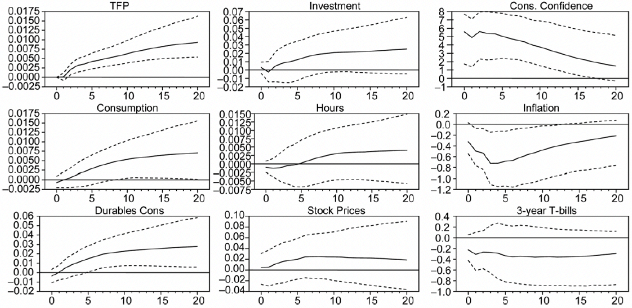

Fig. 3. Impulse responses to orthogonalized news shocks (9-variable VAR, hybrid specification).

Interpretation: improved specification is not consistent with TFP news as a main driver

Outline

Investment technology shocks

- Alternative: maybe shocks specific to **investment technology** (rather than neutral technology shocks) are important

  Goes back to discussions in Keynes (1936).

- Will briefly review two contributions studying the importance of **IST shocks**

  1. Fisher (2006): IST analogue of Galí (1999), identifies IST shocks as the only shocks that move relative IST prices and labor productivity in the long run
  2. Ben Zeev & Khan (2015): use medium-run restrictions to identify IST news shocks, similar to Barsky et al. (2014)

Results

Fisher (2006)

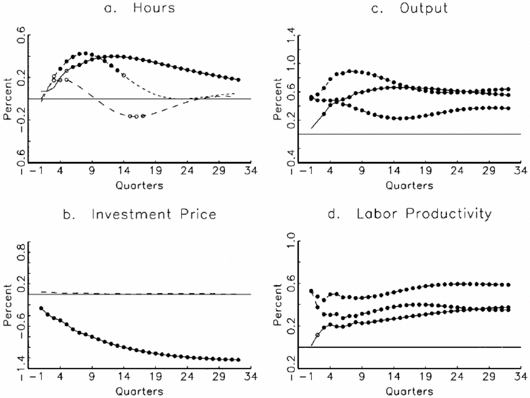

Fig. 5.—Responses to technology shocks, 1982:III–2000:IV. Solid lines denote responses to I-shocks in the five-variable system, long-dashed lines responses to N-shocks in the five-variable system, and short-dashed lines responses to N-shocks in the four-variable system. Solid circles indicate that the response is significant at the 5 percent level and open circles at the 10 percent level.

Results

Ben Zeev & Khan (2015)

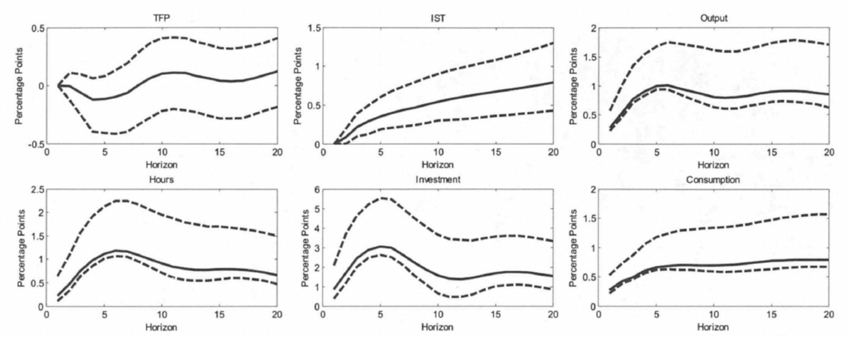

Ben Zeev & Khan (2015)

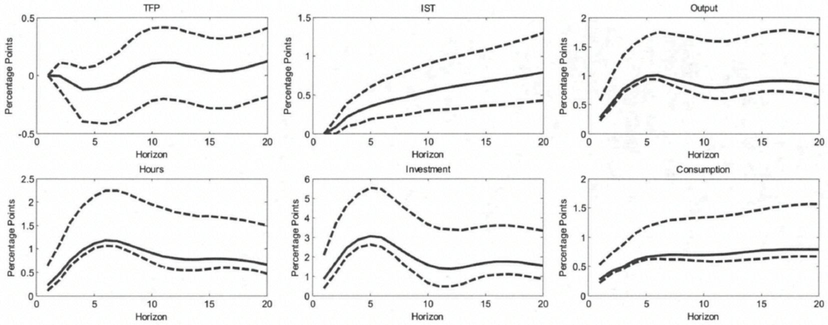

Interpretation: IST news shocks can account for typical business-cycle comovements

Summary

So what do our semi-structural time series strategies teach out about the origins of aggregate business-cycle fluctuations?

