

Demand-System Asset Pricing: Theoretical Foundations*

William Fuchs[†] Satoshi Fukuda[‡] Daniel Neuhann[§]

April 16, 2025

Abstract

Recent approaches to asset pricing use structural methods to estimate investor-level demand functions for financial assets. We show that cross-asset complementarities and price spillovers can significantly bias these estimates: if close substitutes exist, measured elasticities are near one even if true elasticities are near infinite. This reconciles low demand-system elasticities with higher theoretical benchmarks. Biases are smaller for less substitutable assets, such as broad portfolios or asset classes. Control variables lead to estimates of *residual* demand elasticities which may offer limited information about asset-level demand. We caution against interpreting estimated demand elasticities as structural parameters which remain stable under counterfactuals.

*Special thanks to Andres Almazan, Zhiguo He, Marco Loseto, Christian Opp (discussant), Andrey Ordin, and Andy Skrypacz for detailed conversations. We also thank Manuel Adelino, Hengjie Ai (discussant), Aydogan Altı, Philippe van der Beck (discussant), Alon Brav, Nicolae Garleanu (discussant), Valentin Haddad, Lars Hansen, Ben Hebert, Ali Hortacsu, Amy Huber, Paul Huebner, Travis Johnson, Peter Kondor, Jian Li, Erik Loualiche, Deborah Lucas, Aaron Pancost, Stavros Panageas (discussant), Anna Pavlova, Lukas Schmid, Dejanir Silva (discussant), Yannis Stamatopoulos, Sheridan Titman, Jonathan Wallen, Milena Wittwer, and seminar attendees at ASU Sonoran Winter Finance Conference, Bocconi, Cornell, Duke/UNC Asset Pricing Conference, 2024 Finance Theory Group Spring Meeting, 2024 FTG Conference on Bridging Theory and Empirical Research, 2024 FIRS conference, Finance Theory Webinar, Stanford GSB, UT Austin and Wharton for useful feedback. We thank Kevin Mei for excellent research assistance. Fuchs gratefully acknowledges the support from Grant PGC2018-096159-B-I00 financed by MCIN-AEI-10.13039-501100011033, and Comunidad de Madrid (Spain) through grants EPUC3M11 (V PRICIT) and H2019-HUM-5891.

[†]University of Texas at Austin, UC3M, FTG, and CEPR. william.fuchs@mcombs.utexas.edu.

[‡]Bocconi University and IGIER. satoshi.fukuda@unibocconi.it.

[§]University of Texas at Austin and FTG. daniel.neuhann@mcombs.utexas.edu.

1 Introduction

Recent approaches to asset pricing following [Kojen and Yogo \(2019\)](#) involve the structural estimation of investor-level demand functions for financial assets. Advocates of this approach argue that detailed data on portfolio holdings can be used to structurally identify investor-level demand parameters for specific assets, and that granular descriptions of individual demand functions offer new insights into the functioning of financial markets, including the equilibrium response to a wide array of counterfactuals, such as shocks to the wealth distributions, investor preferences, or policy interventions. Perhaps the most striking claim in this literature is that demand elasticities for financial assets are orders of magnitude lower than in standard models ([Kojen and Yogo, 2021](#)).

While compelling in its motivation, the promise of the demand-system approach ultimately hinges on its ability to accurately identify structural parameters of interest. Yet existing demand estimation techniques, including the logit approach in [Kojen and Yogo \(2019\)](#), were originally developed for settings that differ substantially from portfolio choice and asset pricing. For example, the prototypical demand estimation in industrial organization considers market-level analyses of discrete choices over consumption goods. Such settings generally do not feature many considerations which are central to asset pricing, including cross-asset demand complementarities within portfolios, variable quantities, dynamic trading, and general equilibrium price determination. Hence it is an open question whether current demand approaches can indeed identify structural parameters in settings where these features are critical. We answer this question in a canonical asset pricing model ([Lucas, 1978](#)) enriched with heterogeneous tastes for financial assets.¹

Our main finding is that current asset demand systems do not adequately account for cross-asset complementarities in portfolio choice, whereby the marginal value of an asset depends, in an asset-specific manner, on the investor's holdings of other assets. As such, asset demand systems may yield low *measured* demand elasticities even when true elasticities are near infinite. This offers a simple explanation for the striking difference between the low demand elasticities documented by leading demand systems estimates and the high demand elasticities obtained in standard models. More constructively, we

¹Taste differences are critical because they generate cross-sectional heterogeneity in portfolios, as in the data. They also make it possible to construct demand shocks to some investors that are suitably orthogonal to the demand of other investors. This is required for identification of demand functions.

also show that this bias is smaller when assets are not particularly substitutable, as may be the case for highly aggregated asset classes with little overlap in the cash-flow distribution. Nevertheless, we caution that demand elasticities, even if well-measured, can be interpreted as structural parameters only under stringent additional assumptions.

