A Note on Koijen and Yogo's Cross-sectional Demand Estimator

William Fuchs* Satoshi Fukuda† Daniel Neuhann‡

June 9, 2025

Abstract

Fuchs, Fukuda, and Neuhann (2025) [FFN25] present a model to argue that logit asset demand systems produce biased estimates of demand elasticities in the presence of cross-asset demand complementarities. However, Koijen and Yogo (2025) contend that the cross-sectional estimator from Koijen and Yogo (2019) correctly estimates the structural elasticity even in FFN25's model. We show that this conclusion rests on a knife-edge parameter restriction that ensures perfect symmetry between assets. For *arbitrary violations of this restriction*, their cross-sectional estimator (i) has the wrong sign, (ii) does not vary with the true elasticity, and (iii) can exhibit unbounded bias. In contrast, the estimator in FFN25 (i) has the right sign, (ii) is monotone in the true elasticity, and (iii) is unbiased in a benchmark with symmetric substitution (as in logit).

*University of Texas at Austin, FTG, and CEPR. william.fuchs@mccombs.utexas.edu.

†Bocconi University and IGIER. satoshi.fukuda@unibocconi.it.

‡University of Texas at Austin and FTG. daniel.neuhann@mccombs.utexas.edu.

1 Introduction

Koijen and Yogo (2019) [hereafter, KY19] introduced the influential demand system approach to asset pricing, whereby structural models are combined with portfolio data to infer demand elasticities for financial assets. The resulting estimates are then used to inform our understanding of asset price determination and policy counterfactuals. A striking claim from this literature is that estimated demand elasticities are near one even though classical models suggest that they should be three orders of magnitude higher.

Fuchs, Fukuda, and Neuhann (2025) [hereafter, FFN25] argues that low estimated elasticities may instead be due to estimation biases driven by cross-asset spillovers. When investors have portfolio-level preferences, shocks to individual assets can spill over to the demand and prices for other assets. This in turn creates omitted variable bias in the class of asset-level demand regressions implied by KY19-style demand systems.

Koijen and Yogo (2025) [hereafter, KY25] disagrees with FFN25. In particular, they argue that the cross-sectional estimation strategy proposed in KY19 is *not* susceptible to spillover biases even in the model employed by FFN25. Their abstract states that

> [...] FFN25 use an incorrect estimator for their central claim that "measured elasticities are near one even if true elasticities are near infinite." The cross-sectional instrumental variables estimator correctly identifies the demand elasticities in KY19 and FFN25's three-asset model.

They arrive at this claim by applying their estimator to FFN25's model under a knife-edge parameter restriction that imposes ex-ante symmetry between assets. This note shows that their claim is *not* valid in FFN25's model outside of this symmetry restriction. In particular, even for *arbitrarily small violations of symmetry*, the KY estimator (i) has the wrong sign, (ii) does not vary with the underlying elasticity, and (iii) can exhibit unbounded bias. Hence even small asymmetries can result in large and systematic biases.

The reason for this knife-edge effect is that their estimator is a ratio of two functions, say F/G. Instead of computing this object directly, KY25 approximates it as $F/G \approx F'/G'$. This approximation is valid if (i) the true model is log-linear (which FFN25 is not), or if F and G are both equal to zero at the initial equilibrium point. This latter condition holds if and only if the two inside assets are perfectly symmetric. For any other parameter values, KY25's approximation approach is inaccurate. In practice, such knife-edge conditions are unlikely to hold, and verifying them is infeasible with finite data.

We end by showing that the within-asset estimator used in FFN25 has more desirable properties: it (i) has the right sign, (ii) is monotone in the underlying elasticity and

(iii) is unbiased in a limiting case where the substitution matrix is symmetric (as in logit). Hence it is useful benchmark to study the scope for spillover bias, as is FFN25's intent.

2 The FFN25 Three-Asset Economy

We use the baseline three-asset economy from FFN25. Uncertainty is represented by two aggregate states $z \in \{1, 2\}$, and a distributional shock $\iota \in \{g, r\}$ which further affects payoffs. The probability of state z is π_z, and the probability of shock r is $\Pr(\iota = r) = \rho$.