- **Main conclusions**
  - Find little support for the classical TFP shocks stressed by the old RBC literature as a source of cyclical fluctuations
  - Somewhat more evidence in favor of investment-specific technology shocks, but evidence is far from conclusive
- Instead, most recent published work on business-cycle origins has leveraged additional **model structure**

Outline

Overview

- The alternative approach is **likelihood-based estimation** of a structural model
  - Write down structural model (= internal propagation mechanism) subject to several exogenous shocks (incl. TFP and IST)
  - Estimation approach: try to find model parameterization that matches second-moment properties of U.S. time series as well as possible

  $\Rightarrow$ Intuition: origins of business cycles = shocks that induce co-movement among aggregate time series that look like the U.S. business cycle

- What we'll do here:
  1. Brief review of **likelihood-based model estimation**
  2. A diagnostic device: "**business-cycle anatomy**"
  3. Frontier estimated models: **RANK & HANK**

Outline

Likelihood-based estimation

- Both Bayesian and classical ML estimation of structural business-cycle models rely on **likelihood evaluation**
- From **state-space model** to **likelihood**
  - A model is a parameter vector ψ giving rise to a state-space system for observables y_t:

$$\begin{aligned} y_t &= \Psi(s_t;\psi) + u_t, \quad u_t \sim F_u(\bullet;\psi) \\ s_t &= \Phi(s_{t-1},\varepsilon_t;\psi), \quad \varepsilon_t \sim F_\varepsilon(\bullet;\psi) \end{aligned}$$

    Note that for now this is slightly more general than before, allowing for non-linearities.

  - We observe data $y_{1:T}$. **Q**: What is the likelihood of $y_{1:T}$ given a parameter vector ψ? This is the key input to standard ML or Bayesian model estimation.
  - Standard approach to likelihood evaluation: filtering

A generic filter

- We can always factorize the likelihood as follows:

$$p(y_{1:T} \mid \psi) = \prod_{t=1}^{T} p(y_t \mid y_{1:t-1}, \psi)$$

- This object can in general be evaluated by proceeding as follows:

0. Let $p(s_0) = p(s_0 \mid y_{1:0})$ be an initial distribution for the states, e.g. the stat. distribution
1. Forecasting t given $t-1$
   a) Transition equation
   $$p(s_t \mid y_{1:t-1}) = \int p(s_t \mid s_{t-1}, y_{1:t-1}) p(s_{t-1} \mid y_{1:t-1}) ds_{t-1}$$
   b) Measurement equation
   $$p(y_t \mid y_{1:t-1}) = \int p(y_t \mid s_t, y_{1:t-1}) p(s_t \mid y_{1:t-1}) ds_t$$
2. Updating with Bayes' theorem

$$p(s_t \mid y_{1:t}) = p(s_t \mid y_t, y_{1:t-1}) = \frac{p(y_t \mid s_t, y_{1:t-1}) p(s_t \mid y_{1:t-1})}{p(y_t \mid y_{1:t-1})}$$

Linear model estimation & beyond

- **Linear models**
  - In linear state-space models with normal errors, likelihood evaluation is straightforward, using the Kalman filter
  - The Kalman filter turns $p(s_{t-1} \mid y_{1:t-1})$ into $p(s_t \mid y_{1:t-1})$, $p(y_t \mid y_{1:t-1})$ and $p(s_t \mid y_{1:t})$. Stringing together T steps, we can evaluate the likelihood.
    See the appendix of Lecture Note 4, or any standard reference on Kalman filtering.
  - Model estimation then proceeds by finding ψ to maximize $p(y_{1:T} \mid \psi)$ (plus perhaps a prior over the ψ's, in the Bayesian case). See Fernandez-Villaverde & Schorfheide (2016) for details.
- **Non-linear models**
  - Estimation of non-linear structural macro models is beyond this course
  - If you're interested: a general-purpose (but demanding) method to solve the generic filtering problem is the so-called particle filter
- **Main concern**: requires entire model to be correctly specified

Outline

Business-cycle anatomy

- Before presenting results from estimated business-cycle models, I will first introduce a useful diagnostic device: "**business-cycle anatomy**"
  - Developed in Angeletos et al. (2021)
  - Purpose: will yield an interesting perspective on a) the origins of cycles in general & b) the limitations of the model estimation results that we'll review later

- What we'll do here:
  1. Sketch the **econometric procedure**
  2. Present the **main results**
  3. Draw some **general lessons** for theories of business cycles

