A Trilemma for Asset Demand Estimation

William Fuchs   Satoshi Fukuda   Daniel Neuhann

2026 Frontiers in Finance Conference

Motivation

Influential literature aiming to estimate **demand functions** for financial assets from portfolio data.

- **Promise**: better understand price formation, policy transmission, counterfactuals.
- Unlike micro-structure, focus on lower frequencies in which markets have time to adjust.

Approach follows the **Marshallian paradigm**: recover properties of asset demand from supply shocks,

$$\mathcal{E}_{jk} = -\frac{\partial a_j(\vec{a}_{-j}, \vec{p})}{\partial p_k} \times \frac{p_k}{a_j(\vec{a}_{-j}, \vec{p})}.$$

Background

Koijen and Yogo (2019) use model and claim that elasticities are **much lower** than we thought.

FFN (2025) showed that their assumed structure cannot capture "demand complementarities": the standard finance notion that the value of one asset depends on holdings of other assets.

As a result, estimated elasticities can be low even if the truth is that they are very large.

Takeaway Graph from FFN 2025 (log scale)

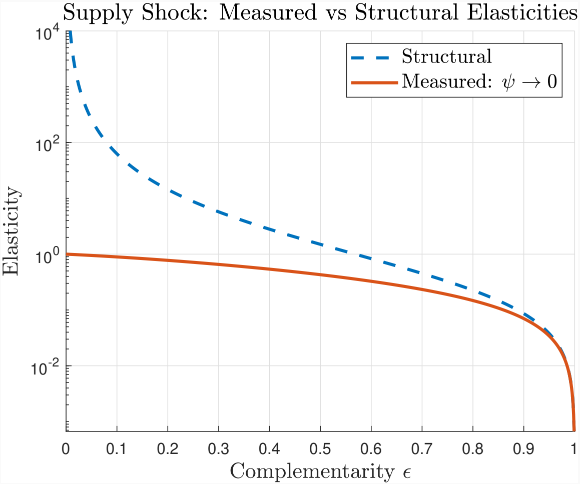

The KY19 measured elasticity is always low; even if the true structural elasticity is near infinite.

Today: broader questions

1. Is asset-level demand estimation the right paradigm for empirical analysis of portfolio data?

2. When can asset demand be identified from observational data **without unverifiable restrictions**?

Finance presents distinct issues from "standard" demand estimation

1. Cross-asset linkages through portfolio choice: investors care about total payoffs, not assets directly.
2. Prices bound together by the specific geometry of no arbitrage.
3. Fundamental asset attributes – their future *payoffs* – are unobserved and variable.

We establish limits to demand identification in asset markets

It is generically impossible to jointly maintain that:

1. Investors have preferences over payoffs

2. Prices satisfy no arbitrage,

3. Structural asset demand functions are identifiable without strong and **untestable** assumptions.

We establish limits to demand identification in asset markets

It is generically impossible to jointly maintain that:

1. Investors have preferences over payoffs

2. Prices satisfy no arbitrage,

3. Structural asset demand functions are identifiable without strong and **untestable** assumptions.

$\Rightarrow$ Estimated demand slopes **always reflect** – rather than validate – a-priori theoretical assumptions.

The Marshallian graph: ideal versus arbitrage-free asset markets

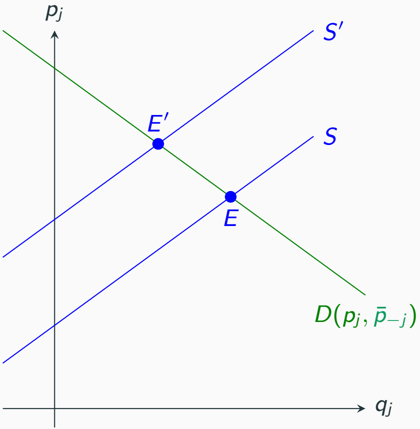

(a) Ideal: Identification given stable demand for asset j.

