

Heterogeneous Agent Macroeconomics

Simulating economies with millions of agents:
dynamic programming, learning, and policy evaluation

Ralph Luetticke

Professor of Economics, University of Tübingen
`ralph.luetticke@uni-tuebingen.de`

Cluster Colloquium “Machine Learning”
University of Tübingen

About me, in one slide

- Macroeconomist: monetary and fiscal policy, inequality, **computational methods**
- PhD Bonn → University College London / Stanford → Tübingen (2022–)
- What I build: **simulators of entire economies** — millions of heterogeneous, optimizing households — fit to micro and macro data
- Two funded agendas, both overlapping with machine learning:
 - ERC Starting Grant *AIRMAC*: aggregate risk $\times$ inequality
 - Thyssen project *HE-HANK*: replace rational expectations with **learning agents**

The problem macroeconomists face

Questions we are asked:

- What if the ECB raises rates by 1pp?
- Stimulus checks in a recession — to whom?
- Did Germany's 2020 VAT cut work?
- Who pays for inflation?

The constraint:

- **No A/B tests** on economies
- One history, no counterfactual
- Correlations break when policy changes

Our answer

Build a structural **simulator** of the economy, fit it to data, run the counterfactual inside it.

The end of the representative-agent paradigm

- The average household does not exist: top 10% own half of all wealth; many households hold **zero liquid savings**
- Responses to policy differ: marginal propensity to consume ≈ 0.2 – 0.5 if liquidity-constrained, ≈ 0 if wealthy
- Recessions redistribute: unemployment risk, asset prices, inflation hit people differently

HANK: Heterogeneous-Agent New Keynesian models

A continuum of households — each with its own income history, wealth, portfolio — interacting through markets. **Aggregate dynamics and distributional questions in one framework.**


A household is a Markov decision process

$$V(s_t, X_t) = \max_{a_t} u(a_t) + \beta \mathbb{E}[V(s_{t+1}, X_{t+1}) \mid s_t, a_t, X_t]$$

- **State** s : liquid wealth, illiquid wealth, stochastic labor income (Markov chain)
- **Actions** a : consume, save, rebalance portfolio, work
- **Aggregate state** X : prices — wages, interest rates, inflation

In ML terms: an **MDP with known reward and known transition kernel** — dynamic programming, not model-free RL. Solved globally on grids: **borrowing constraints are kinks**, and the kinks are the economics.

The economy is a mean-field game



Equilibrium = **fixed point**. Formally: a **mean-field game with common noise**.

The curse: the distribution is a state variable

Rational expectations: agents forecast prices — but prices are functions of the distribution:

$$p_t = P^*(G_t, z_t), \quad G_{t+1} = A_\pi^\top G_t$$

- Prices alone are **not Markov**; the Bellman equation is $V(s, G, z)$ — the *Master equation*, infinite-dimensional
- Discretized: my two-asset model has **~660,000 state variables**

The methods question of the field

Solve, simulate, and *estimate* a fixed point over a 660k-dimensional system — fast enough for a Bayesian likelihood loop.

My two research agendas

Agenda I — Solve and estimate the simulator (ERC *AIRMAC*)

Compress the 660k-dim system, differentiate it, estimate it on data — then push beyond linearization to study **aggregate risk $\times$ inequality**.

Agenda II — Fix the expectations assumption (Thyssen *HE-HANK*)

Rational expectations is the source of the curse — and partially rejected by the data. Replace it with **measured, heterogeneous learning rules** — more realistic and cheaper at once.

Plan: methods and applications of Agenda I first, then Agenda II.