Econometric procedure

- Basic idea: try to find "main shocks" that drive the various aggregate macro var's
- This is implemented using the **max-share** approach to shock identification
  - Let y_t denote a list of macro aggregates (output, consumption, unemployment, investment, inflation, interest rates, TFP ...)
  - Assume invertibility, estimate a VAR in y_t, and invert it to get

    $$y_t = C(L)Q\varepsilon_t$$

    where ε_t denotes orthogonalized unit variance shocks and Q is a rotation matrix
  - Then, for each variable y_k, find the "shock" that accounts for the largest possible share of volatility over frequencies $[\underline{\omega}, \bar{\omega}]$. That is, choose q (where $||q|| = 1$) to maximize

    $$\int_{\omega \in [\underline{\omega}, \bar{\omega}]} \left(\overline{C_k(e^{-i\omega})q}\, C_k(e^{-i\omega})q \right) d\omega = q' \int_{\omega \in [\underline{\omega}, \bar{\omega}]} \left(\overline{C_k(e^{-i\omega})}\, C_k(e^{-i\omega}) \right) d\omega q$$

    Note: no pretense that these are "structural" shocks, but implied patterns may still be interesting

Results

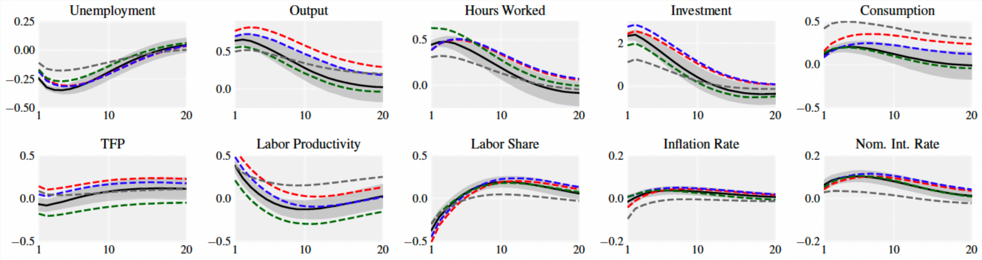

u shock; Y shock; I shock; h shock; C shock; Shaded area: 68% HPDI.

Results

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
u	73.71	58.51	47.72	62.09	20.38	5.86	23.91	27.02	6.96	22.27
Y	56.24	80.13	44.73	67.13	33.03	4.24	41.31	40.20	10.47	16.89
h	49.84	47.54	70.45	47.99	21.78	11.62	22.61	19.47	7.23	22.38
I	59.03	66.60	45.20	80.29	19.01	3.81	33.74	36.44	7.69	21.51
C	19.19	31.59	20.15	17.10	68.30	1.57	12.93	10.31	9.93	4.50

Results: summary

1. The procedure identifies a **common shock** to y, i, c, u
   - "Interchangeability": no matter what variable is targeted, we $\approx$ recover the same shock
   - Note: this result is somewhat less pronounced for c than for the others
2. This identified common shock has some **particular features**:
   - At business-cycle frequencies, it shows little comovement with TFP
   - It has little effect on long-run fluctuations of y, i, c, u
   - It is largely disconnected from inflation dynamics

Interpretation

Why are those findings useful for business-cycle analysis?

1. The data are not inconsistent with **single-shock theories** for cyclical fluctuations
   - This single shock should lead to co-movements of standard real variables, without affecting TFP or prices much
   - Lucas (1977): *"[W]ith respect to the qualitative behavior of comovements among series, business cycles are all alike."*
2. Several candidates for this main business-cycle driver are **ruled out**
   - TFP (news) shocks would map into fluctuations into TFP (which we don't see) Note: same would hold for many typical financial or uncertainty shocks—they map into TFP "wedges".
   - Textbook demand shocks would lead to y-π comovements (which we don't see)
3. Even **multi-shock accounts** of cycles should be consistent with the anatomy as a **reduced-form moment** [Will see next: this is an informative test.]

Outline

Estimated NK models

- Let's now consider results from **structural NK model estimation**
  - Environment: RBC core + various frictions (nominal rigidities, adjustment costs, ...)
  - Then add various candidates for "shocks" as business-cycle drivers (TFP, investment technology, policy, ...) and try to account for agg. fluctuations in a likelihood sense
  - **Q**: which shocks are picked out as main cyclical drivers?
    Key logic: will pick the shocks that generate "business cycle-like" impulse responses.