A simple thought experiment is instructive. Consider an investor who must choose between three assets: two closely substitutable “inside assets,” say Microsoft and Apple stocks, and a less substitutable “outside asset,” say a bond. Now consider a supply shock to Microsoft which raises its price. All else equal, the investor would like to buy Apple and finance this trade by selling Microsoft (a *demand complementarity*). Given this substitution within inside assets, she then finds it optimal to leave her holdings of the outside asset roughly unchanged (*heterogeneous substitution*). Yet if many investors pursue the same strategy, the price of Apple must increase (a *price spillover*), taking away any individual investor’s incentive to switch from Microsoft to Apple. In equilibrium, the investor’s portfolio is thus relatively unresponsive to exogenous variation in the price of Microsoft even though she would have responded very rapidly had the price of Apple remained fixed. That is, demand complementarities and price spillovers create a disconnect between observed elasticities (which incorporate all equilibrium adjustments) and structural elasticities (which counterfactually presume that other asset prices remain fixed).

Given this disconnect, identifying structural elasticities from observational data requires appropriately accounting for heterogeneous substitution and price spillovers. Yet current approaches place stark restrictions on substitution patterns and price spillovers. For example, logit demand systems following [Kojien and Yogo \(2019\)](#) assume that complementarities and spillovers can be accounted for by measuring demand for inside assets *relative* to the outside asset. Given this restriction, price spillovers between inside assets are immaterial, and observed changes in relative portfolio shares identify the elasticity of relative demand. Naturally, biases arise when this restriction fails, as is the case when inside assets are more substitutable with each other than with the outside asset.

Figure 1 illustrates the nature of this bias. In both panels, we consider an exogenous supply shock from S to S' , and identical observed equilibrium prices and quantities E and E' . The left panel depicts a demand system following the approach in [Kojien and Yogo \(2019\)](#) in which the demand curve determining the *relative* portfolio share of asset j is invariant in the quantities or prices of other inside assets. Given this assumption,

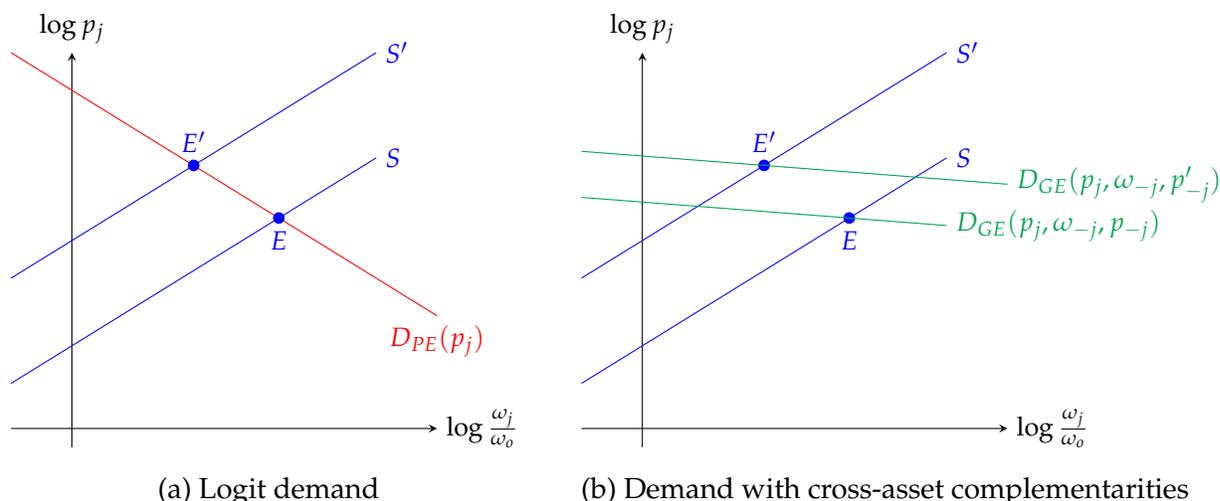


Figure 1: Elasticity measurement based on different demand system specifications. The left panel corresponds to logit demand for financial assets (Kojien and Yogo, 2019), whereby relative demand for a given asset is invariant in the prices and quantities of other assets (as in Partial Equilibrium). The right panel allows for cross-asset complementarities (as in General Equilibrium). The supply curves and the observed equilibrium allocations are identical in both panels. We use ω_j to refer to portfolio share of asset j , and ω_{-j} to denote the vector of portfolio shares of assets other than j . Analogous definitions hold for asset prices p_j . Demand is measured in units of portfolio shares relative to ω_o , the portfolio share of the outside asset.