There are three assets: a green tree, a red tree, and tree 2. Tree 2 is the "outside asset" with price normalized to $p_2 = 1$. The green and red "inside assets" have prices p_g and p_r, respectively. The payoff matrix is in Table 1. Parameter $\epsilon \in (0, 1]$ determines the degree of complementarity between red and green trees. The two inside assets are perfect substitutes for each other in the limit $\epsilon \to 0$.

	State 1 (π_1)		State 2 ($1 - \pi_1$)
	Green shock ($1 - \rho$)	Red shock (ρ)	
Green tree	$1 + \epsilon$	$1 - \epsilon$	0
Red tree	$1 - \epsilon$	$1 + \epsilon$	0
Tree 2	0	0	1

Table 1: Payoff structure in the FFN25 three-asset economy.

There is a unit mass of identical log-utility investors. They receive identical per-capita endowments $e_2 = 1$, $e_r = \frac{1}{2}$ and $e_g = \frac{1}{2} + \psi$. Supply parameter ψ is used to generate exogenous price variation. As in FFN25 and KY25, we study small shocks ($\psi \approx 0$).

Every investor i chooses asset holdings a_j^i for all assets $j \in \{g, r, 2\}$ to maximize expected utility over consumption. The market-clearing conditions are

$$a_j = e_j \qquad \text{for all } j \in \{g, r, 2\}.$$

Equilibrium. FFN25 characterizes this model in detail. For the purposes of this note, we summarize key properties of equilibrium. We write demand in terms of asset j's portfolio weight $\omega_j^i(p) \equiv \frac{p_j a_j^i}{W^i}$, $p \equiv (p_g, p_r)$ is the price vector, and W^i is the investor's wealth. This leads to the following demand functions and prices.

- Taking prices as given, demand functions for inside assets satisfy

$$\omega_g^i(p_g, p_r) = \pi_1 p_g \cdot \frac{(p_g + p_r)\epsilon^2 - (p_g - p_r) + 2p_r\epsilon(1-2\rho)}{(p_g+p_r)^2\epsilon^2 - (p_g - p_r)^2};$$
$$\omega_r^i(p_g, p_r) = \pi_1 p_r \cdot \frac{(p_g + p_r)\epsilon^2 + (p_g - p_r) - 2p_g\epsilon(1-2\rho)}{(p_g+p_r)^2\epsilon^2 - (p_g - p_r)^2}.$$

- Given a value of the supply shifter ψ, market-clearing prices satisfy

$$p_g(\psi) = \frac{\pi_1}{\pi_2}\frac{1 + (1-2\rho)\epsilon + (1-\epsilon^2)\psi}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)};$$
$$p_r(\psi) = \frac{\pi_1}{\pi_2}\frac{1 - (1-2\rho)\epsilon + \psi\left((1-(1-2\rho)\epsilon)^2 + 4\rho(1-\rho)\epsilon^2\right)}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)}.$$

- The equilibrium portfolio shares associated with market-clearing prices are

$$\omega_g(\psi) = \pi_1\frac{1+2\psi}{2}\cdot\frac{1 + (1-2\rho)\epsilon + (1-\epsilon^2)\psi}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)};$$
$$\omega_r(\psi) = \frac{\pi_1}{2}\cdot\frac{1 - (1-2\rho)\epsilon + \psi\left((1-(1-2\rho)\epsilon)^2 + 4\rho(1-\rho)\epsilon^2\right)}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)}.$$

Estimation target. The goal of KY25's estimation procedure is to identify asset-level demand elasticities. A demand elasticity reflects a thought experiment where an investor responds to a price change for a given asset while all other prices are held fixed. This thought experiment is used to trace out the demand curve for a particular asset. In the model at hand, this *structural* elasticity can be computed in closed form. We focus on the own-price elasticity of the green asset. It is given by

$$\mathcal{E}_{gg} \equiv -\frac{\partial\omega_g}{\partial p_g}\frac{p_g}{\omega_g} = \frac{(1-\epsilon^2)(1-\epsilon(1-2\rho))}{8\rho(1-\rho)\epsilon^2} \geq 0. \tag{1}$$

This elasticity is strictly positive for all $\epsilon < 1$, attains its minimum of zero when $\epsilon = 1$, and diverges to infinity as $\epsilon \to 0$. The key question is whether the KY25 estimator recovers this object (up to a scalar multiplier determined by portfolio shares.)

