

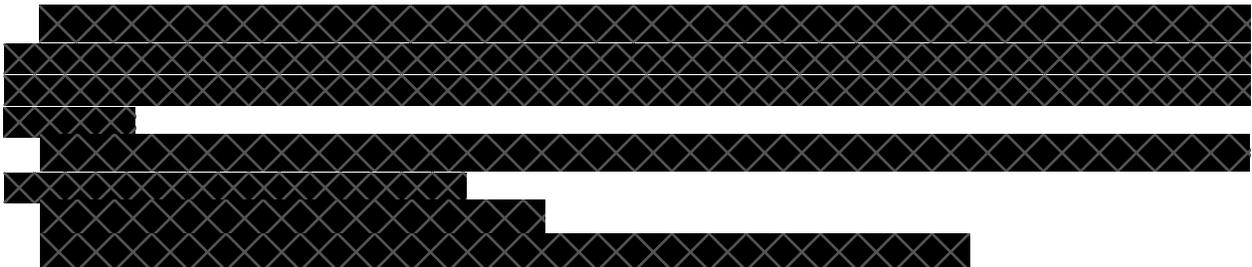
A Trilemma for Asset Demand Estimation*



November 18, 2025

Abstract

This paper derives fundamental limits to identifying asset demand from observational data. We establish a trilemma: it is impossible to simultaneously maintain that (i) prices satisfy no arbitrage, (ii) investors value assets for their payoffs, and (iii) asset-level demand elasticities can be inferred from supply shocks to individual securities. Resolving the trilemma requires as many independent quasi-experiments as the dimensionality of the asset span, or unverifiable theoretical restrictions. These results define the feasible scope of empirical identification in asset demand estimation and provide guidance for credible research design.



1 Introduction

Can we identify investors' demand functions for financial assets from their observed portfolios? This question underlies much of the recent "demand-system" approach to asset pricing. Building on methods from industrial organization, this approach typically uses supply shocks to estimate asset-level demand elasticities, interpreting them as the slopes of asset-specific demand functions. Such an interpretation requires the econometrician to isolate the effect of a single asset price, holding *all others fixed*.

In financial markets, however, such *ceteris paribus* variation is unattainable. Because investors choose portfolios, asset prices move together through the common pricing kernel and the principle of no arbitrage. A shock to the supply of one asset inevitably induces changes in the prices of others. As a result, the observed demand response to a price change does not isolate a structural demand slope, but instead reflects a mixture of reallocations across assets dictated by equilibrium spillovers.

We formalize this problem as a trilemma for asset demand estimation: given observational data, one cannot simultaneously maintain that (i) prices respect no arbitrage, (ii) investors care about asset payoffs, and (iii) asset-level demand elasticities can be recovered from supply shocks to individual assets. That is, given standard preferences and the weak requirement of no arbitrage, the ideal experiment which would identify an asset-level demand elasticity is infeasible. This implies strict theoretical limits on the identification of asset demand functions from observational data.

Neither no arbitrage nor payoff-based preferences are easily discarded. Preferences over payoffs (as opposed to direct preferences over assets) form the basis of portfolio choice theory. Given such preferences, no arbitrage is necessary to ensure the existence of smooth demand functions that do not jump discontinuously in response to small price changes—as is required for demand estimation. Empirical applications also rely on no arbitrage to derive empirically tractable demand systems based on a small number of common characteristics and risk factors (Ross, 2004). This dimension reduction is necessary because financial markets may feature hundreds or thousands of assets.

The problem becomes clear when we write asset prices in terms of *state prices*, which measure the cost of state-contingent consumption. Let p be the vector of asset

prices, Y the payoff matrix, and q the vector of state prices. Arbitrage free pricing requires

$$p = Yq.$$

Now consider the canonical notion of an asset-level demand elasticity. This object corresponds to an *ideal experiment* in which we vary the price of a single asset j while holding all other prices fixed. Under no arbitrage, a shock to a single asset price implies a specific set of state price changes. For simplicity, assume for now that markets are complete and that there are no redundant assets. Then state prices are related to asset prices by $q = Y^{-1}p$, and a shock to a given asset price affects state prices according to

$$\frac{\partial q}{\partial p} = Y^{-1}.$$

The ideal experiment in which we exogenously vary a single asset price can be equivalently interpreted as an experiment in which we induce a specific set of state price changes that are fully characterized by the *inverse* payoff matrix Y^{-1} . Measuring the elasticity thus requires researchers to generate precisely these implied state price changes.

This presents a fundamental challenge. Under standard risk averse preference, an increase in the supply of state-contingent consumption in a given state reduces the associated state price. Asset-level supply shocks thus alter state prices in all states in which the asset pays off, and these state price changes are proportional to payoff matrix Y . Except in the knife-edge case where Y is diagonal (i.e., assets are Arrow securities), the state price variation induced by supply shocks thus differs from the state price variation in the ideal experiment, which is proportional to Y^{-1} . Indeed, except in the case of a diagonal payoff matrix, the induced state price changes may not even have the same *sign*. Hence individual asset-level supply shocks generically fail to generate the ideal experiment that would allow for direct identification of the demand elasticity, and may induce substitution in the wrong direction compared to the ideal experiment.

While we state our trilemma under strict definitions of no-arbitrage and payoff-based preferences, the resulting limits to identification are not knife-edge. With respect to preferences, investors may care about cash flows *and* non-pecuniary aspects of an asset, such as its social responsibility score. Yet even if such preferences partially decouple valuations from state prices, they do not do so entirely. Hence, cross-asset restrictions im-

plied by no-arbitrage continue to constrain identification. Alternatively, a financial market may exhibit limits to arbitrage, such as segmentation or leverage constraints. While these frictions dampen price spillovers across assets, they rarely eliminate them. Thus, the tendency for asset prices to reflect common state prices remains even in these settings.

We formalize our argument in a general portfolio choice model, showing that it also generalizes to the case of incomplete markets and redundant assets. In the case of incomplete markets, supply shocks create appropriate state price variation only if *for each state, there is a unique asset with a positive payoff*. This is a much stronger condition than requiring assets to be uncorrelated conditional on risk factors, as is commonly assumed in the literature. We also use a two-asset example to describe how the direction and magnitude of the implied consumption shifts diverge between the ideal experiment and the actual equilibrium allocation. Across both complete and incomplete markets, the results reveal a robust obstacle to estimating asset-specific elasticities using supply shocks.

We then turn to potential solutions to the identification problems posed by our trilemma. To do so, we first consider an idealized setting in which the researcher observes N independent “experiments” which reveal an investors’ demand response to exogenous price variation. Under the simplifying assumption of linear demand (which, depending on the setting, may be very restrictive), we prove that one can point-identify the entire matrix of demand parameters *only if* N is no smaller than J , the dimensionality of the asset span. This provides a constructive benchmark: with as many independent shocks as the dimensionality of the asset span, (linear) demand functions are point-identified. However, with fewer than J experiments, identification collapses to projections onto the subspace of observed price changes. Since these projections are not chosen by the econometrician but are endogenously determined by equilibrium spillovers, we show that demand parameters are arbitrarily unconstrained outside the span of observed shocks. Allowing for non-linear demand further increases the number of required experiments.

These data requirements are stringent: in many applications, the asset span is high-dimensional and J may be in the hundreds or thousands. Moreover, the researcher must observe J *independent* changes in the vector of equilibrium prices and all associated quantities. This is a much stricter requirement than assuming the existence of J independent supply shifters which occur within a given observation period—simultaneous price shifters generate only a *single* experiment comprising of a single shock to the price vector.

Given this disconnect between supply shocks and the ideal experiment, we stress that structural assumptions (e.g., on payoffs, preferences, or the pricing kernel) are, in practice, likely necessary to recover demand elasticities from equilibrium data. The key challenge is that these assumptions are inherently untestable in observational data because the required *ceteris paribus* price variation is never observed. Moreover, because (arbitrarily) many theoretical models may be consistent with the observed data, these models must be judged on *ex-ante* theoretical considerations and plausibility, not on their empirical fit. Nevertheless, we view the combination of portfolio data with sound models of portfolio choice as the promising direction for future research.

The remainder of the paper is organized as follows. Section 2 introduces the model and defines the key objects of interest. Section 3 clarifies the link between demand functions and no arbitrage in the context of our framework. Section 4 formalizes the inconsistency among no arbitrage, state prices, and asset-level elasticities. Section 5 discusses how researchers can overcome the trilemma by relying either on a very large number of independent quasi-experiments or structural assumptions. Section 6 illustrates the logic using a fully-solved general equilibrium model. Section 7 discusses implications for estimation and provides concluding remarks. The proofs are delegated to the Appendix.

Related literature

Our paper contributes to a growing literature on asset demand estimation, particularly the empirical study of asset-level demand elasticities following [Kojien and Yogo \(2019\)](#). Their framework—and much subsequent work—relies on the interpretation that changes in asset prices induced by supply shocks can be used to infer structural demand elasticities. However, while [Kojien and Yogo \(2019\)](#) explicitly rely on a structural model to discipline demand estimation, other authors pursue more reduced-form strategies—see [Gabaix and Kojien \(2020\)](#) or [van der Beck \(2021\)](#). Our key result is that the same no arbitrage restrictions which facilitate smooth, low-dimensional demand functions greatly complicate the identification of demand elasticities in observational data.

