Aggregate Uncertainty

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.

This lecture covers Part 1; the next one covers Parts 2 and 3.

Part 1: approximate equilibrium

The KS economy

- The stochastic growth model with a continuum of households who cannot insure their idiosyncratic risk. A period is a quarter.
- Technology: $Y = z\,K^{\alpha}L^{1-\alpha}$. Productivity $z \in \{z_b, z_g\}$ follows a Markov chain with transition probabilities $\pi_{zz'}$.
- Employment status: $e \in \{0, 1\}$. An employed household supplies $\tilde{l}$ units of labor; an unemployed one earns nothing.
- e also follows a Markov chain, **but** its transition probabilities depend on today's and tomorrow's z:

$$\pi_{ee'}(z, z') = \Pr(e' \mid e, z, z').$$

- Given z and z', individual shocks are independent, and a law of large numbers holds. The chain is chosen so that the unemployment rate, u_z, and aggregate labor, $L(z) = (1 - u_z)\,\tilde{l}$, depend only on z: z is the only aggregate risk.

Calibrating the two chains

- Parameters: 2 for the chain of z, π_{gg} and π_{bb}; and, for each of the four moves (z, z'), a 2×2 matrix for e whose rows add up to one, 2 each. In all, **10**.
- Ten conditions pin them down:

Condition	Number
Booms and recessions last 8 quarters on average	2
Unemployment moves exactly between 4% in booms and 10% in recessions, for each move (z, z')	4
Unemployment spells last 1.5 quarters in booms and 2.5 in recessions	2
Staying unemployed is 25% more likely when a recession starts than while it lasts, and 25% less likely when a boom starts than while it lasts	2

Households hold capital

- In KS, capital is the only asset. Households own capital, a, and rent it to firms. The budget constraint reads

$$c + a' = w(z, K)\,\tilde{l}\,e + \big(1 + r(z, K)\big)\,a, \qquad a' \geq 0,$$

  with $r = z\,F_K\big(K, L(z)\big) - \delta$ and $w = z\,F_L\big(K, L(z)\big)$. There is no separate asset market: $K = \int a \,\mathrm{d}\lambda$.

- K was chosen yesterday. Knowing K and z is enough for today's prices, **exactly**: the distribution matters only to forecast K', and with it future prices.

- The unemployed earn nothing, so a household in debt could not always repay: zero is the **natural** borrowing limit.

The household's problem with aggregate risk

- The household's problem is

$$V(a, e; z, \lambda) = \max_{c,\, a'} \; u(c) + \beta\, \mathbb{E}\big[V(a', e'; z', \lambda') \mid e, z\big]$$

$$\text{subject to} \quad c + a' = w(z, K)\, \tilde{l}\, e + \big(1 + r(z, K)\big)\, a, \qquad a' \geq 0,$$

$$K = \int a \, \mathrm{d}\lambda, \qquad \lambda' = T(\lambda, z, z').$$

- Four state variables: two individual, (a, e), and two aggregate, (z, λ); two exogenous, (e, z), and two endogenous, (a, λ).

- T depends on z' although a' is chosen today, because the law of e' depends on (z, z').

Recursive competitive equilibrium

Recursive competitive equilibrium (KS)

A law of motion T, a value function V and policy $a^\star$, and pricing functions $r(z, K)$ and $w(z, K)$ such that

1. given T, r and w, the functions V and $a^\star$ solve the household's problem;
2. r and w are competitive: marginal products at $K = \int a \, \mathrm{d}\lambda$ and $L(z)$, which is how prices depend on λ;
3. T is generated by $a^\star$: for every set of assets $\mathcal{A}$ and every e',

$$\lambda'(\mathcal{A} \times \{e'\}) = \sum_e \int \mathbf{1}\{a^\star(a, e; z, \lambda) \in \mathcal{A}\} \, \pi_{ee'}(z, z') \, \lambda(\mathrm{d}a, e).$$

- Condition 3 is the hard one: T maps distributions into distributions.

Why the distribution matters now

- To decide how much to save, the household needs tomorrow's prices, which depend on tomorrow's capital,

$$K' = \int a^\star(a, e; z, \lambda)\, \lambda(\mathrm{d}a, \mathrm{d}e),$$

  which in general depends on the whole distribution, not only on K.

- And now λ never settles: aggregate shocks keep moving it. The state is **infinite-dimensional** and does not stay put.

- We cannot solve this problem exactly. Two options:
  - assumptions under which the distribution drops out (complete markets, deck 1);
  - an **approximate equilibrium**.

KS's idea

- Suppose households perceive prices as depending only on z and a few **moments** m of λ.
- They forecast the moments with a law of motion from a given **family**, $m' = H(m, z, z'; b)$, indexed by parameters b.
- Their question: do households who use such a rule make large forecast errors? If not, the economy is close to a true equilibrium.
- So the task is to find, within the family, a law that forecasts well and is cheap to compute with. If the errors are large, add a moment or use a more flexible family.