- Will look at two prominent contributions to the literature:
  1. Justiniano-Primiceri-Tambalotti (2010): canonical estimated RANK model
  2. Auclert-Rognlie-Straub (2020): JPT + high household MPCs ("HANK")

Justiniano et al. (2010): investment shock IRFs

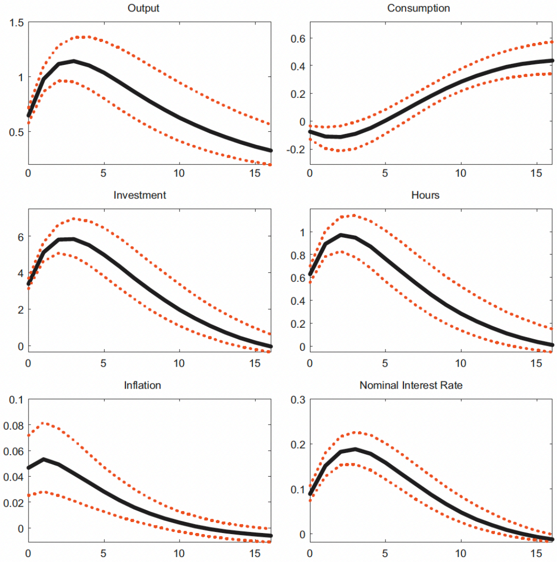

Justiniano et al. (2010): cycle decomposition

Table 1
Posterior variance decomposition at business cycle frequencies in the baseline model.

	Medians and [5th, 95th] percentiles						
Series\shock	**Policy**	**Neutral**	**Government**	**Investment**	**Price mark-up**	**Wage mark-up**	**Preference**
Output	0.05	0.25	0.02	0.50	0.05	0.05	0.07
	[0.03, 0.08]	[0.19, 0.33]	[0.01, 0.02]	[0.42, 0.59]	[0.03, 0.07]	[0.03, 0.08]	[0.05, 0.10]
Consumption	0.02	0.26	0.02	0.09	0.01	0.07	0.52
	[0.01, 0.04]	[0.20, 0.32]	[0.02, 0.03]	[0.04, 0.16]	[0.00, 0.01]	[0.04, 0.12]	[0.42, 0.61]
Investment	0.03	0.06	0.00	0.83	0.04	0.01	0.02
	[0.02, 0.04]	[0.04, 0.10]	[0.00, 0.00]	[0.76, 0.89]	[0.02, 0.06]	[0.01, 0.02]	[0.01, 0.04]
Hours	0.07	0.1	0.02	0.59	0.06	0.07	0.08
	[0.04, 0.10]	[0.08, 0.13]	[0.02, 0.03]	[0.52, 0.66]	[0.04, 0.09]	[0.04, 0.11]	[0.06, 0.12]
Wages	0.00	0.4	0.00	0.04	0.31	0.23	0.00
	[0.00, 0.01]	[0.30, 0.52]	[0.00, 0.00]	[0.02, 0.07]	[0.23, 0.41]	[0.16, 0.32]	[0.00, 0.01]
Inflation	0.03	0.14	0.00	0.06	0.39	0.34	0.02
	[0.02, 0.06]	[0.09, 0.21]	[0.00, 0.00]	[0.02, 0.13]	[0.29, 0.50]	[0.26, 0.42]	[0.01, 0.04]
Interest rates	0.17	0.09	0.01	0.47	0.05	0.04	0.16
	[0.13, 0.22]	[0.06, 0.12]	[0.00, 0.01]	[0.37, 0.56]	[0.03, 0.07]	[0.03, 0.07]	[0.11, 0.23]

Note: Business cycle frequencies correspond to periodic components with cycles between 6 and 32 quarters. The decomposition is obtained using the spectrum of the DSGE model and an inverse first difference filter for output, consumption, investment and wages to reconstruct the levels. The spectral density is computed from the state space representation of the model with 500 bins for frequencies covering that range of periodicities. Medians need not add up to one.

Main result: investment technology shock as important cyclical driver

Discussion

- Why does the **investment shock** play such a big role?
  - It wouldn't in vanilla RBCs: better investment opportunities $\rightarrow$ save today $\rightarrow$ $c \downarrow\downarrow$ today
    Note: you may sometimes see this referred to as the "Barro-King curse".
  - Here various model features help: habit formation leads to slow movements of c, and sticky prices/variable capacity utilization enable $i \uparrow$ to be largely accommodated by $y \uparrow$
- Yet it still does not look like the **main business-cycle driver** of Angeletos et al. (2021)
  - Note: consumption is essentially flat at the beginning, and the i shock accounts for little of the business-cycle fluctuations in c
  - Can formalize this through business-cycle anatomy: consumption still needs its own shock (the impatience shock), see next slide
  - Interpretation: second-moment properties of the data are matched poorly along the dimension that matters for the Angeletos et al. decomposition