While these considerations are critical for financial markets, they are less central in other settings. In the canonical industrial organization framework (see [Berry and Haile \(2021\)](#) for a review), a shock to the supply of apples might affect the price of bananas, but assuming banana prices are fixed does not undermine the validity of the consumer

problem. In financial markets, the analogous assumption dramatically alters basic properties of demand and market equilibrium. The crucial difference is that financial assets are not consumption goods; they are claims to bundles of state-contingent payoffs. Thus, if investor preferences pertain to state-contingent wealth or consumption, their demand functions pertain to asset *portfolios* that jointly deliver a desired payoff process, not to individual assets per se.

This observation also clarifies the relationship of our paper to important recent work in industrial organization which estimates demand systems with complementarities (e.g., [Iaria and Wang, 2020](#); [Wang, 2024](#); [Fosgerau, Monardo, and de Palma, 2024](#); [Ershov, Laliberté, Marcoux, and Orr, 2024](#)). These approaches typically study settings in which consumers make discrete choices over a limited number of bundles, or where substitution patterns are governed by exogenous functional-form parameters. This contrasts with asset-market settings, where investors choose continuous portfolio allocations and the degree of complementarity or substitutability is endogenously determined by asset holdings and general-equilibrium interactions. As such, the source of the complementarities and the methods to deal with them are distinct from the ones arising from no arbitrage. The spirit of the exercise we conduct is also different. Whereas the recent IO literature develops empirical tools for estimating demand with complementarities under exogenous price variation, our paper asks: whether such elasticities can be identified at all once equilibrium restrictions like no-arbitrage endogenously couple prices across assets.

Our analysis also relates to work which emphasizes quantity effects, such as portfolio balance ([Tobin, 1969](#)), convenience yields ([Krishnamurthy and Vissing-Jorgensen, 2012](#)), intermediary asset pricing ([He and Krishnamurthy, 2013](#); [Adrian, Etula, and Muir, 2014](#)), and capital flow studies ([Shleifer, 1986](#); [Harris and Gurel, 1986](#)). These approaches highlight that asset demand shocks or cross-investors capital flows can have important effects on asset prices. Our contribution is to show that, under no-arbitrage and payoff-based preferences, the price elasticities measured in these approaches do not directly reveal structural slopes of demand functions: supply shocks trace general equilibrium reallocations, not structural elasticities.

There are a number of potential paths to circumventing the issues we raise. The first is to restrict attention to settings where one can estimate demand even without supply variation. This typically requires additional data on demand functions, as in the Cana-

dian bond market studied by [Allen, Kastl, and Wittwer \(2025\)](#). The second is to abandon no arbitrage. As discussed above, this makes it more difficult to reduce the dimensionality of the portfolio choice problem, and therefore likely limits the scope of asset demand estimation to settings with a limited number of assets. The third is to assume that preferences are defined directly over assets (as opposed to state-contingent payoffs), for example because investors have non-pecuniary *tastes* over assets. Because preferences are not defined over cash flows, there is no need to worry about cross-asset spillovers through the common pricing of cash flows. However, [Fuchs, Fukuda, and Neuhann \(2025\)](#) show that heterogeneous tastes may also invalidate no arbitrage pricing.

The fourth is to change the target of estimation, possibly under additional assumptions. One example is in [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#), who show that a difference-in-difference estimator based on a single supply shock can identify a *relative* elasticity (the own minus cross-price elasticities) under a set of stringent symmetry assumptions. Another example is [An and Huber \(2024\)](#), who propose measuring demand elasticities over aggregated factor portfolios (as opposed to individual assets). Our results here suggest that these factors must be designed to generate a diagonal payoff matrix. This is a stricter requirement than just requiring factors to be mutually orthogonal. Moreover, constructing exogenous supply variation at the portfolio level from asset-level shocks typically requires knowledge of the payoff matrix, which is a latent variable.¹

Overall, our results suggest an important role for model-based inference in asset markets. When the ideal experiment is infeasible, restrictions imposed by structural models allow researchers to learn about investor behavior from equilibrium data. This interpretation supports the general approach laid out in [Kojien and Yogo \(2019\)](#). However, because these restrictions can never be fully validated in the data, the assumed models must be plausible ex-ante. However, [Fuchs, Fukuda, and Neuhann \(2025\)](#) also show that logit demand systems of the type employed in [Kojien and Yogo \(2019\)](#) may not appropriately account for cross-asset interactions in portfolio choice.

¹Constructing a diagonal payoff matrix will typically require forming long-short portfolios. See [An \(2025\)](#) for an analysis of demand elasticities for long-short portfolios.

2 Setup

We consider a canonical model with a set I of potentially heterogeneous investors. Each investor $i \in I$ must choose how much to consume at date 0 and across Z states of the world at date 1. To acquire a desired state-contingent consumption profile, the investor can invest in J assets. Investor i 's *portfolio* is a vector $a^i \equiv (a_j^i)_{j=1}^J \in \mathbb{R}^J$ of asset positions, where each element a_j^i denotes the investor's holdings of asset j . Asset j has payoff $y_j(z)$ in state z . We denote by $Y \equiv (y_j(z))_{j,z}$ the $J \times Z$ matrix of cash flows. In line with the literature, we assume that the payoff matrix is known to the investor but unobserved by the econometrician. This is because the payoff matrix reflects expected returns, which is latent. Prices are observed by both the investor and the econometrician.

We treat time-zero consumption as the numeraire good (or, equivalently, as the *outside asset*) whose price is normalized to 1. Investor i is endowed with e_j^i units of asset j and e_0^i units of the numeraire. The budget constraints at date 0 and in state z are given by

$$c_0^i = e_0^i - \sum_{j=1}^J p_j(a_j^i - e_j^i) \quad \text{and}$$

$$c_z^i = \sum_{j=1}^J y_j(z)a_j^i \quad \text{for all } z.$$

Each investor i has standard preferences over consumption given by

$$U^i(a^i) \equiv (1 - \delta^i)u^i(c_0^i) + \delta^i \sum_{z=1}^Z \pi_z u^i(c_z^i),$$

where $\delta^i \in (0, 1)$ is the discount factor, u^i is a strictly increasing and strictly concave von-Neumann Morgenstern utility function, and $\pi_z \in (0, 1)$ is the probability of state z .

Investors may face constraints on portfolio formation beyond the budget constraint. Let \mathcal{A}^i denote the set of feasible portfolios of investor i , and assume that \mathcal{A}^i is a closed convex subset of \mathbb{R}^J . The investor's decision problem is:

$$\sup_{a^i \in \mathcal{A}^i} U^i(a^i). \tag{1}$$

When necessary, we close the model using the standard notion of competitive equilibrium, whereby investors form optimal portfolios given prices and asset markets clear.

The general solution to the portfolio choice problem consists of J asset-level demand functions $a_j^i(p)$ each of which depends on the entire vector of asset prices p . These asset-level demand functions can be characterized by the price elasticity of demand, which is the *partial derivative* of demand for asset j with respect to an asset price p_k , holding all other prices fixed. That is,

$$\mathcal{E}_{j,k}^i \equiv -\frac{\partial a_j^i(p)}{\partial p_k} \frac{p_k}{a_j^i(p)}.$$

The main difficulty in measuring this object is that the demand function depends on the entire vector of asset prices. Hence identifying the demand elasticity requires a setting in which there is exogenous variation in one price but others remain fixed.

Remark 1 (Dependence on multiple prices even in special settings) *The dependence of demand functions on all multiple asset prices exists even in settings that purportedly induce asset demand functions that depend only on the asset’s own price. For example, [Kojien and Yogo \(2025, Appendix A\)](#) study a model with CARA preferences, normally distributed payoffs, a diagonal covariance matrix conditional on factors, and unconstrained borrowing at an interest rate normalized to zero. The first two features generate linear marginal utility, and the combination with the third and fourth features yields separable asset demand functions that depend only the excess expected return and volatility of a given asset. However, the independence of other prices is achieved by “normalizing” the risk free rate to a fixed number. Yet in equilibrium, the risk-free rate is the inverse sum of state prices, and thus depends on the vector of asset prices.*

More generally, specific settings such as CARA-normal with diagonal covariance matrices generate separable demand functions only under knife-edge assumptions on the payoff structure. Hence we study the generic case where demand depends on the price vector.

3 The Importance of No Arbitrage for Demand Analysis

Demand analysis in financial markets faces two basic challenges. The first is the large number of assets under consideration. For example, in US equities markets alone, investors can choose among many *thousands* of assets, which creates a curse of dimensionality in demand estimation. The second is that demand functions must be sufficiently

well-behaved. For example, demand elasticities are partial derivatives of demand with respect to an asset price. Hence the demand elasticity can be used to describe demand only if demand functions are smooth functions of asset prices.

Both challenges can be addressed by relying on the principle of no arbitrage. With respect to the first challenge, [Kojien and Yogo \(2019\)](#), and many papers based on their approach, implicitly rely on Ross's arbitrage pricing theory to argue that asset demand can be summarized by a relatively small number of asset characteristics and common risk factors, leading to a low-dimensional representation. (We leave aside here the concern that several common characteristics, such as book-to-market ratios, are themselves endogenous to demand.) With respect to the second challenge, it is well-established that arbitrage opportunities can lead to discontinuous changes in demand functions with respect to arbitrarily small price changes. Imposing no arbitrage rules out such discontinuities, thereby facilitating an analysis of demand elasticities. For demand analysis, no arbitrage is thus not only a constraint on *equilibrium* prices, but an important restriction on investors' decision problems themselves.

