Aggregate Uncertainty: Accuracy and Linear Methods

Manuel Macera | 2026

Where we are

- So far: after an unexpected one-time shock, households have perfect foresight, and we can compute the transition with the whole distribution.

- Now: **aggregate shocks** hit the economy every period, so the distribution never settles and households cannot know future prices.

- Two routes, which approximate in different directions:
  - a **global** solution in which households track a few moments of λ *(Krusell and Smith, 1998; KS from now on)*;
  - a **linear** solution around the stationary equilibrium, which keeps the whole distribution and builds on the transitions of deck 4a.

Plan

1. Approximate equilibrium: KS's method, its result, and why the mean is enough.

2. **Accuracy**: how to test a law of motion, and what approximate aggregation does and does not mean.

3. **Linear methods**: the impulse response as a derivative, and the sequence-space Jacobian.

Last lecture covered Part 1; this one covers Parts 2 and 3.

Part 2: accuracy

Testing accuracy

- After KS, laws of motion were judged mostly by R^2 and $\hat{\sigma}$. Both measure **one-step-ahead** errors: each period, the forecast starts from the simulated K_t.

- Den Haan (2010): they **cannot tell** a good law from a bad one. A law with a slightly wrong slope errs little in one quarter, since $\ln K$ barely moves; iterated over many quarters, its errors **compound**.

Den Haan's warning

- Take a solved economy and its simulated data, $\{z_t, K_t^{\text{sim}}\}$. **Nothing is re-solved**: households' policies, and so K_t^{sim}, stay as they are.
- Lower each fitted slope by the same amount, $b_1(z) = \hat{b}_1(z) - d$, and reset each $b_0(z)$ so that the mean error is zero.
- Evaluate the new law on the same data in two ways:
  - one step ahead, as the R^2 does, starting each period from the simulated K:

$$\ln \hat{K}_{t+1} = b_0(z_t) + b_1(z_t) \ln K_t^{\text{sim}};$$

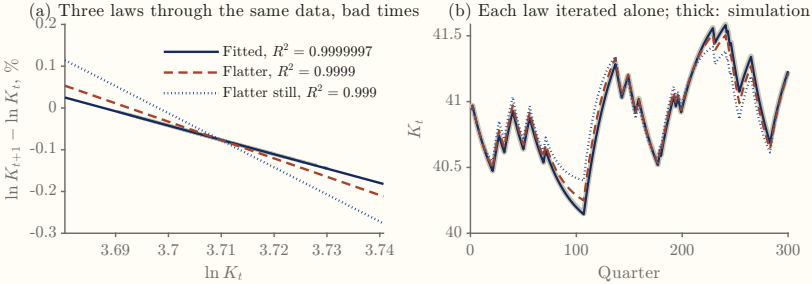
  - iterated alone, never updated with the simulated K:

$$\ln K_{t+1}^{\text{law}} = b_0(z_t) + b_1(z_t) \ln K_t^{\text{law}}.$$

- The new laws are not equilibria. The question is only whether the R^2 would notice that they are wrong.

Output: Den Haan's warning in our economy

Our economy: the fitted law and two flatter laws through the same simulated data



- R^2 stays above 0.999, but iterated alone the flatter laws imply **less persistent** capital.

Den Haan's results

- Goodness of fit does not **discriminate** between laws with very different dynamics.

The same experiment in Den Haan's KS economy

Slope of the law	0.964 (fitted)	0.954	0.932	0.864
R^2 on the simulated data	0.999997	0.9999	0.999	0.99
St. dev. of $\ln K$ implied by the law	0.0248	0.0217	0.0174	0.0113

- In his simulated economy the standard deviation is 0.0248: 14% more than the law with $R^2 = 0.9999$ implies, and 119% more at 0.99.
- KS's case for one moment rests on the solution barely changing when a moment is added, and on **25-year forecast errors**, not on the R^2.
- When could an algorithm deliver such a law? With a family that is too restrictive, a missing moment, or an algorithm stopped too early.