Discussion

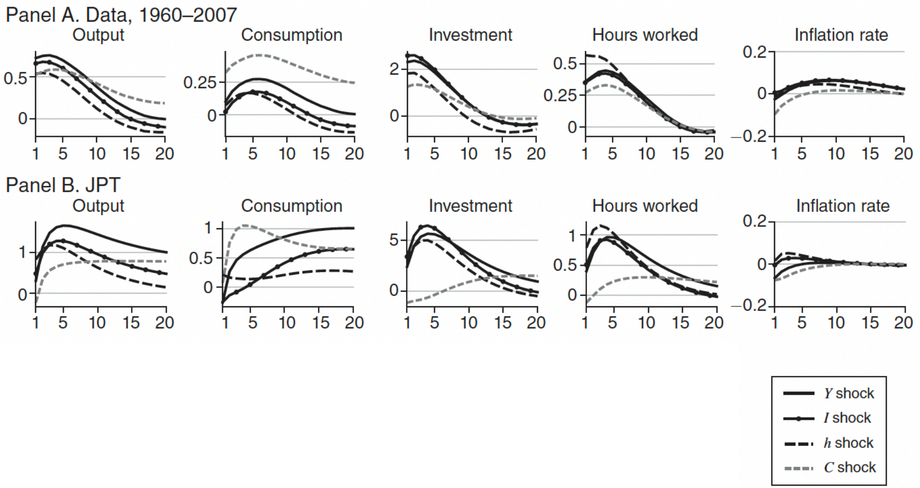

Panel A. Data, 1960–2007
Output
Consumption
Investment
Hours worked
Inflation rate
Panel B. JPT
Output
Consumption
Investment
Hours worked
Inflation rate
Y shock
I shock
h shock
C shock

Auclert et al. (2020): investment shock IRFs

Figure 14: Impulse response to a 1-standard deviation investment shock

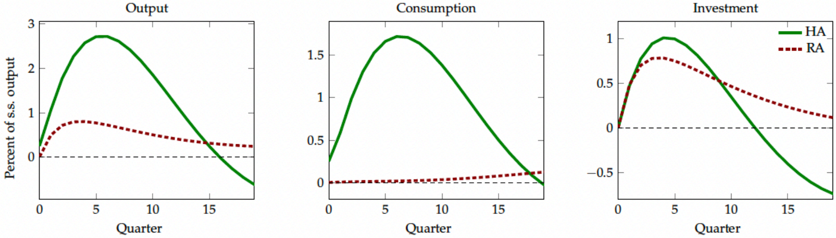

Note. Impulse responses are to a one-time negative shock to risk premia ϵ_t^I, using the estimated standard deviation and persistence.

Mechanism: high MPCs endogenously tie consumption to (labor) income

Auclert et al. (2020): cycle decomposition

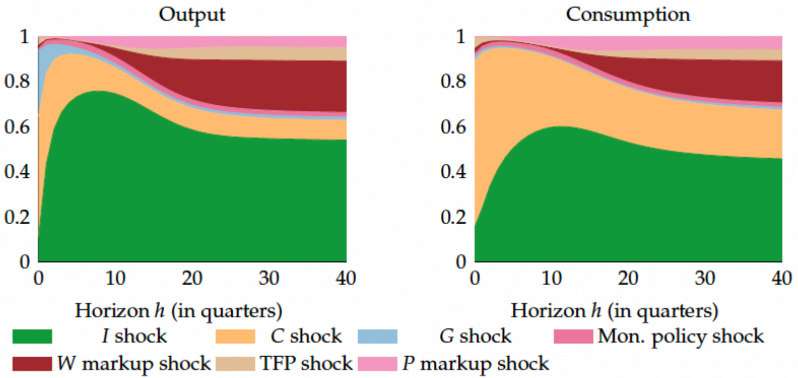

Main result: due to high MPCs, IST shock can become the key driver of the cycle

Conjecture: their model probably also passes the Angeletos et al. "interchangeability" test.

Summary

- **Main takeaways**

  (i) Data are not inconsistent with a single (type of) shock as the main cyclical driver

  (ii) Classical TFP shocks don't work, but investment technology/demand-type shocks do

  (iii) Promise of HA: tie C endogenously to the cycle

- Note: hard to credibly establish much using our purely semi-structural approaches. Tend to still need quite a bit of **model structure** for business-cycle origin exercises.