Agenda I, step 1: compress the model

Bayer & Luetticke (*Quantitative Economics*, 2020) — now a field standard (“BASEforHANK”)

- ① Solve the **stationary** economy exactly, once
- ② Fluctuations are smooth deviations $\Rightarrow$ compress:
 - Value functions: **discrete cosine transform** — JPEG for policy functions
 - Distributions: **marginals + copula**
- ③ 660,000 $\rightarrow \sim 3,000$ coordinates, keeping 99.99% of the variation
- ④ Eigendecompose fitted co-movement $\Rightarrow \sim 400$ **states**: a learned latent space

Then **autodiff** the equilibrium system (Julia) $\Rightarrow$ a structured, differentiable local surrogate.

Agenda I, step 2: fit the simulator to data

Bayer, Born & Luetticke (*American Economic Review*, 2024)

- Compressed model = linear-Gaussian state space $\Rightarrow$ **Kalman-filter likelihood** of 70 years of US data, incl. inequality series
- Full **Bayesian estimation** (MCMC), ~ 50 structural parameters

Findings:

- **Income risk is itself a shock**: risk rises in recessions $\rightarrow$ precautionary saving $\rightarrow$ deeper recessions (Bayer et al., *Econometrica*, 2019)
- Business cycles *cause* inequality dynamics
- Portfolio liquidity shapes monetary transmission (*AEJ: Macro*, 2021)

Agenda I, frontier: risk $\times$ inequality (ERC AIRMAC)

Linearization $\Rightarrow$ **certainty equivalence**: agents ignore aggregate risk itself. But:

- Exposure to aggregate risk is heterogeneous (banker vs. civil servant); income risk rises in recessions
- Questions: cost of aggregate risk, and for whom? Do recessions deepen with inequality? Fiscal vs. monetary stabilization — who prefers which?

Method goal: higher-order perturbation of the 660k-dim system

2nd/3rd-order expansions break certainty equivalence. Derivative tensors explode $\Rightarrow$ **sparsity** by model structure + **autodiff** + the 400-dim latent space.

Byproduct: fast **differentiable simulators** of mean-field economies up to third order.

Agenda I applied: banks, spreads, who bears volatility

Faccini, Lee, Luetticke, Ravn & Renkin (*American Economic Review*, 2026)

- Danish administrative data: **consumer credit spreads** vary over time and across households
- Rising spreads in recessions $\Rightarrow$ constrained households cut hard: **countercyclical marginal propensity to consume**
- HANK + banking sector: frictional finance amplifies shocks, concentrates them on the poor

Macro vs. micro volatility trade-off

Safer, less leveraged banks stabilize the *aggregate* — but shift volatility onto *households*, with welfare losses at the bottom.

Agenda II: the expectations problem

Rational expectations assumes every household forecasts prices by **solving the model** — distribution included. Nobody does that.

- Survey data: household inflation expectations dispersed over **20+ pp** (professionals: clustered at 2%), biased upward, updated sluggishly and heterogeneously
- And the expectations assumption is what *creates* the curse: G is a state only to keep beliefs model-consistent

Moll's "free lunch" (*The Trouble with Rational Expectations*, EJ 2026)

Replace rational expectations with *measured* learning rules $\Rightarrow$ the fixed point disappears, the distribution drops out. **More realistic and cheaper at the same time.**

Agenda II applied: Germany's 2020 VAT cut

Bachmann, Born, Goldfayn-Frank, Kocharkov, Luetticke & Weber (*Review of Economic Studies*, forthcoming)

- July–Dec 2020: VAT temporarily cut by 3pp — buy now, prices rise later
- No RCT? Use **beliefs as treatment assignment**: only households who *perceived* the cut in prices had an incentive to buy early (surveys + scanner data)
- Perceiving households: **~37% more durable spending**; effect rises with durability and toward the deadline
- HANK simulator for the aggregate: **+4.4% consumption** on impact, cumulative multiplier ≈ 1.8 — beats a comparable rate cut
- It worked because it was **salient, comprehensible, actionable**

Agenda II applied: measure expectations, don't assume them

Luetticke, Roth, Wiederholt & Wohlfart (2025/26)

