

What is structural estimation?

Applied Microeconomics

Hans-Martin von Gaudecker & Florian Zimmermann

The goal

Reduced-form work identifies an **effect** (a LATE, an ATT, a regression coefficient) under as few assumptions as possible.

Structural work has a different target: the **parameters of an economic model of choice**.

- preferences (risk aversion, discounting, ...)
- constraints (budget sets, borrowing limits, ...)
- beliefs and information

Once you have those, you can ask what the agent would do *under a policy you have never observed*.

Why it matters: risk and timing have a price

A reduced-form estimate gives you a *mean* effect. If markets are incomplete and agents are prudent, the mean is not the whole story:

- a higher-variance income stream is *worth less* per dollar of expected income
- a steeper income profile is *worth less* if you cannot borrow against it
- taxes and transfers attenuate the gross return but also insure it

So the *welfare-relevant* return can differ sharply from the gross return. Pinning down that gap needs a life-cycle model with precautionary saving, and the model is only as credible as the moments you feed it.

A different bargain with the data

	Reduced-form / IV	Structural
Identified by	a research design (an instrument, a discontinuity)	a model + functional-form & distributional assumptions
Delivers	one credibly-identified effect	a full set of parameters → counterfactuals, welfare
Robust to	mis-specification of the model	weak instruments, design failures
Vulnerable to	"that is only one number, for one margin"	"that answer is only as good as the assumptions"

Neither dominates. They answer **different questions**, and trade different assumptions to get there.

What a structural model usually needs

To write down "what the agent would choose", you typically need:

- a **dynamic decision problem**: savings, labour supply, schooling, and so on, where today's choice changes tomorrow's state. Solved by **dynamic programming**.
- an explicit stand on **expectations**: what does the agent know, and what do they believe about the future? Rational expectations is the common default, but it *is* an assumption.
- distributional assumptions on the **unobservables** (shocks, latent types) the data never show you directly.

None of these are testable from the reduced-form moments alone. That is the cost.

How the parameters get pinned down

Given the model, the parameters are chosen so that the model **reproduces features of the data**:

- maximum likelihood: match the full likelihood
- GMM / simulated method of moments: match selected moments
- minimum distance: match a vector of reduced-form estimates

The recurring pattern: **reduced-form estimates become the targets**, and the structural step asks *which parameters could have generated them*.

The honest trade-off

- The standard criticism, *"more assumptions, less robust"*, is true but incomplete.
- The objects structural work is after (willingness to pay, the welfare cost of risk, behaviour under a counterfactual policy) are **simply not in the reduced-form moments**.
- So the choice is not "robust answer vs fragile answer".
- It is: answer a different question, under stated assumptions — or do not answer it at all.
- This week's paper sits exactly on that interface: a careful reduced-form step, feeding a structural model. Keep this distinction in mind as you read it.