observed portfolio changes are interpreted as a move *along* a relatively *inelastic* demand curve. The right panel depicts a demand system in which demand for asset j depends on the holdings and prices of other inside assets, and these respond endogenously to the supply shock. The observed demand response is now rationalized by *high* elasticities and a *shift* of the demand curve. Because standard models naturally generate demand curves with complementarities (even when measured in relative terms), this mechanism can account for the dramatic difference between low measured elasticities in the demand-system approach and much higher elasticities in standard models.

We derive this bias formally by decomposing the difference between measured and structural elasticities into the product of *demand complementarities* (i.e., the cross-elasticity between the focal asset and all potential substitutes and complements) and *price spillovers* (the equilibrium response of other asset to a shock to a given asset). Hence the bias is large whenever close substitutes are available and substitute assets are quickly repriced in response to shocks. That is, the bias is large precisely when markets are elastic.

One suggested approach to the problem of heterogeneous substitution and spillovers is the use of control variables. A given supply shock may trigger spillovers primarily to other assets with similar factor exposures. In this case, controlling for common factor

exposures mitigates the resulting asymmetry in substitution. However, controls also directly alter the degree of substitutability between choices: two assets may be substitutable precisely *because* they have common exposures. The use of controls then yields demand elasticities defined over the *residual* cash flows unaccounted for by controls. If these residual cash flows are less substitutable than the asset itself, the resulting demand elasticities are naturally lower and may carry little information about asset-level elasticities.

In the final part of our analysis, we ask whether financial demand elasticities, even if well-measured, should be interpreted as structural parameters that are likely to be invariant under counterfactuals. To do so, we incorporate another feature that distinguishes financial markets from many goods markets, which is that assets are investment goods whose current value critically depends on their resale price. Using a dynamic variant of our model, we derive demand functions for financial assets that depend both on the investor's private tastes and her expectations of market returns. As in a beauty contest (Keynes, 1936), demand elasticities alone thus cannot distinguish whether an investor's demand is due to her own tastes or her expectations of others' tastes. Yet estimating counterfactuals in many cases requires assigning tastes to a particular investor.

We establish a related result for unobservable portfolio constraints. For a range of parameters, tastes for a given asset are observationally equivalent to unobserved mandates that constrain an investor's portfolio choice. Yet an unconstrained investor will respond differently to a counterfactual price shock than a constrained investor.

Related Literature

Demand-system asset pricing is grounded in structural estimation of investor-level portfolio choice functions. This is a sharp break from neoclassical asset pricing, which has little interest in asset quantities and instead focuses on price data disciplined by no arbitrage (Ross, 2004). It also differs from existing approaches that do emphasize quantities, such as classical theories of portfolio balance (Tobin, 1969), convenience yields in Treasury markets (Krishnamurthy and Vissing-Jorgensen, 2012), intermediary asset pricing (He and Krishnamurthy, 2013; Adrian, Etula, and Muir, 2014), capital flows due to index inclusion or other market frictions (Shleifer, 1986; Harris and Gurel, 1986), which emphasize *aggregate* demand effects in certain asset classes or markets, but stop short of

structurally estimating investor-level demand functions for specific assets.² The fact that estimated demand systems appear to reveal that financial institutions exhibit low demand elasticities has also been used to argue that financial markets as a whole are inelastic, with implications for the equity premium (Gabaix and Koijen, 2020).

Demand elasticities implied by the aforementioned literature on capital flows are broadly similar to those found in demand-based approaches. However, the goal of this literature is not to isolate structural elasticities (in which all other prices are held fixed) from general equilibrium elasticities (which incorporate all adjustments). That both approaches find similar elasticities can therefore be explained by the fact that they ultimately estimate similar objects (namely, the general equilibrium elasticity). Our analysis here focuses on approaches that aim to estimate structural elasticities and parameters.

Several papers build on the logit demand system to study substantive questions, including effects of counterfactual wealth distributions (Koijen, Richmond, and Yogo, 2024), global imbalances and currencies (Jiang, Richmond, and Zhang, 2023), corporate bond markets (Bretscher, Schmid, Sen, and Sharma, 2022; Darmouni, Siani, and Xiao, 2023), asset purchase programs (Breckenfelder and De Falco, 2023), bond market substitution (Nenova, 2025), and stock market competitiveness (Haddad, Huebner, and Loualiche, 2025). Davis, Kargar, and Li (2025) share our interest in accounting for low measured elasticities but use a partial equilibrium approach without endogenous price spillovers. While these studies extend the scope of asset demand systems in important ways, they do not address the specific issues of complementarities and spillovers we discuss here.