A better test

- Den Haan makes KS's 25-year idea the standard test.
- Use the same path of z to generate two series of K, from the same K_0:
  - aggregating households' choices in the simulation,

$$K_{t+1}^{\text{sim}} = \int a^{\star}\left(a, e; z_t, K_t^{\text{sim}}\right) \lambda_t(\mathrm{d}a, \mathrm{d}e);$$

  - iterating the law of motion alone, **never updated** with K_t^{sim},

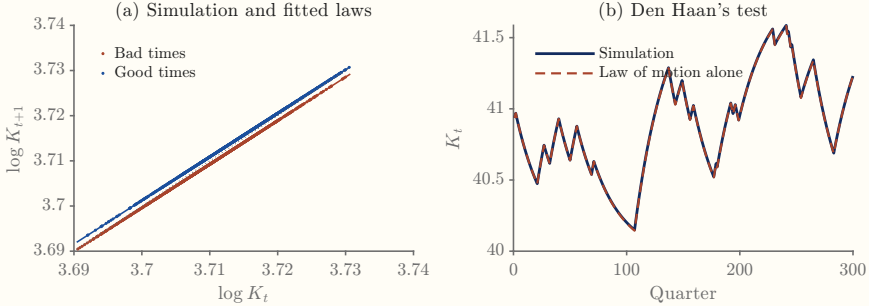
$$\ln K_{t+1}^{\text{law}} = b_0(z_t) + b_1(z_t) \ln K_t^{\text{law}}.$$

- Report the **maximum** error, $\max_t \left| \ln K_t^{\text{law}} - \ln K_t^{\text{sim}} \right|$. Averages can hide large errors in rare states, such as deep recessions.

Output: KS in an Aiyagari economy

- To compare methods, we solve one economy with all of them: Aiyagari's economy, with households holding capital, plus productivity shocks.

The law of motion in a simulation of 5,500 quarters, and Den Haan's test



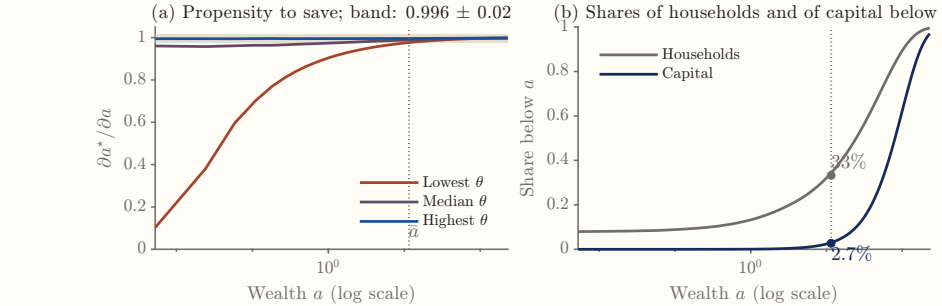
- Den Haan's test: the law of motion alone stays within 0.022% of the simulated K over 5,500 quarters, 0.006% on average.

The meaning of approximate aggregation

- For us, approximate aggregation had a **strong** meaning: a representative-agent model reproduces the aggregates.
- In KS, it has a **weaker** meaning: the mean of wealth is enough to forecast prices.
- Aren't they the same? Strong implies weak: a representative agent's K' depends only on K and z. The converse **fails**:
  - KS's benchmark has both: its aggregates are almost those of the complete-markets economy, with 0.6% more capital.
  - KS's model with random discount factors, built to match the wealth distribution, has only the weak one. The law for K still forecasts almost perfectly, but corr(c, y) is 0.83, against 0.69 with complete markets.

Output: why the mean is enough

Our economy without aggregate risk: the propensity to save, and who holds the capital



- Last lecture: with a **common propensity to save** out of wealth, K is enough to forecast K'. Here, above $\bar{a} \approx 11.5$, every propensity is within 0.02 of 0.996; below $\bar{a}$ are 33% of households, but only 2.7% of the capital.
- Aggregate shocks keep the economy close to this stationary equilibrium.

Part 3: linear methods

Two kinds of approximation

- **KS**: policies are solved globally and the distribution is simulated exactly. What is approximated is the **forecast**: households use

$$\ln K' = b_0(z) + b_1(z) \ln K \qquad \text{instead of} \qquad \lambda' = T(\lambda, z, z').$$

- **Linear methods**: policies and the distribution are exact at the stationary equilibrium. What is approximated is how policies and the distribution, and so every aggregate, **respond** to shocks, to first order in $\mathrm{d}Z$:

$$X_t = X_{ss} + \sum_{s \geq 0} \frac{\partial X_t}{\partial Z_s} \mathrm{d}Z_s.$$

- Linear methods impose **certainty equivalence**: the economy's average equals the stationary equilibrium without aggregate risk. They require responses **proportional to the size of the shock**.

Why other methods

- KS's method needs a law of motion for every moment that matters, and solves the household problem on grids for every aggregate state. With many aggregate shocks or prices, this becomes costly.