Parameter restrictions in KY25. KY25 analyzes the special case where the probability of the red state, ρ, is exactly equal to $\frac{1}{2}$. This ensures that red and green assets are *ex-ante symmetric*. We will show that their estimator is unbiased only when this knife-edge restriction holds. To establish this, it is useful to state two implications of this restriction.

Corollary 1 (Equilibrium given $\rho = \frac{1}{2}$) *If $\rho = \frac{1}{2}$, the green own-price elasticity is*

$$\mathcal{E}_{gg}\big|_{\rho=\frac{1}{2}} = \frac{1}{2}\left(\frac{1}{\epsilon^2} - 1\right).$$

Moreover, if $\rho = \frac{1}{2}$, then green and red assets have the same price in the limit as $\psi \to 0$:

$$\lim_{\psi \to 0} p_g(\psi) = \lim_{\psi \to 0} p_r(\psi).$$

Proof. The statements follow directly from the equilibrium properties stated above. ■

3 KY25's Cross-Sectional Estimator

We now describe KY25's cross-sectional estimator and provide conditions under which it is biased. Equation (B11) in KY25 posits a log-linear demand function for asset j, namely

$$\log(\omega_j) = -\beta_0 \log(p_j) + \alpha + \eta_j, \tag{2}$$

where η_j is latent demand.[1] The key parameter determining the demand elasticity is β_0, which is assumed to be the same across all assets. (As KY25 correctly points out, β_0 does not directly correspond to the structural elasticity $\mathcal{E}_{gg}$ derived above. In this particular model, it instead corresponds to $\beta_0/2$ because it must be multiplied by portfolio shares.)

There are two main identification challenges. The first is the standard concern that price p_j may be endogenous to latent demand η_j. To circumvent this, KY25 estimates β_0 using an instrumental variables approach, where z is the vector of asset-level instruments. In the current model, the instrument is the supply shifter ψ, so that $z = (\psi, 0)$.

The second is what FFN25 call spillover bias, whereby η_j in fact includes the prices of other assets, and these are endogenously affected by the supply shock. This is reflected in the true demand functions reported in Section 2, which depend on *both* asset prices simultaneously. KY25 argue that their estimator also addresses this challenge.

The estimator used in KY25 is reported in their equation (B12). It is given by

$$\hat{\beta}_0 = -\frac{\mathrm{Cov}(\log \omega_j(\psi), z_j)}{\mathrm{Cov}(\log p_j(\psi), z_j)}. \tag{3}$$

[1] We write η_j instead of ϵ_j in order to avoid confusion with the payoff process of the FFN25 model.

In the model at hand, this estimator has an *exact* solution. By the definition of covariance,

$$\begin{aligned}\hat{\beta}_{0,\text{exact}} &= -\left((\log p_g(\psi) - \log p_r(\psi))\psi\right)^{-1}(\log(\omega_g(\psi)) - \log(\omega_r(\psi)))\psi \\ &= -\frac{\log(\omega_g(\psi)) - \log(\omega_r(\psi))}{\log(p_g(\psi)) - \log(p_r(\psi))}. \end{aligned} \tag{4}$$

Under the proposed estimator, β_0 is thus identified from the ratio of the log portfolio share difference and the log price difference between red and green assets.

The challenge is that the data reflects an *equilibrium allocation* which jointly depends on all underlying structural own- and cross-elasticities. Hence, *if* the true demand functions include both own and other prices (as is the case in FFN25's model) it is difficult to separately identify own- *and* cross-price elasticities using the cross-section alone.

KY25's Implementation. Despite the identification challenges described above, KY25 shows that their proposed estimator can identify the structural elasticity in the FFN25 model under the parameter restrictions they impose (namely, $\rho = \frac{1}{2}$). Why is this?