The Marshallian graph: ideal versus arbitrage-free asset markets

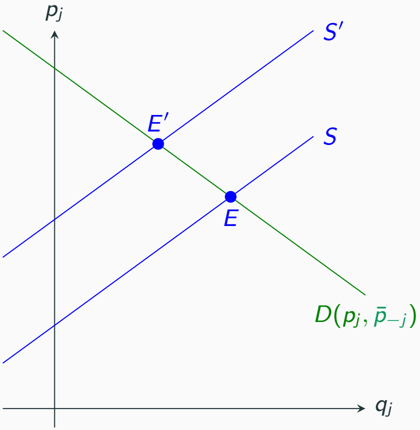

(a) Ideal: Identification given stable demand for asset j.

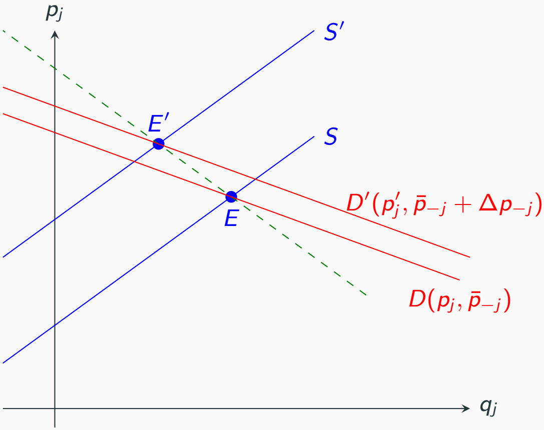

(b) Identification failure with endogenous demand shifts.

Conclusion: asset demand functions are fragile, **theory-dependent objects**, not portable moments.

Related literature

Equilibrium effects from supply shocks

Shleifer (1986), Harris–Gurel (1986), ...

Estimation of structural asset demand elasticity for policy / counterfactual

Koijen–Yogo (2019) & (2021) and many others.

Most closely related

FFN25: low estimates can be due to missing demand complementarities—even with ideal instruments.

HHHKL (2025): aim to recover relative elasticities from supply shocks "without" a structural model.

We show this requires implicit structure and strong assumptions on Y.

Binsbergen–David–Opp (2025); He–Kondor–Li (2025): Biases if shocks used for estimation also change Y

Framework

Asset pricing in an endowment economy

- Two dates, $t = 0, 1$. At date 1, one of Z states is realized.

- Investors choose among J assets. Y is the $J \times Z$ payoff matrix, $y_j(z) \geq 0$ is j's payoff in state z.
- **Markets can be complete or incomplete.**

- p is the vector of asset prices, q is the vector of state prices (need not be unique).

- Potentially heterogeneous investors indexed by i.
- Investor i has endowment e_j^i of asset j. Aggregate endowment of asset j is $E_j = \sum_i e_j^i$.

- Equilibrium prices determined by market clearing.

Why no arbitrage? Consistent pricing of portfolios

A defining feature of portfolio choice: investors can flexibly **bundle and unbundle** assets.

Example. Two assets with payoffs $[1, 1]$ and $[1, 0]$ can be combined into $[0, 1]$—or any payoff in $\mathbb{R}^2$.

Continuous choice over assets implies a *continuum* of feasible payoff vectors.

Why no arbitrage? Consistent pricing of portfolios

A defining feature of portfolio choice: investors can flexibly **bundle and unbundle** assets.

Example. Two assets with payoffs $[1, 1]$ and $[1, 0]$ can be combined into $[0, 1]$—or any payoff in $\mathbb{R}^2$. Continuous choice over assets implies a *continuum* of feasible payoff vectors.

Problem. The mapping from payoff vectors to portfolios depends on the unobserved Y. To infer preferences from portfolio holdings, we need a pricing rule that covers *all* feasible portfolios consistently.

No arbitrage provides exactly this: a consistent pricing system for all tradeable payoff bundles.