- Can we get something out of our transitional dynamics, which were much simpler to compute?

The impulse response as a numerical derivative

- Boppart, Krusell and Mitman (2018; BKM from now on): compute, nonlinearly, the transition after one **small** MIT shock: $\ln z$ rises by ε at $t = 0$ and returns at rate ρ_z, the path expected after an innovation. Call x_t the resulting path of any aggregate.
- Its deviation from the steady state, per unit of shock, is a **numerical derivative**:

$$\hat{x}_t \equiv \frac{x_t - x_{ss}}{\varepsilon} \approx \frac{\partial x_t}{\partial \varepsilon_0}.$$

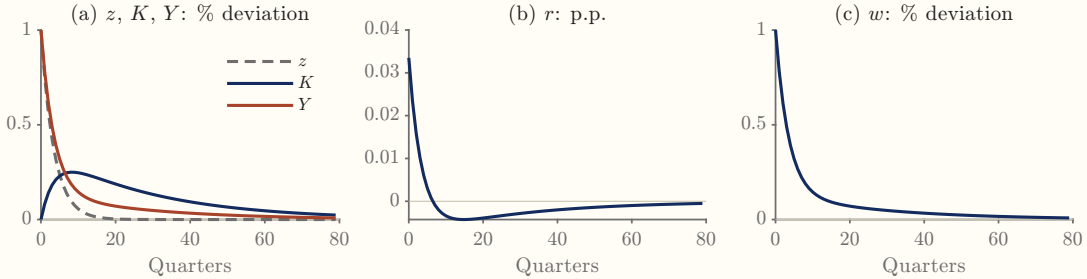
  Since z returns at rate ρ_z, the persistence of the shock is built into $\hat{x}_t$.
- If the economy is **approximately linear** in aggregate shocks, responses scale and add up, and one deterministic path simulates the economy with risk. With an innovation ε_t every period, build

$$X_t = x_{ss} + \hat{x}_0\,\varepsilon_t + \hat{x}_1\,\varepsilon_{t-1} + \hat{x}_2\,\varepsilon_{t-2} + \cdots$$

- Under **certainty equivalence**, the response to a surprise equals the average response to the same innovation when shocks are expected.

Output: the response to a small MIT shock

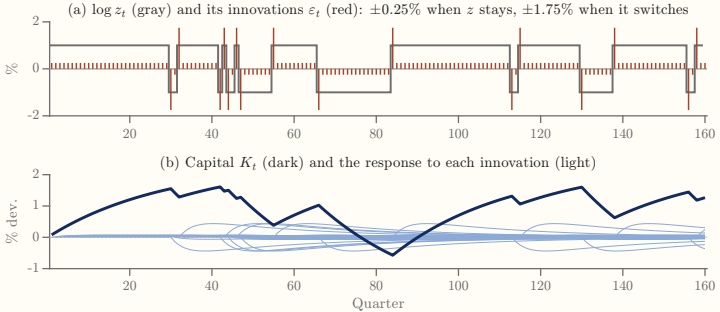
Our economy: nonlinear transition after $\varepsilon = 0.1\%$, scaled to a 1% innovation in z



- Capital peaks after 8 quarters at 0.25% above its steady state. The interest rate rises on impact and then falls below its steady state as capital accumulates.

Output: shocks every period

160 quarters of the linear solution: capital as a sum of scaled and shifted responses



- $\log z$ is expected to revert, so **staying put is a surprise**: with $\ln z \in \{\pm 1\%\}$,

$$\varepsilon_t = \ln z_t - \rho_z \ln z_{t-1} = \begin{cases} \pm(1-\rho_z)\,1\% = \pm 0.25\% & \text{if } z \text{ stays,} \\ \pm(1+\rho_z)\,1\% = \pm 1.75\% & \text{if } z \text{ switches.} \end{cases}$$

- Capital sums all past responses, $\hat{K}_0\varepsilon_t + \hat{K}_1\varepsilon_{t-1} + \cdots$; each fades slowly.

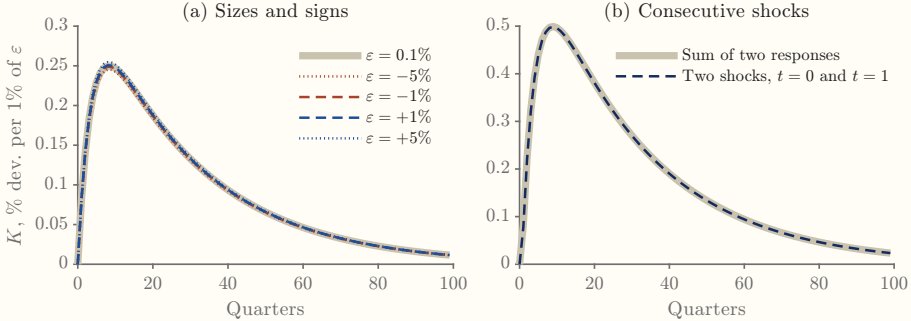
Is linearity a good approximation?