The key observation is that KY25 does not actually use the exact estimator to compute β_0. Instead, they approximate asset prices and portfolio shares around $z = 0$. This is shown in their equations (B9) and (B10). Adapted to our notation, they approximate

$$\omega_j \approx \omega_j^* + \left.\frac{d\omega_j}{d\psi}\right|_{\omega_j=\omega_j^*}\psi;$$

$$\log p_j \approx \log p_j^* + \left.\frac{d\log p_j}{d\psi}\right|_{p_j=p_j^*}\psi.$$

These *approximate* solutions are then used to compute the estimator. Since ω_j^* and p_j^* are constants, this yields an approximate value of the estimator in *changes* rather than *levels* of prices and portfolio shares (see the second line of equation (B18)):

$$\begin{aligned}\hat{\beta}_0 \approx \hat{\beta}_{0,\text{approx}} &\equiv -\frac{\text{Cov}\left(\frac{d\log\omega_j}{d\psi}\psi, z_j\right)}{\text{Cov}\left(\frac{d\log p_j}{d\psi}\psi, z_j\right)} \\ &= -\frac{\frac{d(\log(\omega_g(\psi))-\log(\omega_r(\psi)))}{d\psi}}{\frac{d(\log(p_g(\psi))-\log(p_r(\psi)))}{d\psi}}.\end{aligned} \tag{5}$$

While equation (B18) reads as if this solution is exact, it is exact only *conditional* on the approximations in (B9) and (B10).

The approximate estimator in (5) differs from the exact estimator in (4) because it can infer a high demand elasticity from a small quantity change given an even smaller price change. However, the critical question is whether this approximation is *accurate*.

To see this, observe that estimator $\hat{\beta}_0$ is a ratio of two functions, say $\hat{\beta}_0 = F(\psi)/G(\psi)$, while the approximate estimator in (5) is $F'(\psi)/G'(\psi)$. This approximation is accurate if

$$\frac{F(\psi)}{G(\psi)} \approx \frac{F'(\psi)}{G'(\psi)}.$$

This holds either (i) if the true model is as in (2), or (ii) if $F(0) = G(0) = 0$ and $G'(0) \neq 0$. Since the FFN25 model does not satisfy (2), the critical condition is that $G(0) = F(0)$.

In the model at hand, the requirement $F(0) = G(0) = 0$ has a precise interpretation: the prices and portfolio shares of red and green assets must be identical in the initial equilibrium where $\psi = 0$. By Corollary 1, **this is the case if and only if** the knife-edge condition $\rho = \frac{1}{2}$ holds, so that the inside assets are *perfectly symmetric*. Conversely, the KY25 approximation is invalid if $\rho \neq \frac{1}{2}$. We next show that this can create severe bias. However, we also exactly replicate their identification result under perfect symmetry.

Proposition 1 (Bias in the cross-sectional estimator) *If $\rho \neq \frac{1}{2}$, the KY approximation is invalid and the exact value of the KY25 estimator $\hat{\beta}_0$ satisfies*

$$\lim_{\psi \to 0} \hat{\beta}_{0,\text{exact}}(\psi) = -1, \tag{6}$$

which is of the opposite sign as, and does not vary with, the structural elasticity $\mathcal{E}_{gg}$.

If $\rho = \frac{1}{2}$, then the KY approximation is accurate and equation (B18) of KY25 holds,

$$\lim_{\psi \to 0} \hat{\beta}_{0,\text{exact}}(\psi) = \frac{1}{\epsilon^2} - 1. \tag{7}$$

Proof of Proposition 1. Suppose first that $\rho \neq \frac{1}{2}$. By properties of the logarithm,

$$\begin{aligned}
\hat{\beta}_{0,\text{exact}}(\psi) &= -\frac{\log(\omega_g(\psi)/\omega_r(\psi))}{\log(p_g(\psi)/(p_r(\psi))} \\
&= -\frac{\log\left(p_g(\psi)a_g(\psi)/p_r(\psi)a_r(\psi)\right)}{\log(p_g(\psi)/(p_r(\psi))} \\
&= -\frac{\log\left(a_g(\psi)/a_r(\psi)\right)}{\log\left(p_g(\psi)/p_r(\psi)\right)} - 1.
\end{aligned}$$

By market clearing, equilibrium quantities given $\psi = 0$ satisfy $a_g(0) = a_r(0) = \frac{1}{2}$. Since

$p_g(0) \neq p_r(0)$ given $\rho \neq \frac{1}{2}$, it follows that $\lim_{\psi \to 0} \hat{\beta}_{0,\text{exact}}(\psi) = -1$.