In the following, we briefly recapitulate the link between demand functions and no arbitrage in the context of our model. Since much of the empirical literature emphasizes institutional investors' constraints on portfolio choice when designing instruments, we explicitly incorporate these into our analysis as well. We also establish the standard result that no arbitrage allows for an analysis of asset prices (and thus demand) using *state prices*.

To this end, we begin by defining *unbounded* arbitrage opportunities as those that can be exploited using arbitrarily large asset positions. Standard definitions of arbitrage always consider unbounded arbitrage opportunities ([Duffie, 2001](#)). Hence this definition differs only in that we permit *bounded* arbitrages. (We discuss this case below.)

Definition 1 (No unbounded arbitrage for investor i) *Investor i has an unbounded arbitrage opportunity if, for any $m > 0$, there exists a portfolio $a^i \in \mathcal{A}^i$ such that either (i) $p \cdot a^i \leq 0$, $Y^T a^i \geq 0$, and $(Y^T a^i)_z \geq m$ for some z or (ii) $p \cdot a^i \leq -m$ and $Y^T a^i \geq 0$. Otherwise, investor i has no unbounded arbitrage opportunity.*

Proposition 1 shows that a well-defined decision problem requires the absence of unbounded arbitrage opportunities. The simple reason is that unbounded arbitrage precludes the existence of a solution to the investor's problem. Hence the absence of unbounded arbitrage is thus a minimal requirement for any analysis of investor demand

functions. This is a well-known result based on textbook treatments (e.g, [Duffie, 2001](#)).

Proposition 1 (Duffie (2001): No arbitrage and the investor’s problem) *If there is a solution to (1), then investor i has no unbounded arbitrage opportunity. If U^i is continuous and investor i has no unbounded arbitrage opportunity, then there is a solution to (1).*

Proof. See Appendix. ■

Under weak additional assumptions, asset prices and demand can then be analyzed using *state prices*, which measure the marginal cost of a unit of state-contingent consumption. In particular, if the union of investors’ feasible sets covers the entire space of feasible portfolios \mathbb{R}^J , the absence of unbounded arbitrage implies the existence of state prices such that asset prices are payoff-weighted sums of state prices.

Lemma 1 (Existence of state prices) *If there exists a subset I_0 of investors such that every investor $i \in I_0$ does not have an unbounded arbitrage opportunity and $\mathbb{R}^J = \bigcup_{i \in I_0} \mathcal{A}^i$, then there exist state prices $q \in \mathbb{R}_{++}^Z$ such that asset prices are payoff-weighted sums of state prices:*

$$p = Yq. \tag{2}$$

Taken together, these results show that the absence of unbounded arbitrage is required for well-defined demand functions, and that, along with a weak condition on portfolio constraints, it implies the existence of strictly positive state prices. Hence we can recast both the interpretation and measurement of demand elasticities in terms of the cost of state-contingent consumption. This is useful because, given preferences over consumption, asset prices affect consumption plans to the extent that they affect state prices.

Bounded arbitrage. The same basic consideration applies to bounded arbitrage, whereby investors can only exploit mispricing up to an exogenous constraint on asset positions. In particular, it remains optimal to exploit the arbitrage to the extent possible, and, as the example below illustrates, this can lead to discontinuous changes in demand in response to arbitrarily small price changes.

Example 1 (Discontinuous demand functions) *Suppose there are two states of the world at*

date 1, and three assets. Given some $\epsilon \in (0, \frac{1}{2})$, let a cash flow matrix Y be given by

$$\begin{bmatrix} \frac{1}{2} + \epsilon & \frac{1}{2} - \epsilon \\ \frac{1}{2} - \epsilon & \frac{1}{2} + \epsilon \\ 1 & 1 \end{bmatrix}.$$

Now consider the demand functions for some investor i with continuous utility function U^i .

- (i) Suppose $\mathcal{A}^i = \mathbb{R}^3$. The absence of unbounded arbitrage requires that $p_3 = p_1 + p_2$. Given this restriction on prices, well-defined demand functions exist for all three assets, with the investor taking weakly positive quantities in all three assets. Now suppose that, starting from an initial benchmark where no arbitrage pricing holds, p_3 increases slightly. Then, investor i 's problem (1) is no longer well-defined, and well-defined demand functions no longer exist.
- (ii) Suppose instead that investor i faces the short-sale constraint $a_j^i \geq -\chi$ for some $\chi > 0$. Given $p_3 = p_1 + p_2$, well-defined demand functions still exist for all three assets, with the investor taking weakly positive quantities in all three assets. Now suppose that p_3 increases slightly. Then it is optimal for the investor to jump to a portfolio allocation where $a_3^i = -\chi$. This can trigger discontinuities in optimal demand.

It is clear that such discontinuities prevent a fruitful analysis of demand elasticities. Assuming portfolio constraints is thus in general not sufficient to have a well-posed estimation problem in the presence of redundant assets. Since any infinitesimal price change triggers an arbitrage for redundant assets, for the remainder we focus on the more interesting case without redundant assets.

Assumption 1 (No redundant assets) $Z \geq J$ and $\text{rank}(Y) = J$.

4 The Trilemma

We now turn to our main result, which pertains to the difficulty in measuring demand elasticities for individual assets in settings governed by no arbitrage restrictions. In line with the literature, we focus primarily on the case where the demand elasticity for a given asset is estimated by instrumenting for its price using a shock to the (residual) supply curve for that asset. (Remark 3 considers instruments constructed from multiple shocks.)

Ideal experiment. The notion of an asset-level demand elasticity pertains to an *ideal experiment* in which one traces out an investor's demand response to *ceteris paribus* variation in a single asset price. As such, the demand elasticity is widely interpreted as the slope of an asset-specific demand function. We show that no arbitrage restrictions imply sharp conditions on the exact nature of state price changes associated with the ideal experiment, and that asset-level supply shocks do not satisfy these restrictions.

It is useful to describe the ideal experiment in terms of state prices, as these ultimately determine the investor's optimal consumption plans through the cost of consumption. Say that the investor observes asset prices p and payoff matrix Y . Then equation (2) allows the investor to infer the vector of state prices implied by prevailing asset prices:

$$q = Y^+ p, \quad (3)$$

where Y^+ is the Moore-Penrose pseudo-inverse of Y . If Y is square, as when markets are complete, then $Y^+ = Y^{-1}$ and there is a unique vector of state prices. If markets are incomplete ($J < Z$), then there are many feasible state price vectors. We focus on the minimum norm solution, whereby the pseudo-inverse is $Y^+ = Y^T(Y Y^T)^{-1}$.

Implied state prices. Equation (3) determines the vector of state price changes that occur when we vary asset price p_j . That is, if we impose no (unbounded) arbitrage on the investor's problem to obtain well-behaved demand functions, then the thought experiment in which there is *ceteris paribus* variation in a given asset price is formally equivalent to one in which we vary the set of *state prices* that determine the cost of consumption.

Lemma 2 (State price changes from an ideal experiment) *Let v_j denote the unit vector in \mathbb{R}^J with 1 in the j -th position and zeros elsewhere. Then the changes in state prices given the exogenous variation in a single price p_j are*

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

Proof. The assertion follows immediately from equation (3). ■

These state price changes then induce a change in the optimal consumption plan, which can then be mapped into a change in desired portfolio holdings. Estimating the demand elasticity therefore requires an empirical setting in which one can generate the state

price variation $\Delta \mathbf{q}_j^{\text{ideal}}$ associated with the ideal experiment. The challenge is that these restrictions can be tight: in particular, when markets are complete, the ideal experiment requires a *unique* set of state price changes.

Measurement using supply variation. In practice, one does not generally observe exogenous price shocks. Instead, one observes shocks to an economic environment that might trigger equilibrium price changes. As such, empirical approaches to estimating asset demand elasticities typically rely on suitably exogenous variation in asset supply (e.g., through flows or institutional holdings) and interpret the resulting change in an asset price as identifying a local demand response. However, in demand systems where the demand for one asset depends on the prices of others (as in financial markets, either through preferences or arbitrage relations), obtaining exogenous variation in a single price is insufficient. Instead, one must ensure that other asset prices remain unchanged.

We now show that these conditions are generically not satisfied even when researchers have access to quasi-experimental variation in asset supply. The key problem is that asset prices are linked to state prices through no arbitrage restrictions, and that supply shocks alter state prices in all states of the world in which the asset pays off. Hence, the prices of other assets that pay off in those states will also change. Unless the asset has a unique state-contingent payoff (i.e., Y contains a diagonal matrix), there is no reason to expect only p_j to change.

To establish this result, we must describe how supply shocks affect state prices. To do so, we impose the weak requirement that a positive supply shock to asset j reduces state prices in all states where asset j has a strictly positive payoff. To provide favorable conditions for identification, we further assume that these state price changes are *separable* across states, in the sense that a supply shock does not directly induce state price changes in states of the world where the shocked asset does not affect payoffs. This assumption also aligns with the canonical expected utility framework in which utility over consumption is additively separable across states. Remark 2 below argues that settings with non-separable state prices generically present more difficult identification challenges.