KS's choice

- One moment: mean wealth, which is aggregate capital K.
- A log-linear family:

$$\ln K' = b_0(z) + b_1(z)\ \ln K.$$

  K' does not depend on z': it is decided today.

- The parameters cannot be anything: they must be consistent with households' decisions. The law $b_0 = 0$, $b_1 = 1$, which says capital never moves, would be refuted by the simulated economy.

The household's problem with a perceived law of motion

- Replace λ with K in the household's problem:

$$V(a, e; z, K) = \max_{c,\, a'} \; u(c) + \beta\, \mathbb{E}\big[V(a', e'; z', K') \mid e, z\big]$$

$$\text{subject to} \quad c + a' = w(z, K)\, \tilde{l}\, e + \big(1 + r(z, K)\big)\, a, \quad a' \geq 0,$$

$$\ln K' = b_0(z) + b_1(z)\, \ln K.$$

- Given $\{b_0(z), b_1(z)\}$, this is a standard problem with two more states, z and K, and a grid for K. Harder than the stationary one, but solvable.

- The new fixed point is in the parameters b: with the log-linear law and two aggregate states, four numbers.

Approximate equilibrium

Approximate equilibrium (KS)

Households perceive that a vector m of moments of λ follows $m' = H(m, z, z')$, with H in a given **family** $\mathcal{S}$. An approximate equilibrium is a law H, a value function V and policy $a^\star$, and pricing functions $r(z, K)$ and $w(z, K)$ such that

1. given H, r and w, the functions V and $a^\star$ solve the household's problem;
2. r and w are competitive;
3. H is the **best fit**, within $\mathcal{S}$, to the path of m that $a^\star$ implies, and the fit is **almost perfect**.

- Compare with the recursive competitive equilibrium: **only condition 3 changes**. KS: it satisfies "all the standard equilibrium conditions except the agents' ability to make perfect forecasts."

How to read an approximate equilibrium

- **Bounded rationality**: households ignore information about the distribution. The bound on what they can forecast hardly constrains them.

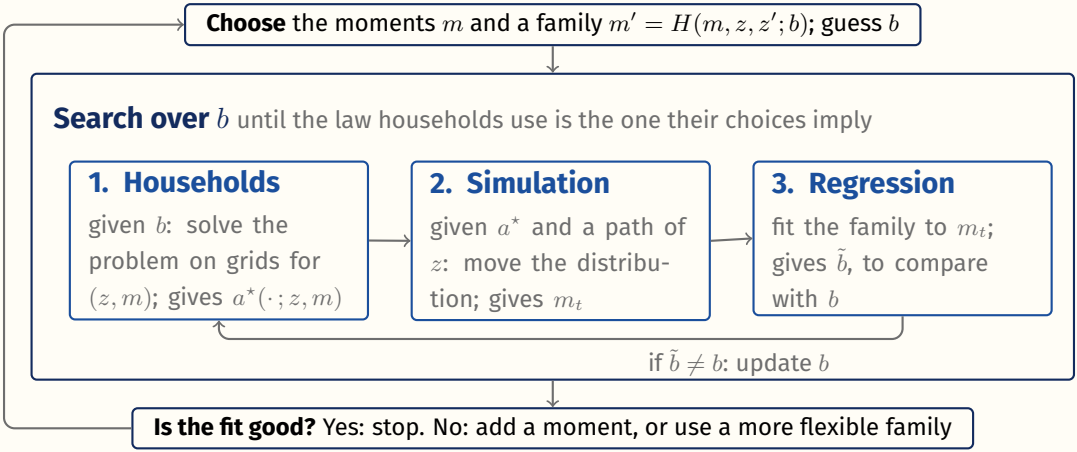
- **Rational, given costs**: better forecasts would raise a household's welfare by a vanishingly small amount. If gathering and processing more information has any cost, following the simple rule is the sensible choice.