Theorem (Fundamental Theorem of Asset Pricing)
There is no arbitrage if and only if there exist state prices $q \in \mathbb{R}^Z_{++}$ such that

$$p = Yq.$$

Why no arbitrage? Three additional roles

1. No arbitrage ensures the **existence of well-behaved optimal portfolios**
2. **Dimension reduction** using factor structures to handle many assets.
3. **Consistency of counterfactuals** under alternative price vectors.

Discarding no arbitrage likely to invalidate the basic premises of much of asset demand analysis.

Step 1: Properties of Asset Demand

Start from the canonical investor decision problem

$$U^i = \sup_{a^i \in \Phi^i(p)} \quad u^i(c_0^i) + \delta^i \sum_{z=1}^{Z} \pi_z u^i(c_z^i) \qquad \text{(Utility over payoffs + constraints)}$$

$$\text{s.t.} \quad c_0^i = e_0^i - \sum_{j=1}^{J} p_j (a_j^i - e_j^i) \qquad \text{(Budget constraint)}$$

$$c_z^i = \sum_{j=1}^{J} y_j(z) a_j^i + w^i(z) \quad \text{for all } z \qquad \text{(Endowments= demand shifters)}$$

Start from the canonical investor decision problem

$$U^i = \sup_{a^i \in \Phi^i(p)} \quad u^i(c_0^i) + \delta^i \sum_{z=1}^{Z} \pi_z u^i(c_z^i) \qquad \text{(Utility over payoffs + constraints)}$$

$$\text{s.t.} \quad c_0^i = e_0^i - \sum_{j=1}^{J} p_j (a_j^i - e_j^i) \qquad \text{(Budget constraint)}$$

$$c_z^i = \sum_{j=1}^{J} y_j(z) a_j^i + w^i(z) \quad \text{for all } z \qquad \text{(Endowments= demand shifters)}$$

Standard ID problem: use observed behavior to recover utility parameters and demand shifters. But:

1. Indirect demand: preferences pertain to *payoffs*, observations to *assets*. But Y and π are latent.
2. Cross-asset linkages: asset-specific demand $a_j^i(\bar{p})$ generically depends on the **entire vector** of prices.

What do these problems mean for asset demand? A decomposition

By the budget constraint, the consumption process induced by a portfolio a is

$$c = Y^{\mathrm{T}} a + w^i.$$

Hence a feasible target consumption vector c^* requires the portfolio:

$$a^*(c^*) = (Y^+)^{\mathrm{T}}(c^* - w^i)$$

By no arbitrage, $p = Yq$. The chain rule then yields

$$\frac{\partial a^*(c^*)}{\partial p^{\mathrm{T}}} = (Y^+)^{\mathrm{T}} \frac{\partial c^*}{\partial p^{\mathrm{T}}} = (Y^+)^{\mathrm{T}} \underbrace{\left(\frac{\partial c^*}{\partial q^{\mathrm{T}}}\right)}_{\equiv \mathcal{D}^i} \underbrace{\left(\frac{\partial q^{\mathrm{T}}}{\partial p^{\mathrm{T}}}\right)}_{= Y^+ \text{by NA}}.$$

Why is this decomposition useful?

$$\boxed{\underbrace{\frac{\partial a^i(p, Y, \text{shifters})}{\partial p}}_{\text{Asset demand}} = \underbrace{(Y^+)^T}_{\text{Mapping of } c \text{ into assets}} \left[\underbrace{\frac{\partial c^i(q, \text{shifters})}{\partial q}}_{\text{Fundamental demand } \mathcal{D}^i} \right] \underbrace{Y^+}_{\text{Mapping of asset prices to state prices}} .}$$

Fundamental consumption demand $\mathcal{D}^i \equiv -\frac{\partial c^i}{\partial q^{\mathrm{T}}}$ is **recognizable as a structural demand function.**

- It reflects internal preference parameters, not "external" characteristics of the choice set.
- It is invariant with respect to standard perturbations of the environment.