If instead $\rho = \frac{1}{2}$, the above expression for $\lim_{\psi \to 0} \hat{\beta}_{0,\text{exact}}(\psi)$ is indeterminate because $p_g(0) = p_r(0)$, so that $\lim_{\psi \to 0} \log(p_g(\psi)/p_r(\psi)) = 0$. Applying L'Hôpital's rule then yields

$$\lim_{\psi \to 0} \hat{\beta}_{0,\text{exact}}(\psi) = -\lim_{\psi \to 0} \frac{\frac{d}{d\psi}\left[\log \omega_g(\psi) - \log \omega_r(\psi)\right]}{\frac{d}{d\psi}\left[\log p_g(\psi) - \log p_r(\psi)\right]} = \frac{1}{\epsilon^2} - 1.$$

■

This result shows that Koijen and Yogo (2025) estimator accurately identifies β_0 only in the knife-edge case where $\rho = \frac{1}{2}$, and is severely biased otherwise. Specifically, the estimator $\hat{\beta}_{0,\text{exact}}$ has the *wrong sign*, and is constant even though the true elasticity $\mathcal{E}_{gg}$ ranges from 0 to infinity as ϵ ranges from 1 to 0. The reason is that KY25 does not use an exact solution, and that their approximation is inaccurate outside the special case $\rho = \frac{1}{2}$. In real-world data, knife-edge restrictions are unlikely to hold exactly, and verifying them empirically is infeasible with finite data. At least in FFN25's model, this implies that slight differences in risk exposures or investor beliefs can result in large and systematic biases.

Remark 1 *Given our results, one might wonder whether it is preferable to start from a cross-sectional estimator defined in differences rather than levels. Haddad, He, Huebner, Kondor, and Loualiche (2025) propose such a difference-in-differences estimator, and show that it identifies the* relative *elasticity between two assets (i.e., the own-price minus cross-price elasticity), but not the own-price elasticity directly. While their estimator has better properties than the KY25 estimator, it still needs (weaker) symmetry assumptions to be unbiased. FFN25 provides additional discussion.*

4 Comparison with FFN25's estimator

We now briefly describe the estimator used in FFN25. The goal of that paper is to establish that elasticity estimates based on logit asset demand systems can be susceptible to biases from cross-asset spillovers. To do so, one must define an estimator that is consistent with the assumptions of logit demand and performs well if there are no cross-asset spillovers.

Like the KY25 estimator, the FFN25 estimator is built on the assumptions of logit demand. However, it uses within-asset variation rather than the cross-section of assets. This has the simple benefit that, absent strong symmetry assumptions, an asset is a better "control" for itself than other assets. While clean "time series" variation of this sort may be hard to come by in practice, it can be constructed within a model to evaluate properties of logit demand, such as its ability to account for cross-asset complementarities.

Specifically, under the hypothesis of logit demand, parameter β_0 can be identified from observed portfolio shares of an inside asset *prior to* and *after* a supply shock (appropriately scaled by the outside asset's portfolio share to account for the budget constraint). Thus, the estimator is the *observed* elasticity of demand relative to the outside asset,

$$\hat{\mathcal{E}}_{gg} \equiv -\frac{d\omega_g/\omega_2}{dp_g}\frac{p_g}{\omega_g/\omega_2}, \tag{8}$$

where ω_2 is the portfolio weight of the outside asset (i.e., tree 2).

FFN25 shows that, in their model, this estimator has zero bias in the case where $\epsilon = 1$. This is an important benchmark because $\epsilon = 1$ ensures that all three assets are ex-ante symmetric. Consistent with the assumptions of logit demand, this in turn implies that the substitution matrix is perfectly symmetric. Furthermore, the estimator has the same sign as, and is monotone in, the structural elasticity for all $\epsilon < 1$. This makes it a suitable benchmark to assess the scope for spillover biases in the context of logit demand.