Definition 2 (Separable downward-sloping consumption demand) Let $E \equiv (E_j)_{j=1}^J \in \mathbb{R}_{++}^J$ denote the vector of aggregate asset endowments. An economy has separable downward-sloping consumption demand if there exists a $Z \times Z$ diagonal matrix V with strictly positive

elements such that

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

where y_j^T is the transpose of the j -th row $y_j \equiv (y_j(z))_{z=1}^Z$ of Y .

Under this definition, price changes are proportional to the induced change in consumption, as determined by the payoff matrix Y , multiplied by the marginal change in valuations induced by this shift, as measured by V . The assumption that V is diagonal reflects the notion that asset supply shocks have *separable* marginal effects on state prices. (Because we are interested in marginal effects, this does not imply that the total response to any given shock must satisfy the same diagonal structure.)

In standard models, V can be interpreted as the marginal utility of the marginal investor. Example 2 provides an illustration.

Example 2 (V in a representative-agent model) *In a standard representative-agent model, state prices relate to marginal utility over aggregate consumption,*

$$\frac{\partial q_z}{\partial E_j} = \frac{\delta}{1-\delta} \pi_z \frac{u''(C_z)}{u'(C_0)} y_j(z) < 0,$$

where C_0 and C_z are aggregate consumption at date 0 and in state z . Thus

$$V = -\frac{\delta}{1-\delta} \text{diag} \left(\pi_1 \frac{u''(C_1)}{u'(C_0)}, \dots, \pi_z \frac{u''(C_z)}{u'(C_0)}, \dots, \pi_Z \frac{u''(C_Z)}{u'(C_0)} \right).$$

Remark 2 (Non-separable state prices and non-diagonal V) *While the standard expected-utility framework implies that V is diagonal because utility is additively separable in state-contingent consumption, this property need not hold in general. Recursive preferences in the tradition of [Epstein and Zin \(1989\)](#) or [Kreps and Porteus \(1978\)](#) allow for non-separable valuations across states, implying that the shock-response matrix V may have non-zero off-diagonal elements.*

From the perspective of identification, such non-diagonal V greatly complicates the mapping from supply shocks to state prices: a single supply shock can affect multiple state prices simultaneously, mechanically inducing endogenous price responses in other assets and thereby violating the ideal ceteris paribus condition. In this sense, a general V corresponds to a fully non-separable substitution matrix in the language of industrial-organization demand models, where the slope of demand for one good depends on the levels of all others.

A knife-edge exception arises if the cross-state interactions in V exactly offset the cross-asset restrictions implied by no arbitrage, restoring separability. However, because V is determined by preferences while the no-arbitrage relation depends only on the payoff matrix Y , there is no economic reason for such an exact offset to hold in equilibrium.

Formal conditions for identification. We can now state two formal conditions which ensure that supply shocks create the type of state price variation required for the ideal experiment underlying a demand elasticity. The first is that the supply shock generates precisely required price variation, up to a scalar transformation that allows for a change in the size of the shock.

Condition 1 (Identical variation) *A supply shock to asset j generates the ideal state price variation for asset j if there exists some scalar k_j such that*

$$\Delta \mathbf{q}_j^{\text{ideal}} = k_j \Delta \mathbf{q}_j^{\text{supply}}.$$

This condition holds for all assets if

$$Y^+ = -VY^T K, \quad \text{where } K \equiv \text{diag}(k_1, \dots, k_J).$$

While natural, one might argue that Condition 1 is too strict. The supply shock may still provide useful variation if it does not depart too much from the ideal experiment. Hence we also consider a much weaker condition, namely the state price variation generated by a supply shock has the same *sign* as the state price changes in the ideal experiment.

Condition 2 (Variation of the same sign) *The supply shock generates state price variation of the same sign if $\Delta \mathbf{q}_j^{\text{ideal}}$ has the same sign as $\Delta \mathbf{q}_j^{\text{supply}}$ element by element. Given that Y has only weakly positive entries, this condition holds for all assets if Y^+ has only weakly positive entries.*

We can state our main result, which states that Conditions 1 and 2 are satisfied only under stringent conditions on the payoff matrix. In particular, if supply shocks are to induce useful variation in state prices, then for every state of the world there must exist a *unique* asset which offers a positive payoff in the world. Strikingly, both conditions require the *same stringent restrictions*. That is, as long as one wants to be sure to satisfy the

minimal requirement that the induced state price variation is of the same sign as in the ideal experiment, then there must be no assets with overlapping payoffs.

Definition 3 (Overlapping payoffs) *Assets j and j' have overlapping payoffs if there exists at least one state of the world z such that $y_j(z) > 0$ and $y_{j'}(z) > 0$.*

Theorem 1 (Trilemma) *If Condition 1 or Condition 2 is satisfied, then YY^T is diagonal, and:*

- (i) *If YY^T is diagonal, then there are no assets with overlapping payoffs.*
- (ii) *If markets are complete, then YY^T is diagonal if and only if Y is diagonal up to permutations.*

The trilemma has direct implications for recent empirical strategies that estimate asset demand from supply shocks. Under no arbitrage and preferences over state-contingent payoffs, supply shocks induce general equilibrium price changes that deviate from the ceteris paribus variation required to identify demand elasticities. As a result, observed demand elasticities reflect portfolio reallocation across all assets, not the partial derivative of demand with respect to own price. This limits the structural interpretability of observed asset-level demand elasticities even with quasi-experimental shocks to asset supply. In particular, observed behavior should be interpreted as reflecting *composite* demand functions which reflect interactions between multiple asset-level demand functions. Importantly, the same considerations apply even when assets offer small “convenience yields”—while such effects can induce return differences between assets, they do not eliminate the potential for spillovers. Similarly, the same issues would also arise in a dynamic setting where investors can trade securities referencing different states and dates, as these would also have to be priced by a common pricing kernel and governed by no arbitrage. This broader view helps connect our findings to those in [Binsbergen, David, and Opp \(2025\)](#).

Remark 3 (Constructing instruments from multiple shocks) *One potential solution to the problems discussed above is to construct price instruments from multiple supply shocks. For example, if a researcher has access to supply shocks for every asset, then a combination of these shocks may generate the right state price variation. To fix ideas, suppose that, for a given asset j , there exists a vector of coefficients $\psi \in \mathbb{R}^J$ such that the ideal variation $\Delta \mathbf{q}_j^{\text{ideal}}$ is a linear combination of the vectors of supply shocks:*

$$\Delta \mathbf{q}_j^{\text{ideal}} = \Delta \mathbf{q}^{\text{supply}} \psi.$$

Then, an instrument constructed from a ψ -weighted combination of asset-level supply shocks would be suitable for estimating the demand elasticity for asset j . However, while this is theoretically possible, the main practical challenge is that the weights ψ necessarily depend on the payoff matrix Y , which is unobserved to the econometrician. Hence, constructing instruments in this manner does not present a practical solution unless one is willing to rely on strong assumptions about the payoff matrix.

5 Overcoming the Trilemma

Our results thus far highlight a fundamental disconnect between the state price variation required for the ideal experiment and the state price variation which is generated by asset-level supply shocks. In this section, we discuss how researchers can overcome these challenges using either rich quasi-experimental data or stricter structural assumptions.

5.1 Multiple independent experiments

We begin by analyzing whether asset-level demand functions can be identified in an idealized setting where the researcher has access to multiple, *independent* quasi-experimental shocks to asset prices, and demand functions are approximately linear.²

A first-order approximation of investor i 's asset demand function around \bar{p} yields

$$a_i = \bar{a}_i + S_i(p - \bar{p}) + \varepsilon_i,$$

where $a_i \in \mathbb{R}^J$ is the vector of portfolio holdings, $p \in \mathbb{R}^J$ is the price vector, and ε_i is the vector of residual demand shocks. The asset-level *substitution matrix* $S_i \in \mathbb{R}^{J \times J}$ is the object of interest. Row k of matrix S_i collects the loadings of the demand for asset k on all prices, $(S_i)_k = \left(\frac{\partial a_{i,k}}{\partial p_1}, \dots, \frac{\partial a_{i,k}}{\partial p_J} \right)$, while column j captures the derivative $\partial a_i / \partial p_j$. Under the assumption of linear demand, this derivative determines how the vector of asset quantities responds to changes in price p_j .

Suppose that the researcher has access to N distinct “experiments” indexed by n . Each experiment consists of a purely exogenous shock to the supply of a given asset (or

²Linear demand can be a very strong assumption in financial markets. For example, in settings where close substitutes are available, even small price changes can lead to highly non-linear correlated quantity responses across multiple assets. Hence the results in this section apply more naturally to asset menus with limited substitution.

combination of assets) which creates exogenous price changes. This generates a matrix of observable price changes G and quantity changes A_i for each investor, defined as:

$$G \equiv [\Delta p^{(1)}, \dots, \Delta p^{(N)}] \in \mathbb{R}^{J \times N} \quad \text{and} \\ \Delta A_i \equiv [\Delta a_i^{(1)}, \dots, \Delta a_i^{(N)}] \in \mathbb{R}^{J \times N}.$$

We emphasize that the researcher is assumed to observe N *independent* changes in the vector of equilibrium prices and all associated quantity changes. This is a much stricter requirement than assuming the existence of N independent supply shifters within a given observation period. Simultaneous shocks generate only a single experiment—one induced set of price changes and one set of equilibrium demand responses to these price changes. For example, [Kojien and Yogo \(2019\)](#) rely on a single cross-section of prices and quantities, which corresponds to a single independent experiment.