The algorithm

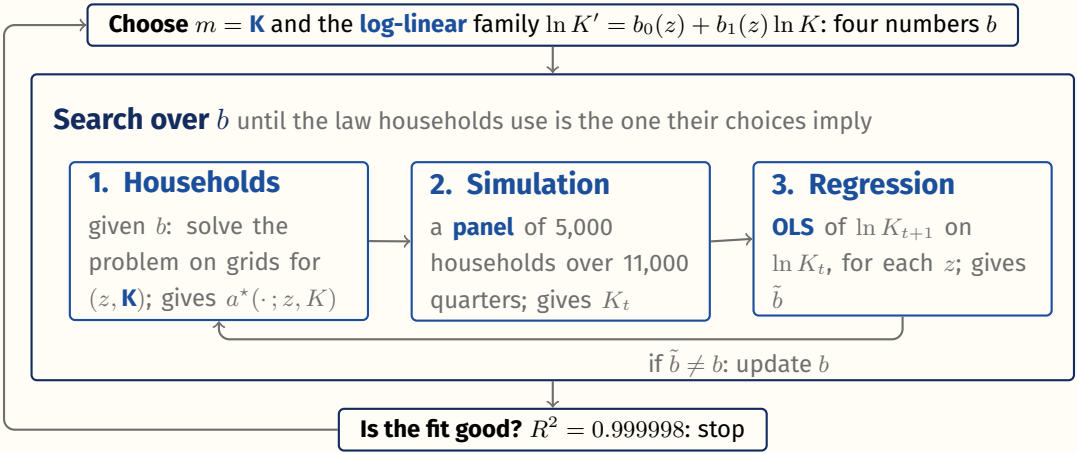
Algorithm: KS

1. Choose m and a family for H; guess b (KS: $m = K$, log-linear).
2. Solve the household's problem, with grids for the aggregate states.
3. Draw a path of length T for z, and paths of e for N households, with exactly a share u_{z_t} unemployed each period.
4. Simulate the panel; compute m_t in every period. Discard the first periods.
5. Fit the family to the simulated m_t to get $\tilde{b}$. Log-linear law: for each z, regress $\ln K_{t+1}$ on $\ln K_t$ by OLS over the periods with $z_t = z$.
6. If $\tilde{b}$ is close to b, go to step 7; otherwise update b and go back to step 2.
7. If the fit is good, stop. Otherwise add a moment, or use a more flexible family, and go back to step 1.

The algorithm in a picture



The algorithm in a picture: KS's choices



Comments on the algorithm

- Two nested fixed points: the household's problem, given the law of motion; and the law of motion, given households' choices.
- KS do not say how they update the guess in step 6. A damped update, $b \leftarrow \omega \tilde{b} + (1-\omega)\, b$ with $\omega \in (0, 1]$, as in last lecture's transition algorithm, is common practice.
- Step 3 simulates a panel of households, with the sampling noise we discussed last lecture. KS use 5,000 households and 11,000 quarters, and discard the first 1,000.
- Imagine computing this in the mid-1990s: the cost is in steps 2 and 3, repeated for every guess.

Simulating without sampling noise

- Young (2010) replaces the panel in step 3 with last lecture's **histogram**: carry the distribution on a fixed grid along the simulated path of z.
- Recall that each period, the mass at (a, e) is split between the two grid points around a' with the lottery weights, and across e' with the transition probabilities given (z, z').
- He starts the simulation from the stationary distribution without aggregate shocks, already close to the economy's ergodic set.
- No sampling error, and it is much faster than simulating a panel.

Prediction: only the mean matters

- KS's benchmark: quarterly, $\beta = 0.99$, log utility, $\delta = 0.025$, capital share 0.36; $z \in \{0.99, 1.01\}$, and the chains as above.
- With only the mean, the law of motion they find is

$$\text{good times:} \quad \ln K' = 0.095 + 0.962 \ln K, \quad R^2 = 0.999998,\ \hat{\sigma} = 0.0028\%;$$
$$\text{bad times:} \quad \ln K' = 0.085 + 0.965 \ln K, \quad R^2 = 0.999998,\ \hat{\sigma} = 0.0036\%.$$

- $\hat{\sigma}$ is the standard deviation of the regression error, in percent. Households that forecast prices with aggregate capital alone make almost no forecast errors.
- KS call this **approximate aggregation**: all aggregates, including consumption, capital and prices, are almost perfectly described by the mean of the wealth distribution and z.

Why approximate aggregation?

- Suppose saving were linear in own wealth, with a **common slope**: $a^\star(a, e; z, K) = p_0(e; z, K) + p_1(z, K)\, a$. Intercepts may differ across e; slopes may not.
- Aggregating,

$$K' = \int p_0(e; z, K)\, \mathrm{d}\lambda + p_1(z, K)\, K.$$

  The first term depends only on the share employed, which is known given z. So K' is a function of (z, K): the mean is a sufficient statistic, though the law need not be log-linear.
- Actual policies differ: the propensity to save out of wealth rises toward one with wealth, and is lowest near the limit. But most capital is held by well-insured households with nearly the same propensity to save (we come back to this next lecture).

Heterogeneity and aggregates

- The finding was influential. Lucas (2003), describing KS's work: for individual behavior and welfare, heterogeneity is everything; but for aggregates, "realistically modeled household heterogeneity just does not matter very much."

- A tempting reading: for business cycles, heterogeneous-agent models behave like representative-agent models. Is that what KS found? We come back to this next lecture.

Summary

- With aggregate risk the distribution is a state that never settles, and households need it to forecast prices.
- KS: households forecast a few moments with a law of motion from a given family. An approximate equilibrium meets every equilibrium condition except perfect forecasts.
- With the mean of wealth and a log-linear law, households make almost no forecast errors: approximate aggregation. The households who hold the capital have nearly the same propensity to save.
- Next lecture: linear methods, which keep the whole distribution and linearize in the aggregate shocks.

References

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

Lucas, R. E., Jr. (2003). “Macroeconomic Priorities.” *American Economic Review* 93(1): 1–14.

Young, E. R. (2010). “Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell–Smith Algorithm and Non-Stochastic Simulations.” *Journal of Economic Dynamics and Control* 34(1): 36–41.