- **Lucas critique** (1976): correlations fail for policy evaluation because policy changes expectation formation — **distribution shift**, 40 years early
- Our fix: **elicit expectations under counterfactual policy scenarios** in large surveys (“if the ECB raised rates to x...”)
- Feed measured, individual-level expectations into the heterogeneous-agent model — no assumption on expectation formation
- Finding: expectation responses are **state-dependent** $\Rightarrow$ so are the effects of monetary policy

Agenda II data: latent types of learners

ECB Consumer Expectations Survey: monthly panel, ~ 19 k euro-area households (with G. Müller and F. Wilczek)

- Updating: revision of 12-month inflation forecast on revision of *perceived* inflation
- **Finite mixture regression, fit by EM**; group number by information criterion

	type 1	type 2	type 3	type 4
updating gain	0.04	0.16	0.53	0.85
share	26%	38%	21%	14%

- Four **learner types**: from “ignore the news” to “extrapolate one-for-one”; long-run anchoring is a second, nearly orthogonal trait
- Observables explain $\sim 2\%$ — the heterogeneity is **genuinely latent**

Agenda II theory: households as learning agents

Online learning replaces the fixed point (constant gain = recursive least squares with forgetting):

$$\mathbb{E}_{it}[\pi] = \mathbb{E}_{i,t-1}[\pi] + g_i(\hat{\pi}_{it} - \mathbb{E}_{i,t-1}[\pi])$$

- Learning rates g_i : **directly estimated** — the mixture slopes
- $\hat{\pi}_{it}$: what household i *observes*. Our twist:

Same algorithm, different data

Different **consumption baskets** $\Rightarrow$ different observed prices $\Rightarrow$ heterogeneous beliefs *emerge*.
Formally: a partially observable mean-field game with **type-specific observation functions**.

- Feedback: policy $\rightarrow$ relative prices $\rightarrow$ observations $\rightarrow$ beliefs $\rightarrow$ demand $\rightarrow$ prices

Agenda II solution: how RL can solve these models

Recent work by Moll and coauthors — not mine, but the tool we will use. Two uses of RL in macro:

- 1 **RL as an equilibrium computation device** — exists: *Structural RL* (Yang–Wang–Schaab–Moll) and *RSPG* (Wibault et al., **ICML 2026 spotlight**). Slogan: RL about prices, dynamic programming for individual states. Fast: HANK in ~ 4 min on one A100.
- 2 **RL as a model of human learning** — Moll's open question, and **our project**: requires online updating and *experience-weighted sampling* — which is exactly what **consumption baskets** deliver.

Validation loop: simulate survey data from the model, run the identical mixture EM — how much learner-type dispersion emerges from experience alone?

Macroeconomics

household problem (Bellman)
heterogeneous-agent equilibrium
rational expectations
state-space reduction (DCT, copula, factors)
perturbation + autodiff
Bayesian estimation of HANK
adaptive learning, constant gain
expectation survey types
Moll's structural RL
narratives in news and surveys

Machine learning

MDP, dynamic programming
mean-field game, common noise
consistency fixed point / master equation
model compression, learned latent space
differentiable simulation
state-space inference
online learning, RLS with forgetting
latent-class models, EM
policy gradient, structured variance reduction
NLP / LLM text-as-data

We are hiring: PhD position in the HE-HANK project

Fully funded PhD position (3 years, start autumn 2026), U Tübingen

Project: *Heterogeneous Expectations in HANK Models* (Thyssen Foundation) — joint supervision by Ralph Luetticke and Gernot Müller.

What the person will actually do:

- **Latent-variable models** on large expectation panels (mixture/EM and beyond: state-dependent types)
- **Structural reinforcement learning** for mean-field economies
- **LLM-based text analysis**: economic narratives from media and open-ended survey text

Who we look for: **MSc in ML / CS / statistics / quantitative economics** — no economics background required; we teach the economics, you bring the methods.