Stacking the data from all N quasi-experiments yields the matrix equation

$$\Delta A_i = S_i G + U_i, \tag{4}$$

relating the set of all observed quantity changes to the observed price changes and matrix of residual demand shocks $U_i \in \mathbb{R}^{J \times N}$. Our assumptions imply that $\mathbb{E}[U_i | G] = 0$.

We begin by establishing a positive identification result. In particular, in the theoretical ideal where the number of independent experiments equals the dimensionality of the asset span, ordinary least squares identifies the investors' entire substitution matrix.

Proposition 2 (Complete identification of the substitution matrix with J experiments) *Let the number of independent experiments equal the dimensionality of the asset span, so that the matrix of observed price changes is full row rank, $\text{rank}(G) = J$. Let $G^+ (= G^{-1})$ denote the Moore-Penrose pseudoinverse of G . Then the unique ordinary least-squares estimator of S_i is*

$$\hat{S}_i = \Delta A_i G^T (G G^T)^{-1} = \Delta A_i G^+, \tag{5}$$

where \hat{S}_i is an unbiased and consistent estimator of S_i . When $U_i = 0$, $\hat{S}_i = S_i$.

Proposition 2 provides a constructive benchmark: with as many independent shocks as the dimensionality of the asset span, demand functions are point-identified. In particular, under the (strong) assumption of linear demand and when $N = J$, a standard ordinary

least squares regression of quantity changes ΔA_i on price changes is sufficient to identify the entire substitution matrix S_i . We refer to this result as *complete identification* because every element of S_i is point-identified. This provides one constructive method for asset demand estimation, which is to find settings in which researchers can observe sufficiently many shocks relative to the number of assets in the investor's choice set.

In many applications, researchers observe far fewer than J independent experiments. The reason is that J experiments require the existence of J independent shocks occurring in different periods or in distinct, entirely segmented markets with the same economic fundamentals. This is a tall order. Hence we must also assess the identification of substitution matrix S_i in the empirically relevant case where $N < J$.

The next result shows that the substitution matrix is not point-identified if $N < J$, and indeed that demand parameters are arbitrarily unconstrained outside the span of observed shocks. In such settings, demand estimates are therefore conditioned on auxiliary theoretical restrictions imposed by the researcher.

Proposition 3 (Incomplete identification with $N < J$ experiments) *Let $P_G \equiv GG^+$ be the orthogonal projector onto $\text{col}(G)$, the column space of the matrix of observed price changes G , where $G^+ \equiv (G^T G)^{-1} G^T$. Then the general solution to the least-squares problem is*

$$S_i = \Delta A_i G^+ + B_i (I_J - P_G),$$

where $B_i \in \mathbb{R}^{J \times J}$ is an arbitrary matrix that is entirely unrestricted by the data and I_J is the identity matrix.

Proposition 3 shows that any component of S_i in the null space of G is not point-identified and cannot be bounded without ex-ante theoretical restrictions which cannot be rejected by the data. What is identified is the *projection* of S_i onto the observed data:

$$S_i P_G = \Delta A_i G^+ P_G = \Delta A_i G^+.$$

An important question is whether this projection fully determines a particular structural parameter. This is the case only if the G contains "clean" variation that can be used to trace out variation in a single asset price, while holding all other prices fixed. Formally, this is equivalent to the requirement that v_j (i.e., the unit vector where element j is the only positive entry) is in the column space of G , so that the observed shocks span

a pure p_j movement. That is, an individual demand parameter $S_{i,j}$ is point identified if $v_j \in \text{col}(G)$. As Theorem 1 shows, this condition generically fails for asset-level supply shocks because they create correlated movements in other asset prices.

Despite these challenges, for specific individual assets and payoff matrices, it may well be the case that a supply shock does generate the ideal variation. In this case, our results offer a way of checking the internal consistency of an empirical test: since the estimated parameters of a model imply a payoff matrix, one can invert it and check whether the induced variation aligns with the ideal experiment.

Relation to Collinearity and Weak Instruments. When asset prices satisfy $p = Yq$, equilibrium price movements are confined to the low-dimensional space spanned by the state prices q . As a result, the matrix of observed price changes G is typically of rank deficient, implying that instruments constructed from asset-level supply shocks are highly collinear. In instrumental-variable terms, the first-stage regression of individual prices on such instruments is likely weak once other prices are controlled for: the conditional F -statistic is small even if unconditional correlations are large. However, the absence of weak instruments—that is, a strong first stage—is *not* sufficient for credible identification. Even when instruments generate large first-stage variation, they may still induce the wrong direction of price movements relative to the ideal experiment that isolates an own-price effect. In the terminology of Proposition 3, such instruments span an incorrect subspace of the price space, identifying only projections of demand elasticities rather than structural slopes. Hence, strong instruments ensure relevance but not alignment: they are necessary, but not sufficient, for consistent identification of asset-level demand.

5.2 Structural Assumptions

Our results thus establish that demand parameters are not point identified unless the payoff matrix is diagonal or the researcher has access to a large number of independent experiments. Outside of these ideal settings, researchers must therefore rely on theoretical restrictions to obtain point estimates for demand parameters.

One example is the semi-structural approach in [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#), who impose conditional homogeneity restrictions on the substitution matrix. Given this assumption, quasi-experimental variation in asset supply can

identify a “relative elasticity”—the difference between an asset’s own- and cross-price elasticities relative to similar assets—but not the absolute elasticity. The main cost of this approach is that one cannot recover the slope of an asset-level demand function, which may be necessary for many questions of interest.

Fully structural models as in [Kojien and Yogo \(2019\)](#) offer another alternative. By explicitly modeling investor preferences and equilibrium price formation, they can, in principle, recover demand parameters even in the presence of spillovers. However, this comes at the cost of stronger assumptions. As [Fuchs, Fukuda, and Neuhann \(2025\)](#) show, however, misspecification of functional forms or substitution patterns can lead to systematically biased estimates and unreliable counterfactuals. Hence, structural models must be designed to account for the cross-asset spillovers which form the basis of equilibrium portfolio choice and price determination. Moreover, because (arbitrarily) many theoretical models may be consistent with the observed data, these models must be judged on ex-ante theoretical considerations and plausibility, not on their empirical fit.

One can also estimate elasticities over portfolios rather than individual assets. One example of this approach is in [An and Huber \(2024\)](#), who estimate demand for small set of risk factors in foreign exchange markets. Our results suggest that this is a promising approach when assets are combined into portfolios which resemble Arrow securities.³ The key difficulty with estimating demand over portfolios is that it is harder to construct instruments that provide exogenous variation. If shocks originally occur at the asset level, then one requires a specific combination of asset-level shocks to generate a portfolio-level supply shifter. This often requires knowledge of the latent payoff matrix: see Remark 3.

6 Illustration in a general equilibrium model

We now illustrate our results on the inconsistency between state prices in the ideal experiment and states prices given a supply shock using a simple example economy with a log-utility representative investor based on [Fuchs, Fukuda, and Neuhann \(2025\)](#).

There are two assets and two states of the world, both denoted by g (green) and r (red). The probability of state $z \in \{g, r\}$ is $\pi_z \in (0, 1)$. The payoff profile of asset $j \in \{g, r\}$ is $y_j = (y_j(g), y_j(r))$. The aggregate endowments satisfy $E_r = 1$ and $E_g = 1 + s_g$, where

³Given a generic payoff matrix, Arrow securities are long-short portfolios: see [An \(2025\)](#) for an analysis of demand elasticities of long-short portfolios.

s_g is a supply shock to the green asset.

Table 1 depicts the payoff matrix. Markets are complete, and parameter $\epsilon \in (0, 1)$ determines the degree of complementarity between green and red assets. In the limit $\epsilon \rightarrow 0$, green and red assets are perfect substitutes with respect to their cash flows. The assets become more complementary as ϵ increases. In the limit $\epsilon \rightarrow 1$, the green and red assets are Arrow securities paying exactly one unit in one state of the world.

	State g (π_g)	State r (π_r)
Asset g	$\frac{1}{2}(1 + \epsilon)$	$\frac{1}{2}(1 - \epsilon)$
Asset r	$\frac{1}{2}(1 - \epsilon)$	$\frac{1}{2}(1 + \epsilon)$

Table 1: Payoff matrix.