  In particular, depends on Y only through the asset span.

We can then understand properties of asset demand by analyzing the mapping through Y^+.

Asset demand is **derived** from consumption demand using a **latent mapping** based on Y^+.

1. **Non-separability.** The demand for any asset is determined jointly by the payoffs of *all* assets.
   Indeed, whether assets are substitutes or complements depends on other assets in the choice set.

2. **Unobservability.** The mapping that links preferences and behavior is itself unobserved.
   $\Rightarrow$ Fundamental demand $\mathcal{D}^i$ is fundamentally unidentified from portfolio and price data.

Proposition: Y^+ is fundamentally unobservable

Proposition (Non-identification of Y^+)

Consider any finite sample of realized payoffs. Then there exist arbitrarily many candidate payoff matrices Y that are observationally equivalent but have different generalized inverses.

Proof sketch. Take a candidate payoff matrix. Add a state z^* that never realizes in the sample. This changes Y^+ while leaving all historical returns identical. □

Implications: Non-identification of preferences and non-structural asset demand

1. Given **any** $\mathcal{A}^i$, the consumption demand $\mathcal{D}^i$ is generically **unidentified** absent knowledge of Y^+.

   That is, **even with perfect shocks** you cannot identify preferences unless you know payoffs.

2. If preferences are structural, **asset demand is not unless payoffs/ beliefs never change.**
   - Payoff or belief revisions are **routine** in normal market functioning.
   - Since payoffs include resale prices, essentially all standard shocks shift p *and* Y, changing the system.

     See also Binsbergen-David-Opp and He-Kondor-Li.
   - Policy interventions (QE, FX markets) are *designed* to shift Y.

We conclude: Asset demand should not be seen as structural in essentially all settings of interest.

Step 2:

Identification from Supply Shocks

So I cannot learn preferences, but can I learn something?

Say we give up on fundamental consumption demand $\mathcal{D}^i$.

Can we at least identify asset demand $\mathcal{A}^i$ as a "sufficient statistic"?

Now investigate this using: (i) perfectly exogenous asset-level supply shocks, and (ii) many shocks.

Answer: no, unless you have Arrow securities or payoffs do not change.

Case 1: A single asset level supply shock

Suppose we wanted to measure demand for some specific asset.

Canonical approach: trace out asset-level demand function using asset-level supply shock.

Case 1: A single asset level supply shock

Suppose we wanted to measure demand for some specific asset.

Canonical approach: trace out asset-level demand function using asset-level supply shock.

Since asset demand depends on all prices, need **clean variation in a single price.**

Case 1: A single asset level supply shock

Suppose we wanted to measure demand for some specific asset.

Canonical approach: trace out asset-level demand function using asset-level supply shock.

Since asset demand depends on all prices, need **clean variation in a single price.**

Can we get such an ideal experiment? Not under no arbitrage...

Understanding the ideal experiment

We can ask: what are the state price changes induced by a pure price shock to asset j?

No arbitrage says $p = Yq$. So, state prices implied by observed asset prices are:

$$\boxed{q = Y^{+}p.}$$

Then the ideal experiment has a specific state-price representation:

$$\boxed{\Delta \mathbf{q_j}^{\text{ideal}} \equiv \frac{\partial q}{\partial p_j} = Y^{+} v_j.}$$

That is, ideal state price variation is proportional to the **inverse payoff matrix**.

What do supply shocks deliver?

Now ask: do supply shocks deliver (something close to) the ideal experiment?

Approach: **assume we have a perfect supply shock to some asset** j and then verify.

Minimal assumptions on equilibrium pricing

Definition (Downward-sloping consumption demand)

There is downward-sloping consumption demand if there exists a $Z \times Z$ matrix V with **strictly positive diagonal** such that

$$\Delta \mathbf{q}_j^{\mathrm{supply}} \equiv \frac{\partial q}{\partial E_j} = -V y_j^{\mathrm{T}} \qquad \text{for all assets } j.$$

State price changes induced by supply shocks are proportional to the **payoff matrix itself**.