- BKM propose **scalability checks**, each one a few more transitions:
  - shocks of different sizes and signs: are the responses proportional, and symmetric around zero?
  - two consecutive shocks, at $t = 0$ and $t = 1$: is the response the sum of the two single responses?
  - two kinds of shock at once: is the response the sum?
- If they fail, nonlinear methods are needed.

- Linearization does not drop the distribution: its movements enter to first order.

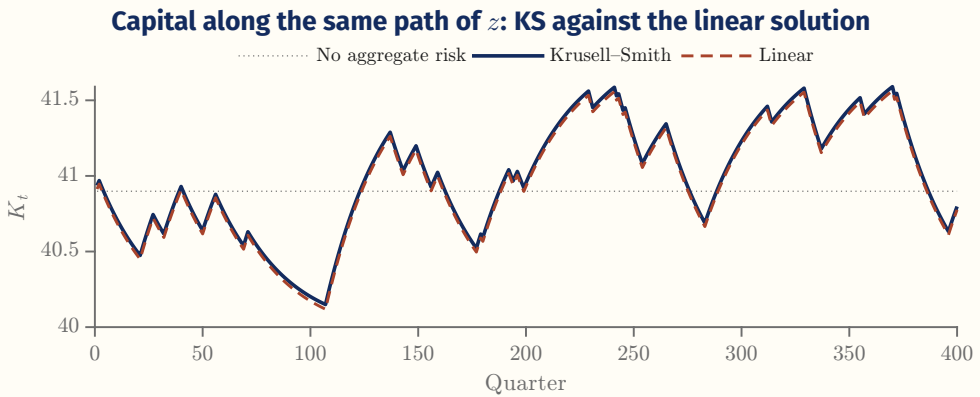
Output: scalability checks

Response of K per 1% of innovation, for shocks of different sizes, signs and timing



- Scaled to the same size, a 1% shock differs from the 0.1% one by at most 0.4% of the peak response; a 5% shock, by 2%. The gap grows with the size of the shock: a second-order term.
- Two consecutive 1% shocks differ from the sum of two responses by 1% of the peak.

Output: one path of z, two solutions



- The two series differ by at most 0.09%, and the standard deviation of $\ln K$ is the same: 0.79% in both.
- The gap is mostly a level: the linear solution misses the 0.06% of extra capital that aggregate risk induces. That is certainty equivalence.

A second approach: the sequence space

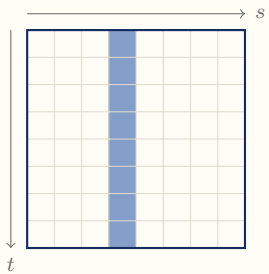
- Auclert, Bardóczy, Rognlie and Straub (2021): write equilibrium as equations in the paths of a few aggregates, $H(U, Z) = 0$, with unknowns U and shocks Z.
- Differentiate these equations **exactly**, around the stationary equilibrium.
- The size of the system depends on the horizon T, not on the grid.
- Once the derivatives are known, any path of shocks costs one matrix multiplication.

The sequence-space Jacobian

- Our economy in sequence space, with last lecture's unknowns, $U = (r_1, \ldots, r_T)$, and shocks $Z = (\ln z_0, \ldots, \ln z_{T-1})$. Market clearing at every date,
$$H_t(U, Z) = A_{t-1}(\{r_s, w_s\}) - K(r_t, z_t) = 0, \qquad t = 1, \ldots, T,$$
where households map price paths into asset paths, and firms give w_t and capital demand.
- To first order, the response to a productivity path $\mathrm{d}Z$ is $\mathrm{d}U = -H_U^{-1} H_Z \, \mathrm{d}Z \equiv G \, \mathrm{d}Z$.
- H_U and H_Z combine the firms' derivatives, date by date, with two **Jacobians** of the household block: $J^r_{ts} = \partial A_t / \partial r_s$ and $J^w_{ts} = \partial A_t / \partial w_s$.
- Column s of J^r is **partial equilibrium**: news at date 0 that r is higher **at date s only**, all other prices at their steady state. z enters only through the firms.
- To first order, households matter only through these Jacobians, which do **not depend on the shock**: a new shock needs only its H_Z; a new path, one multiplication by G.