6.1 Illustration of Trilemma

Prices. Denote by q_g and q_r the state prices measuring the cost of a unit of consumption in states g and r , respectively. Under no (unbounded) arbitrage, asset prices satisfy:

$$\begin{bmatrix} p_g \\ p_r \end{bmatrix} = \begin{bmatrix} y_g(g) & y_g(r) \\ y_r(g) & y_r(r) \end{bmatrix} \begin{bmatrix} q_g \\ q_r \end{bmatrix}. \quad (6)$$

We can invert this expression to solve for state prices as a function of the asset prices:

$$\begin{bmatrix} q_g \\ q_r \end{bmatrix} = \frac{1}{2\epsilon} \begin{bmatrix} (1 + \epsilon) & -(1 - \epsilon) \\ -(1 - \epsilon) & (1 + \epsilon) \end{bmatrix} \begin{bmatrix} p_g \\ p_r \end{bmatrix}.$$

Demand. Since markets are complete, we can solve the decision problem in terms of state-contingent consumption. Let c_z denote quantities of Arrow securities chosen by the investor, and let q_z the associated state prices. The decision problem is:

$$\begin{aligned} \max_{(c_0, c_g, c_r)} & (1 - \delta)u(c_0) + \delta\pi_g u(c_g) + \delta\pi_r u(c_r) \\ \text{s.t.} & c_0 + \sum_{z \in \{g, r\}} q_z c_z = e_0 + \sum_{z \in \{g, r\}} q_z (y_g(z)e_g + y_r(z)e_r). \end{aligned}$$

The necessary and sufficient optimality condition for Arrow security $z \in \{g, r\}$ is

$$q_z = \pi_z \frac{\delta}{1 - \delta} \frac{u'(c_z)}{u'(c_0)}. \quad (7)$$

Given the budget constraint, this condition determines optimal consumption as a function of Arrow prices. Consumption can then be mapped back into asset positions.

State prices in the ideal experiment. Consider the ideal experiment where a given investor faces an exogenous increase in the price of the green asset p_g while p_r remains fixed. Consistently with Lemma 2, the induced change in state prices is

$$\frac{\partial}{\partial p_g} \begin{bmatrix} q_g \\ q_r \end{bmatrix} = \frac{1}{y_g(g)y_r(r) - y_g(r)y_r(g)} \begin{bmatrix} y_r(r) \\ -y_r(g) \end{bmatrix} = \frac{1}{2\epsilon} \begin{bmatrix} 1 + \epsilon \\ -(1 - \epsilon) \end{bmatrix}.$$

A pure shock to p_g thus raises the cost of consumption in state g , but *lowers* it in state r . The reason is that replicating an Arrow security on the green asset requires going long the green asset and shorting the red asset, while replicating a red Arrow security requires going long the red asset and shorting the green asset. Holding p_r fixed, a change in p_g thus has the opposite effect on state prices in the two states of the world. Estimating the demand elasticity associated with this experiment thus requires a shock that triggers precisely this price variation.

State prices after a supply shock. We now show that supply shocks do not create the ideal state price variation. Market clearing requires consumption to equal available resources in every state:

$$c_z = y_g(z)(1 + s_g) + y_r(z).$$

Hence equilibrium state prices as a function of supply s_g are:

$$q_g = \pi_g \frac{\delta}{1 - \delta} \cdot \frac{1}{1 + \frac{1+\epsilon}{2}s_g} \quad \text{and} \quad q_r = \pi_r \frac{\delta}{1 - \delta} \cdot \frac{1}{1 + \frac{1-\epsilon}{2}s_g}. \quad (8)$$

In contrast to the ideal experiment, it is apparent that a negative supply shock to the green asset increases *both* state prices. In particular, differentiating q_z with respect to s_g and

evaluating in the limit $s_g \rightarrow 0$ yields

$$\left. \frac{\partial q_g}{\partial s_g} \right|_{s_g \rightarrow 0} = -\pi_g \frac{\delta}{1-\delta} \frac{1+\epsilon}{2} < 0 \quad \text{and} \quad \left. \frac{\partial q_r}{\partial s_g} \right|_{s_g \rightarrow 0} = -\pi_r \frac{\delta}{1-\delta} \frac{1-\epsilon}{2} < 0.$$

The reason is that the green asset pays off in both states of the world. Unfortunately, this means that the supply shock generates a state price change that is of the *wrong sign* compared to the ideal experiment.

The only exception is when $\epsilon = 1$, so that the payoff matrix is the identity matrix. In line with our theoretical results, this is because the underlying assets *are* Arrow assets, and these do not generate cross-asset spillovers to other assets. However, for all other ϵ even a purely exogenous supply shock does not generate the right variation.

Implications for demand. The fact that the supply shock generates the wrong type of state price variation dramatically affects the observed demand response. We illustrate this effect by computing the response of the consumption ratio c_g/c_r to both the ideal experiment and the supply shock. Given log utility, it follows from the first-order conditions (7) that the relative consumption process satisfies:

$$\frac{c_g}{c_r} = \frac{\pi_g q_r}{\pi_r q_g}. \quad (9)$$

Relative consumption in turn determines the desired holdings of green and red assets.

Consider first the ideal experiment with a pure price shock. Differentiating the relative consumption with respect to p_g and evaluating in the limit $s_g \rightarrow 0$ yields:

$$-\left. \frac{\partial}{\partial p_g} \left(\frac{c_g}{c_r} \right) \right|_{s_g \rightarrow 0} = \frac{1-\delta}{\delta} \frac{(1-\epsilon)\pi_g + (1+\epsilon)\pi_r}{2\pi_g\pi_r\epsilon}. \quad (10)$$

This derivative diverges to infinity as $\epsilon \rightarrow 0$. As the two assets are perfect substitutes in this limit, a small *price* shock triggers a rapid reallocation from green to red assets.

Next, consider the response to the supply shock. In the limit as $s_g \rightarrow 0$,

$$\left. \frac{\partial}{\partial s_g} \left(\frac{c_g}{c_r} \right) \right|_{s_g \rightarrow 0} = \epsilon. \quad (11)$$

which converges to *zero* in the limiting case of perfect substitutes as $\epsilon \rightarrow 0$. When the two

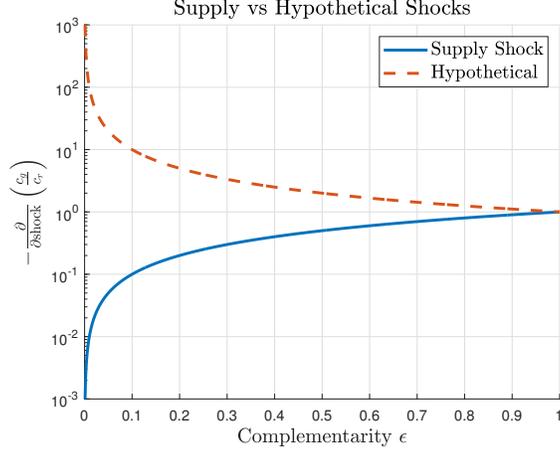


Figure 1: Optimal change in consumption ratio c_g/c_r on log scale. Parameters: $\pi_g = \pi_r = \frac{1}{2}$ and $\delta = \frac{1}{3}$.

assets are perfect substitutes, a supply shock has identical effects in both states. As such, it results in *zero* difference in the optimal consumption ratio across the two states.

Figure 1 depicts the optimal investor-level response to the hypothetical price shock (10) and the response to the supply shock (11) on log scale (Appendix A.4 provides the derivations of these expressions). The difference in responses diverges to infinity as $\epsilon \rightarrow 0$. The only point of overlap occurs when the two assets are both Arrow securities. In line with our theory, this is the case where there can be no spillovers across assets.

7 Conclusion

We show that asset-level supply shocks generally fail to generate the *ceteris paribus* price variation required to identify demand elasticities in asset markets. This reflects a fundamental trilemma: no arbitrage, payoff-based preferences, and identification from supply shocks cannot all hold simultaneously in observational data. As a result, elasticity estimates must rely on structural assumptions that cannot be directly validated. Our findings highlight the limitations of reduced-form approaches and underscore the importance of carefully specified models in empirical asset demand estimation.

References

- ADRIAN, T., E. ETULA, AND T. MUIR (2014): “Financial Intermediaries and the Cross-Section of Asset Returns,” Journal of Finance, 69(6), 2558–2596. [Cited on page 5.]
- ALLEN, J., J. KASTL, AND M. WITTEWER (2025): “Estimating Demand Systems with Bidding Data,” Working Paper. [Cited on page 6.]
- AN, Y. (2025): “Demand Elasticity for Portfolios,” Working paper, Johns Hopkins University. [Cited on pages 6 and 21.]
- AN, Y., AND A. W. HUBER (2024): “Demand Propagation Through Traded Risk Factors,” Discussion paper. [Cited on pages 6 and 21.]
- BERRY, S. T., AND P. A. HAILE (2021): “Foundations of Demand Estimation,” Handbook of Industrial Organization, 4(1), 1–62. [Cited on page 4.]
- BINSBERGEN, J. H. V., B. DAVID, AND C. C. OPP (2025): “How (Not) to Identify Demand Elasticities in Dynamic Asset Markets,” Working paper. [Cited on page 16.]
- DUFFIE, D. (2001): Dynamic Asset Pricing Theory. Princeton University Press, third edn. [Cited on pages 9, 10, and 30.]
- EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” Econometrica, 57(4), 937–969. [Cited on page 14.]
- ERSHOV, D., J.-W. LALIBERTÉ, M. MARCOUX, AND S. ORR (2024): “Estimating Complementarity With Large Choice Sets: An Application to Mergers,” Working Paper. [Cited on page 5.]
- FOSGERAU, M., J. MONARDO, AND A. DE PALMA (2024): “The Inverse Product Differentiation Logit Model,” American Economic Journal: Microeconomics, 16(4), 329–370. [Cited on page 5.]
- 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 6 and 21.]