To see this spillover bias, suppose that the data generating process differs from (2) due to cross-price effects. Log-linearizing the demand functions from Section 2 yields

$$\log(\omega_g) = -\beta_0 \log(p_g) - \beta_r \log(p_r) + \alpha + \eta_g.$$

Given such a demand function, the FFN25 estimator is misspecified and evaluates to

$$\underbrace{-\frac{d\log(\omega_g)}{d\log(p_g)}}_{\hat{\mathcal{E}}_{gg}} = \underbrace{-\frac{\partial\log(\omega_g)}{\partial\log(p_g)}}_{\mathcal{E}_{gg}} \underbrace{- \overbrace{\frac{\partial\log(\omega_g)}{\partial\log(p_r)}}^{\text{cross-substitution}} \cdot \overbrace{\frac{d\log(p_r)}{d\log(p_g)}}^{\text{price spillovers}}}_{\text{bias}}.$$

Hence the estimator is biased when the data exhibits cross-substitution and price spillovers. This is precisely the point of FFN25: estimators based on logit demand without cross-asset complementarities will be biased when such complementarities are present.

Figure 1 compares the structural elasticity, the exact cross-sectional estimator from KY25, and the within-asset estimator of FFN25 given a small deviation from perfect symmetry ($\rho = 0.495$). The cross-sectional estimator is biased and has the wrong sign throughout. The within-asset estimator instead exhibits small biases when $\epsilon \approx 1$, but it is strongly biased if ϵ is close to zero. This is because logit demand cannot account for asymmetric substitution patterns, whereas the underlying model does exhibit asymmetric substitution when ϵ is small. (Namely, inside assets are closer substitutes for each other than for the outside asset.) This underscores the core message of FFN25.

5 Conclusion

The cross-sectional estimator from KY25 does not address the problem of spillover bias in the context of FFN25's model, except in the case of a precise symmetry condition. The within-asset estimator from FFN25 is a useful theoretical benchmark because it is less biased and performs well under symmetric substitution.

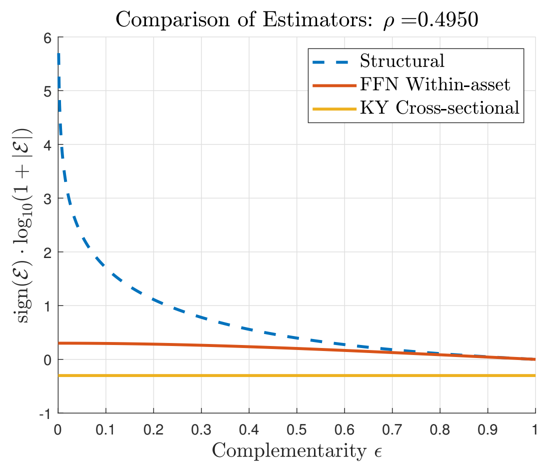

Figure 1: Structural elasticities, and the FFN25 and KY25 estimators given a small perturbation of symmetry ($\rho = 0.495$.) We use an adjusted $\log_{10}(1 + |\mathcal{E}|)$-scale to accommodate negative values.

References

FUCHS, W., S. FUKUDA, AND D. NEUHANN (2025): "Demand-System Asset Pricing: Theoretical Foundations," *Working paper, University of Texas at Austin and Bocconi University*. [Cited on pages 1, 2, 3, 5, 6, 7, 8, 9, and 10.]

HADDAD, V., Z. HE, P. HUEBNER, P. KONDOR, AND E. LOUALICHE (2025): "Causal Inference for Asset Pricing," Discussion paper. [Cited on page 8.]

KOIJEN, R. S. J., AND M. YOGO (2019): "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, 127(4), 1475–1515. [Cited on pages 1 and 2.]

——— (2025): "On the Theory and Econometrics of (Demand System) Asset Pricing," Working paper. [Cited on pages 1, 2, 3, 4, 5, 6, 7, 8, 9, and 10.]