Interested, or know someone? `ralph.luetticke@uni-tuebingen.de`

- ➊ Macro is a **computational simulation science**: millions of MDPs coupled through markets, fit to data
- ➋ The hard part: a **fixed point over an infinite-dimensional state**; compression + differentiation made these models estimable
- ➌ The simulators answer real questions: who bears volatility, whether a VAT cut works (it did), how policy redistributes
- ➍ The frontier: **replace the fixed point with learning** — measured, heterogeneous rules; RL is the natural solver
- ➎ The overlap with this cluster is concrete and funded — **let's collaborate**

`ralph.luetticke@uni-tuebingen.de` · `www.ralphluetticke.com`

Backup: the estimated model in one equation block

Household i , individual states $s = (b, k, h)$ (liquid, illiquid, human capital):

$$V(s; X) = \max_{c, b', k'} u(c, n) + \beta \mathbb{E}[V(s'; X')] \\ \text{s.t. } c + b' + q k' = (1 + r^b) b + (r^k + q) k + w h n - T, \quad b' \geq \underline{b},$$

with occasional (Calvo-type) opportunities to adjust the illiquid asset. Firms set prices/wages subject to adjustment frictions (New Keynesian Phillips curves); central bank follows a rate rule; government issues debt, taxes, transfers.

Estimation: perturbation around the stationary equilibrium; compressed state ~ 900 ; Kalman-filter likelihood on US aggregates *and* distributional series; random-walk Metropolis–Hastings.

Backup: the two-gain learning block and its measurement

Perception (basket exposure): $\hat{\pi}_{it} = \alpha \tilde{\pi}_{it}^b + (1 - \alpha)z_t$, with $\tilde{\pi}^b$ frequency-weighted own-basket inflation.

$$\text{short run: } \mathbb{E}_{it}[\pi_{12m}] = \mathbb{E}_{i,t-1}[\pi] + g_i (\hat{\pi}_{it} - \mathbb{E}_{i,t-1}[\pi])$$

$$\text{long run: } \ell_{it} = \ell_{i,t-1} + \lambda_i (\mathbb{E}_{it}[\pi] - \ell_{i,t-1})$$

Mixture slopes estimate g (short-run gain) and λ (anchoring gain) type by type; rank correlation 0.20 $\Rightarrow$ two-dimensional trait distribution. Validation: identical EM pipeline on model-simulated panels (indirect inference).

Backup: the Lucas critique in ML language

Performative prediction (Perdomo et al., ICML 2020) = the Lucas critique

Deployed model θ shifts the data distribution: $D(\theta)$. *Performative stability* — model optimal against the distribution it induces — is literally a **rational expectations equilibrium**; their “repeated risk minimization” is what economics calls **adaptive learning** (Marcet–Sargent, 1980s), with the same convergence question.

Off-policy evaluation = policy evaluation from historical data

Evaluating a new policy on data generated under an old one works only with *overlap* (importance sampling, doubly robust). A **regime change destroys overlap by construction** $\Rightarrow$ only a model of the environment remains: model-based RL = the structural approach.

Related: Goodhart’s law (1975) = reward hacking at the level of an economy (the 1970s Phillips curve); strategic classification and recommender feedback loops = Lucas at product scale; invariant causal prediction = the search for “deep parameters”.

Backup: structural policy gradient objective

Simulate N price trajectories by rolling out the economy (market clearing inside the rollout), then maximize

$$\hat{v}_\pi = \frac{1}{N} \sum_{n=1}^N \sum_{t=0}^T \beta^t A_{\pi,0 \rightarrow t}^n u(c(z_t^n, p_t^n)),$$

where $A_{\pi,0 \rightarrow t}$ chains the *known* individual transition matrices — expectations over idiosyncratic risk are exact; sampling only over aggregate shocks. Gradient ascent on grid-based or neural policies. Equilibrium concept: restricted perceptions / self-confirming.