- GABAIX, X., AND R. KOIJEN (2020): “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis,” mimeo Harvard and Chicago. [Cited on page 4.]
- HADDAD, V., Z. HE, P. HUEBNER, P. KONDOR, AND E. LOUALICHE (2025): “Causal Inference for Asset Pricing,” Discussion paper. [Cited on pages 6 and 20.]
- HARRIS, L., AND E. GUREL (1986): “Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures,” Journal of Finance, 41(4), 815–829. [Cited on page 5.]
- HE, Z., AND A. KRISHNAMURTHY (2013): “Intermediary Asset Pricing,” American Economic Review, 103(2), 732–70. [Cited on page 5.]
- IARIA, A., AND A. WANG (2020): “Identification and Estimation of Demand for Bundles,” Working Paper. [Cited on page 5.]
- 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 4, 6, 9, 18, and 21.]
- (2025): “On the Theory and Econometrics of (Demand System) Asset Pricing,” Working Paper. [Cited on page 8.]
- KREPS, D. M., AND E. L. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” Econometrica, 46(1), 185–200. [Cited on page 14.]
- KRISHNAMURTHY, A., AND A. VISSING-JØRGENSEN (2012): “The Aggregate Demand for Treasury Debt,” Journal of Political Economy, 120(2), 233–267. [Cited on page 5.]
- PLEMMONS, R. J., AND R. E. CLINE (1972): “The Generalized Inverse of a Nonnegative Matrix,” Proceedings of the American Mathematical Society, 31(1), 46–50. [Cited on page 32.]
- ROSS, S. A. (2004): Neoclassical Finance. Princeton University Press, Princeton, NJ. [Cited on page 1.]
- SHLEIFER, A. (1986): “Do Demand Curves for Stocks Slope Down?,” Journal of Finance, 41(3), 579–590. [Cited on page 5.]

TOBIN, J. (1969): “A General Equilibrium Approach to Monetary Theory,” Journal of Money, Credit and Banking, 1(1), 15–29. [Cited on page 5.]

VAN DER BECK, P. (2021): “Flow-Driven ESG Returns,” Swiss Finance Institute Research Paper Series, 71(21). [Cited on page 4.]

WANG, A. (2024): “Demand with Complementarity: A Market-Share Approach,” Working Paper. [Cited on page 5.]

A Appendix

A.1 Section 2

Proof of Proposition 1. For the first statement, let $a^{*i} \in \mathcal{A}^i$ be a solution to (1). For ease of exposition, we allow 0 to be in the domain of u^i (this is not essential). Suppose to the contrary that there is an unbounded arbitrage opportunity. Since u^i is strictly increasing, there exists $m > 0$ such that

$$U^i(a^{*i}) < (1 - \delta^i)u^i(e_0^i + p \cdot e^i) + \delta^i \pi_z u^i(m) + \delta^i (1 - \pi_z) u^i(0) \text{ for some } z$$

and

$$U^i(a^{*i}) < (1 - \delta^i)u^i(e_0^i + p \cdot e^i + m) + \delta^i u^i(0),$$

where $e^i \equiv (e_j^i)_{j=1}^J$. Since there is an unbounded arbitrage opportunity, for this $m > 0$, there exists $a^i \in \mathcal{A}^i$ such that either (i) $p \cdot a^i \leq 0$, $Y^T a^i \geq 0$, and $(Y^T a^i)_z \geq m$, in which case

$$U^i(a^{*i}) < (1 - \delta^i)u^i(e_0^i + p \cdot e^i) + \delta^i \pi_z u^i(m) + \delta^i (1 - \pi_z) u^i(0) \leq U^i(a^i)$$

or (ii) $p \cdot a^i \leq -m$ and $Y^T a^i \geq 0$, in which case

$$U^i(a^{*i}) < (1 - \delta^i)u^i(e_0^i + p \cdot e^i + m) + \delta^i u^i(0) \leq U^i(a^i).$$

In either way, $a^{*i} \in \mathcal{A}^i$ does not solve (1), a contradiction.

For the second statement, since there is no unbounded arbitrage opportunity, there

exists $m > 0$ such that, for any $a^i \in \mathcal{A}^i$,

$$U^i(a^i) < (1 - \delta^i)u^i(e_0^i + p \cdot e^i + m) + \delta^i u^i(m).$$

Thus, we obtain:

$$\sup_{a^i \in \mathcal{A}^i} U^i(a^i) \leq (1 - \delta^i)u^i(e_0^i + p \cdot e^i + m) + \delta^i u^i(m) < \infty.$$

Then, there exists a sequence $(a^{n,i})_{n \in \mathbb{N}}$ from \mathcal{A}^i such that

$$\sup_{a^i \in \mathcal{A}^i} U^i(a^i) - \frac{1}{n} < U^i(a^{n,i}) \leq \sup_{a^i \in \mathcal{A}^i} U^i(a^i) < \infty \text{ for all } n \in \mathbb{N}.$$

Since $\sup_{a^i \in \mathcal{A}^i} U^i(a^i) < \infty$, it follows that

$$\sup_{n \in \mathbb{N}} |a_j^{n,i}| < \infty \text{ for all } j \in \{1, \dots, J\}.$$

Since \mathcal{A}^i is closed, it follows that there exists a convergent subsequence $(a^{n_k,i})_{k \in \mathbb{N}}$ of $(a^{n,i})_{n \in \mathbb{N}}$ such that $a^{n_k,i} \rightarrow a^{*i} \in \mathcal{A}^i$. Since U^i is continuous, it follows that

$$U^i(a^{*i}) = \sup_{a^i \in \mathcal{A}^i} U^i(a^i),$$

as desired. ■

Proof of Lemma 1. Suppose the conditions in the statement of the lemma. The proof consists of seven steps. First, for each $i \in I_0$, we define a subset M^i of \mathbb{R}^{Z+1} :

$$M^i \equiv \{(-p \cdot a^i, Y^T a^i) \in \mathbb{R}^{Z+1} \mid a^i \in \mathcal{A}^i\}.$$

Then, for each $i \in I_0$, since investor i does not have an unbounded arbitrage opportunity, it follows that

$$M^i \cap \mathbb{R}_+^{Z+1} = \{0\}.$$

Note that \mathbb{R}_+^{Z+1} is a closed convex cone in \mathbb{R}^{Z+1} and does not contain any linear subspace other than $\{0\}$.

Second, let

$$M \equiv \bigcup_{i \in I_0} M^i.$$

It follows from the assumption

$$\mathbb{R}^J = \bigcup_{i \in I_0} \mathcal{A}^i$$

that

$$M = \{(-p \cdot a^i, Y^T a^i) \in \mathbb{R}^{Z+1} \mid a^i \in \mathbb{R}^J\}$$

is a linear subspace.

Third, since

$$M \cap \mathbb{R}_+^{Z+1} = \{0\},$$

it follows from the separating hyperplane theorem (which is referred to as “Linear separation of Cones” in [Duffie \(2001\)](#)), there exists $\bar{q} \in \mathbb{R}^{Z+1} \setminus \{0\}$ such that

$$\bar{q} \cdot t < \bar{q} \cdot x \text{ for all } t \in M \text{ and } x \in \mathbb{R}_+^{Z+1}.$$

Fourth, we show that $\bar{q} \in \mathbb{R}_{++}^{Z+1}$. Since $0 \in M$, it follows that

$$0 = \bar{q} \cdot 0 < \bar{q} \cdot x \text{ for all } x \in \mathbb{R}_+^{Z+1}.$$

Taking x as standard unit vectors in \mathbb{R}_{++}^{Z+1} yields $\bar{q}_z > 0$ for all z .

Fifth, we show that

$$0 = \bar{q} \cdot t \text{ for all } t \in M.$$

Suppose to the contrary that $0 \neq \bar{q} \cdot t$ for some $t \in M$. Since M is a linear subspace, we can assume, without loss, that

$$\bar{q} \cdot t > 0.$$

However, this leads to a contradiction because, for any given $x \in \mathbb{R}_{++}^{Z+1}$, there exists $\lambda \in \mathbb{R}$ such that $\lambda t \in M$ and

$$\bar{q} \cdot x \leq \lambda(\bar{q} \cdot t) = \bar{q} \cdot (\lambda t).$$

Sixth, we show that

$$\bar{q}^T \begin{bmatrix} -p^T \\ Y^T \end{bmatrix} = 0.$$

It follows from the fifth step that

$$\bar{q}^T \begin{bmatrix} -p^T \\ Y^T \end{bmatrix} a = 0 \text{ for all } a \in \mathbb{R}^J = \bigcup_{i \in I_0} \mathcal{A}^i.$$

If

$$\bar{q}^T \begin{bmatrix} -p^T \\ Y^T \end{bmatrix} \neq 0,$$

then letting

$$a = \left(\bar{q}^T \begin{bmatrix} -p^T \\ Y^T \end{bmatrix} \right)^T \in \mathbb{R}^J = \bigcup_i \mathcal{A}^i$$

yields

$$\bar{q}^T \begin{bmatrix} -p^T \\ Y^T \end{bmatrix} a > 0,$$

a contradiction.

Seventh, then, denoting by

$$\bar{q} = (q_0, q_{-0}),$$

we have

$$\bar{q}_0 p^T = q_{-0}^T Y^T, \text{ that is, } p = Y \frac{q_{-0}}{q_0}.$$

Letting $q = \frac{q_{-0}}{q_0} \in \mathbb{R}_{++}^Z$, we finally obtain

$$p = Yq,$$

as desired. ■

A.2 Section 4

Proof of Theorem 1. First, we show that Condition 1 implies that YY^T is diagonal. Suppose $Y^+ = -VY^TK$ for some diagonal matrix $K \equiv \text{diag}(k_1, \dots, k_J)$.