NB: To provide favorable conditions for identification, further assume V is diagonal for now.

Supply shocks produce misaligned state price variation

We will show two alignment conditions both fail generically:

Condition 1 (Identical variation) $\Delta \mathbf{q}_j^{\text{ideal}} = k_j \Delta \mathbf{q}_j^{\text{supply}}$ for some scalar k_j.

Condition 2 (Same sign) $\text{sign}(\Delta \mathbf{q}_j^{\text{ideal}}) = \text{sign}(\Delta \mathbf{q}_j^{\text{supply}})$ element by element.

Supply shocks produce misaligned state price variation

We will show two alignment conditions both fail generically:

Condition 1 (Identical variation) $\Delta \mathbf{q}_j^{\text{ideal}} = k_j \Delta \mathbf{q}_j^{\text{supply}}$ for some scalar k_j.

Condition 2 (Same sign) $\text{sign}(\Delta \mathbf{q}_j^{\text{ideal}}) = \text{sign}(\Delta \mathbf{q}_j^{\text{supply}})$ element by element.

Definition. Assets j and j' **have overlapping payoffs** if $\exists$ a state z such that $y_j(z) > 0$ and $y_{j'}(z) > 0$.

Supply shocks produce misaligned state price variation

Theorem (Misaligned Price Variation). If Conditions 1 or 2 are satisfied, then YY^{T} is diagonal, and:

(i) **there are no assets with overlapping payoffs.**

(ii) If markets are complete, Y **is diagonal up to permutations.**

Proof. Plemmons and Cline (PAMS, 1972). Basic problem: $Y \neq Y^{+}$.

Strengthening. If every state has at least two assets with positive payoffs, every row of Y^{+} contains wrong signs—so misalignment is guaranteed for every asset.

Interpretation

1. Asset supply shocks affect behavior because they change the cost of payoffs (state prices).
2. Under no arbitrage, supply shocks affect the price of other assets with overlapping payoffs.
3. These cross-asset restrictions are generically misaligned with those required by the ideal experiment.
4. Only exception: Arrow securities, which eliminate any cross-asset linkages.

No overlapping payoffs is much stronger than orthogonal payoff distributions.

IO link: Arrow securities eliminate the distinction between asset demand and consumption demand.

Of course, no payoff overlap is unrealistic for essentially all settings of interest.

Case 2: How about multiple supply shocks?

Can we overcome the misalignment problem using many independent supply shocks?

Case 2: How about multiple supply shocks?

Can we overcome the misalignment problem using many independent supply shocks?

Proposition (Necessary conditions for identification of $\mathcal{A}^i$)

Generically, $\mathcal{A}^i$ is identified only if the econometrician:

1. Observes $K = J$ **linearly independent** experiments; **and**
2. Maintains the assumption that Y^+ is **fixed** across all experiments.

These conditions are hard to satisfy...

The rank condition ($K = J$) is standard but stringent with many assets.

The fixed-Y^+ condition is the crux: it is unverifiable and forbids revisions in beliefs or resale prices.

- But: payoff revisions occur in regular market functioning.
- And: shock validity requires not just exogenous variation in prices today, but no effect on payoffs—ruling out central bank interventions, index inclusions, and most policy experiments.

We conclude: identification relies on unrealistic assumptions that cannot be tested.

NB: over-identification fails because the system is non-invertible and the “structural error” cannot be concentrated out.

Step 3: Mapping Y^+ is Ill-behaved

OK, but don't we know something about payoffs?

You may think: yes, Y is unknown, but it's not *that* unknown.

- We know about factor structures, estimated covariance matrices,. . .

Are statistical assumptions on Y enough to ensure well-behaved Y^+?

Answer from random matrix theory: no, not really.

- Assume factor structure of Y, study asymptotic behavior of Y^+.
- Small sample Monte Carlo analysis.