Computing the Jacobian: the direct method

- For column s: raise r_s, solve households **backward** for the policies, and move the distribution **forward**, over T periods.
- The price path is given, so there is **no iteration** on prices: one backward and one forward pass per column.
- Column s is $(A_t - A_{ss})/\mathrm{d}r$ for $t = 0, \ldots, T-1$. With T columns, about T^2 updates of the distribution.
- The **fake news** algorithm gets the same matrix in two steps: first a matrix F, then J from F.

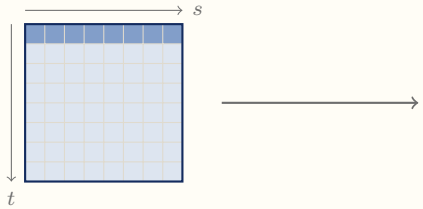


column s: one backward and one forward pass

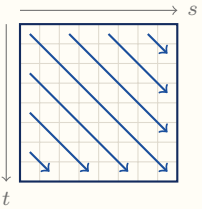
The fake news algorithm

- **Fake news** about date s: announced at date 0, **retracted** at date 1. Its responses F are cheap: **one** backward pass gives row 0 for every horizon; below row 0, only the displaced distribution moves, under steady-state policies.
- True news is fake news **announced again every period**: $J_{ts} = F_{ts} + J_{t-1,s-1}$.
- No column needs its own transition: about T updates of the distribution instead of T^2.

(a) F: row 0, then the displaced distribution

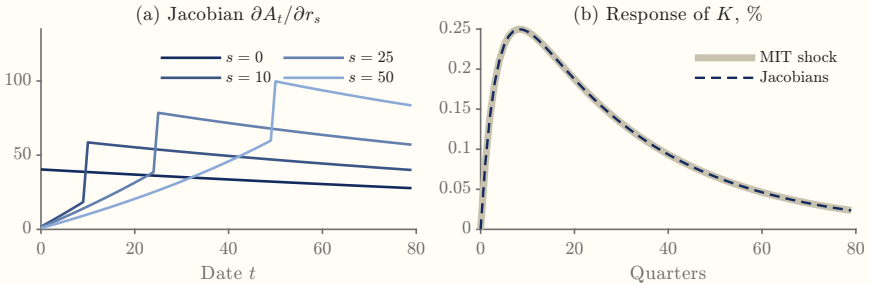


(b) J: sum F down each diagonal



Output: the household Jacobian in our economy

Jacobian of assets with respect to r, and the response of K two ways



- Both Jacobians are 250×250; one column matches the direct method to 3×10^{-5}.
- Panel (b): the response of K from the Jacobians and from the small MIT shock coincide, to within 5×10^{-4} of the peak.
- Column s of $G = -H_U^{-1} H_Z$ is the general-equilibrium response to news, at date 0, that productivity will be higher at date s.

What the linear methods cost

Method	Cost in our economy
BKM	One transition per shock, 23 iterations, about a minute
Sequence-space Jacobian	Jacobians in about 10 seconds; then any shock path in 0.06 seconds

- KS's cost grows with every moment and every aggregate shock added to the state; the linear methods' cost does not.

- Auclert et al. report, in their KS calibration, a 300×300 Jacobian of K in 90 milliseconds, and under 11 seconds with 250,000 idiosyncratic states.

Summary

- Test a law of motion by simulating it without updating, and look at maximum errors: a high R^2 is not enough. In our economy, KS's law passes easily.

- Approximate aggregation means the mean forecasts prices. It need not mean that a representative agent reproduces the aggregates: capital and consumption weight households differently.

- Linear methods keep the distribution and linearize in the aggregate shocks. BKM and the sequence-space Jacobian compute the same impulse responses. They reproduce KS's dynamics up to certainty equivalence.

References

Auclert, A., B. Bardóczy, M. Rognlie and L. Straub (2021). “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models.” *Econometrica* 89(5): 2375–2408.

Boppart, T., P. Krusell and K. Mitman (2018). “Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative.” *Journal of Economic Dynamics and Control* 89: 68–92.

Den Haan, W. J. (2010). “Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents.” *Journal of Economic Dynamics and Control* 34(1): 79–99.

Krusell, P. and A. A. Smith, Jr. (1998). “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy* 106(5): 867–896.