Operating Y on both sides from the left,

$$I = -YVY^TK,$$

where both sides are a $J \times J$ matrix. The (j, j') element of the right-hand side is

$$\begin{cases} -\sum_{z=1}^Z y_j(z)v_z y_{j'}(z)k_{j'} = 1 & \text{if } j = j' \\ -\sum_{z=1}^Z y_j(z)v_z y_{j'}(z)k_{j'} = 0 & \text{if } j \neq j' \end{cases}.$$

This implies that $k_j \neq 0$ for all j . Then,

$$\begin{cases} \sum_{z=1}^Z y_j(z)v_z y_{j'}(z) \neq 0 & \text{if } j = j' \\ \sum_{z=1}^Z y_j(z)v_z y_{j'}(z) = 0 & \text{if } j \neq j' \end{cases}.$$

Since $y_j(z), y_{j'}(z) \geq 0$, and $v_z > 0$, it follows that

$$\begin{cases} \sum_{z=1}^Z y_j(z)y_{j'}(z) \neq 0 & \text{if } j = j' \\ \sum_{z=1}^Z y_j(z)y_{j'}(z) = 0 & \text{if } j \neq j' \end{cases}.$$

Hence, YY^T is diagonal.

Second, we show that Condition 2 implies that YY^T is diagonal. Since Y is a $J \times Z$ matrix with $J \leq Z$ and $\text{rank}(Y) = J$, the Moore-Penrose pseudo-inverse is given by $Y^+ = Y^T(YY^T)^{-1}$. By [Plemmons and Cline \(1972, Theorem 1\)](#), the pseudo-inverse Y^+ is non-negative if and only if there exists a diagonal matrix with positive elements $D \equiv \text{diag}(d_1, \dots, d_Z)$ such that

$$Y^+ = DY^T.$$

Then, operating Y from the left,

$$I = YY^T(YY^T)^{-1} = YDY^T.$$

Then, extracting the (j, k) element (with $j \neq k$) from each of both sides,

$$0 = \sum_{z=1}^Z y_j(z)d_z y_k(z).$$

Since $y_j(z) \geq 0$, $d_z > 0$, and $y_k(z) \geq 0$ for all $z \in \{1, \dots, Z\}$, it follows that

$$y_j(z)y_k(z) = 0 \text{ for all } z \in \{1, \dots, Z\}.$$

This implies that the (j, k) element (with $j \neq k$) of YY^T is 0:

$$0 = \sum_{z=1}^Z y_j(z)y_k(z). \quad (12)$$

Thus, YY^T is a diagonal matrix.

We remark that the converse also holds. Suppose that YY^T is a diagonal matrix. Since YY^T is invertible under Assumption 1, $(YY^T)^{-1}$ is a diagonal matrix with positive entries. Since Y is non-negative, so is Y^T . Then, $Y^+ = Y^T(YY^T)^{-1}$ is non-negative.

Third, we show that, given that YY^T is diagonal, there are no assets with overlapping payoffs. Since YY^T is invertible, it is a diagonal matrix with positive elements. Equation (12) implies that, for any $z \in \{1, \dots, Z\}$, there exists at most one $j \in \{1, \dots, J\}$ such that $y_j(z) > 0$.

Fourth, we show that if markets are complete then YY^T is diagonal if and only if Y has exactly one non-zero element in each row and in each column (so that Y is a diagonal matrix up a re-ordering of rows or columns). If YY^T is diagonal, then its (j, k) element is:

$$\begin{cases} \sum_{z=1}^Z y_j(z)y_j(z) > 0 & \text{if } j = k \\ \sum_{z=1}^Z y_j(z)y_k(z) = 0 & \text{if } j \neq k \end{cases}.$$

Hence, for each row j , there exists exactly one element z such that $y_j(z) > 0$. Thus, Y has J non-zero elements. Since Y is square and invertible, for each column z , there exists exactly one element j such that $y_j(z) > 0$.

Conversely, if Y has exactly one non-zero element in each row and in each column, then

$$\begin{cases} \sum_{z=1}^Z y_j(z)y_j(z) > 0 & \text{if } j = k \\ \sum_{z=1}^Z y_j(z)y_k(z) = 0 & \text{if } j \neq k \end{cases}.$$

Thus, YY^T is diagonal. ■

A.3 Section 5

Proof of Proposition 2. The least-squares objective is

$$Q(S) = \|\Delta A_i - SG\|_F^2,$$

where $\|\cdot\|_F$ denotes the Frobenius (i.e., matrix) 2-norm. Differentiating with respect to S and setting the first-order condition to zero gives

$$-2(\Delta A_i - SG)G^T = 0, \quad \text{that is,} \quad \Delta A_i G^T = S(GG^T).$$

If G has full row rank, then GG^T is invertible. Thus, the unique solution is

$$\hat{S}_i = \Delta A_i G^T (GG^T)^{-1}.$$

Since $G^+ = G^T (GG^T)^{-1}$ when G has full row rank, we have $\hat{S}_i = \Delta A_i G^+$.

Unbiasedness follows from assumed exogeneity of supply shocks,

$$\begin{aligned} \mathbb{E}[\hat{S}_i | G] &= \mathbb{E}[(S_i G + U_i)G^T (GG^T)^{-1} | G] \\ &= S_i GG^T (GG^T)^{-1} + \mathbb{E}[U_i G^T (GG^T)^{-1} | G] \\ &= S_i. \end{aligned}$$

For consistency, assume $(1/N)GG^T \rightarrow Q \succ 0$ for some positive definite Q . Further let $\mathbb{E}[\|U_i\|_F^2] < \infty$, and $\mathbb{E}[U_i | G] = 0$. Then

$$\hat{S}_i - S_i = U_i G^T (GG^T)^{-1} = \left(\frac{1}{N}U_i G^T\right) \left(\frac{1}{N}GG^T\right)^{-1} \xrightarrow{p} 0,$$

by a law of large numbers for the cross-experiment averages. Thus \hat{S}_i is consistent. ■

Proof of Proposition 3. Similarly to the proof of Proposition 3, the first-order condition is $S_i(GG^T) = \Delta A_i G^T$, from which we obtain $S_i G = \Delta A_i$. Multiplying G^+ from the right, we obtain $S_i P_G = \Delta A_i G^+$ as in the main text. Since this is a particular solution, the general solution can be written as

$$S_i = \Delta A_i G^+ + B_i(I_J - P_G),$$

where $B_i \in \mathbb{R}^{J \times J}$ is an arbitrary matrix and I_J is the identity matrix. ■

A.4 Section 6

First, we derive equation (10). Since the Arrow prices q can be expressed as a function of the asset prices p through equation (6), the consumption ratio (9) can be written as:

$$\frac{c_g}{c_r} = \frac{\pi_g (1 + \epsilon)p_r - (1 - \epsilon)p_g}{\pi_r (1 + \epsilon)p_g - (1 - \epsilon)p_r}.$$

Thus, differentiating it with respect to the price p_g , we have:

$$\frac{\partial}{\partial p_g} \left(\frac{c_g}{c_r} \right) = -\frac{\pi_g}{\pi_r} \frac{4\epsilon p_r}{((1 + \epsilon)p_g - (1 - \epsilon)p_r)^2}. \quad (13)$$

In contrast, substituting the Arrow prices (8) into equation (6), we obtain:

$$\begin{aligned} p_g &= \frac{1 + \epsilon}{2} \pi_g \frac{\delta}{1 - \delta} \frac{1}{1 + \frac{1 + \epsilon}{2} s_g} + \frac{1 - \epsilon}{2} \pi_r \frac{\delta}{1 - \delta} \frac{1}{1 + \frac{1 - \epsilon}{2} s_g}; \\ p_r &= \frac{1 - \epsilon}{2} \pi_g \frac{\delta}{1 - \delta} \frac{1}{1 + \frac{1 + \epsilon}{2} s_g} + \frac{1 + \epsilon}{2} \pi_r \frac{\delta}{1 - \delta} \frac{1}{1 + \frac{1 - \epsilon}{2} s_g}. \end{aligned}$$

Substituting the asset prices p at $s_g = 0$ into equation (13), we obtain equation (10). When $\delta = \frac{1}{3}$ and $\pi_g = \pi_r = \frac{1}{2}$, equation (10) reduces to:

$$-\frac{\partial}{\partial p_g} \left(\frac{c_g}{c_r} \right) \Big|_p = \frac{1}{\epsilon}.$$

Second, we derive equation (11). Substituting the Arrow prices (8) into the consumption ratio (9) yields

$$\frac{c_g}{c_r} = \frac{1 + \frac{1 + \epsilon}{2} s_g}{1 + \frac{1 - \epsilon}{2} s_g}.$$

Thus, differentiating it with respect to the supply shock s_g , we obtain

$$\frac{\partial}{\partial s_g} \left(\frac{c_g}{c_r} \right) = \frac{\epsilon}{\left(1 + \frac{1 - \epsilon}{2} s_g \right)^2}.$$

In the limit as $s_g \rightarrow 0$, we get equation (11).