Theorem (Sign Instability of Y^+). Under a factor structure $y_{j,z} = \alpha_j + \beta_j f_z + \varepsilon_{j,z}$ and mild regularity conditions, for almost every realization of factor loadings:

(i) **Individual coin flip.** For each fixed asset j and state z, the sign of an element is a **coin flip:**

$$\lim_{J\to\infty} P\big((Y^+)_{z,j} > 0\big) = \frac{1}{2}.$$

(ii) **Knowledge of factors is insufficient.** Let Y and $\tilde{Y}$ share factor loadings and factor realizations but have independent idio. shocks. Then sign-determining variables are asymptotically independent:

$$q(J) \equiv \frac{1}{JZ} \sum_{j,z} \mathbf{1}\Big(\text{sign}(Y^+)_{z,j} = \text{sign}(\tilde{Y}^+)_{z,j}\Big) \to \frac{1}{2}.$$

Monte Carlo: sign instability converges fast

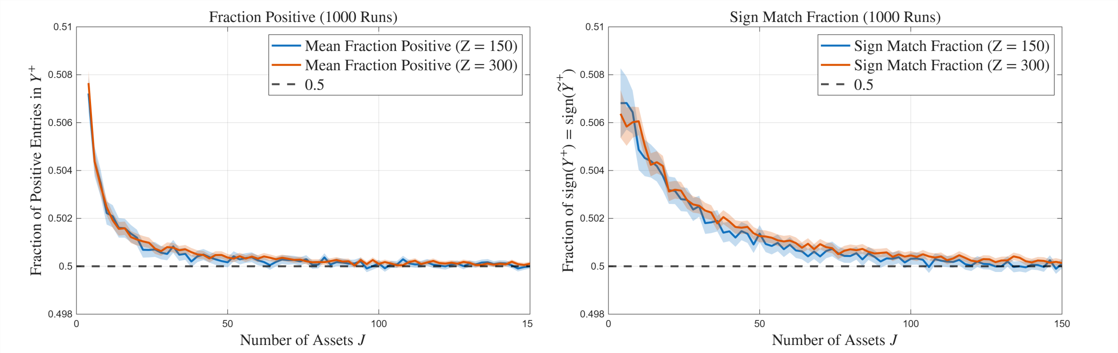

Left: fraction of positive entries in Y^+ converges to $\frac{1}{2}$ quickly.

Right: sign-match frequency between Y^+ and $\tilde{Y}^+$ (same factor model, independent idio. shocks).

S&P 500 data (20×20 submatrices): 50.58% wrong signs, magnitudes equal.

The Trilemma and its implications

The Trilemma

Theorem (Trilemma). Given observational data on portfolios, asset prices, and payoffs, one cannot jointly maintain (i) no-arbitrage asset pricing, (ii) investor preferences over payoffs, and (iii) model-free identification of structural asset demand functions.

$\Rightarrow$ Asset demand elasticities are not model-free moments: they reflect *theoretical assumptions*.

These issues are distinct to finance: no arbitrage, latent mappings, endogenous "attributes."

Implication 1: Theoretical restrictions

One can always achieve parametric identification using a-priori restrictions.

- E.g., KY19's logit demand and HHHKL25's homogeneous substitution conditional on observables.

Our analysis reveals that these assumptions are perhaps stronger than appreciated.

E.g., need payoffs and resale prices to remain fixed even when you shock "factor portfolios."

Since the demand is non-invertible, you cannot rely on overidentification tests to evaluate structure.

Implication 2. Additional data

Allen, Kastl and Wittwer (2025) use bid-level data to avoid needing price instruments (but rely on a structural model).

You could use beliefs data or purer "price shocks" (as opposed to supply shocks), if you can find them.

In any case: a priori restrictions on the latent mapping Y^+ is unavoidable and drives results.

Implication 3. Target different objects

Individual demand functions may be unstable, but maybe certain summary measures are not.