There are three main approaches to addressing the issue of complementarities and spillovers, each of which offers distinct advantages depending on the application. The first approach is using richer structural approaches that directly model complementarities and spillovers. Unfortunately, there are limited methods for dealing with asset-level complementarities in demand estimation (Berry and Haile, 2021).³ In the context of financial markets, Allen, Kastl, and Wittwer (2025) estimate a model of demand complementarities in simultaneous auctions of treasury bonds. In contrast to equity markets, their setting

²See also additional studies of index inclusions (Chang, Hong, and Liskovich, 2015; Pavlova and Sikoraskaya, 2023; Greenwood and Sammon, 2025), and research on fund flows (Gompers and Metrick, 2001; Coval and Stafford, 2007; Lou, 2012; Ben-David, Li, Rossi, and Song, 2022; Hartzmark and Solomon, 2024; Li, 2025) and central bank interventions (Krishnamurthy and Vissing-Jorgensen, 2011; Selgrad, 2023).

³In industrial organization, existing research typically considers small choice sets with limited complementarities. Gentzkow (2007) studies newspaper demand with a choice between print, online, or both.

features a small number of assets and data on bid *schedules*, not just portfolio holdings.

The second approach is to impose additional structure on the substitution matrix. For example, [Kojien and Yogo \(2020\)](#) and [Chaudary, Fu, and Li \(2023\)](#) use nested logit demand systems to study international financial markets and corporate bond markets. Nested logit allows researchers to capture specific forms of heterogeneous substitutions. In the case of corporate bonds, for instance, it is reasonable to assert that substitution within investment grade bonds is easier than across rating groups. Accordingly, [Chaudary, Fu, and Li \(2023\)](#) find much larger elasticities when allowing for heterogeneous substitution. A different restriction is used in [An and Huber \(2024\)](#), who model substitution along a small number of factors in foreign exchange, and [Nenova \(2025\)](#), who allows for more flexible substitution patterns using a rich set of covariates. The main limitation of these approaches is that they require ex-ante restrictions on the substitution matrix, even though substitutability is endogenously determined alongside returns and portfolios. Moreover, these restrictions must be valid for the marginal investor, and spillovers may still occur within groups of assets. In line with our findings, appropriate restrictions may be easier to ascertain for aggregate portfolios than individual securities.

The third approach is the use of reduced-form methods using control variables or difference-in-difference estimators. For example, [van der Beck \(2022\)](#) argues that estimating demand elasticities from trades rather than positions eliminates some endogeneity concerns in instrumental variables, while [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#) show that, if one is willing to make strong symmetry assumptions on the substitution matrix across all assets in the choice set, one can identify *relative* elasticities between pairs of assets. The main limitation of this approach is that one cannot estimate the absolute elasticity, and that the required symmetry assumptions are unlikely to hold without additional conditioning information, such as control variables. As we discuss, the use of control variables generates estimates of a *residual relative* elasticity.

By allowing for tastes over assets, our paper also contributes to a growing literature in which investors hold securities because of non-pecuniary values ([Starks, 2023](#)). The literature includes the study of investment in “green assets” associated with sustainable investments.⁴ [Fuchs, Fukuda, and Neuhann \(2025\)](#) show that “tastes” are orthogonal to returns if investors have deontological preferences, but not if they are consequentialist.

⁴See, for example, [Pastor, Stambaugh, and Taylor \(2021\)](#) and [D’Amico, Klausmann, and Pancost \(2023\)](#).

2 Framework

We begin by describing the general framework we use throughout the paper. Our goal is two-fold: to clarify the theoretical underpinnings of asset demand systems, and to assess whether current approaches to demand estimation can accurately identify structural elasticities. Our approach is to specify a canonical general equilibrium economy based on [Lucas \(1978\)](#) and to evaluate the properties of asset demand systems within this framework. To accommodate features necessary to match heterogeneous portfolios and achieve identification, we enrich this model with investor “tastes,” defined as heterogeneous valuations for financial assets across different investors. We begin by reviewing the rationale for these tastes, and then formally describe our framework.

2.1 The importance of heterogeneous valuations

We start with reviewing key model ingredients necessary to allow for the identification of investor-level demand functions for financial assets from observational data. The main challenge is that quantities and prices are jointly determined in equilibrium. Hence simple regressions of quantities on prices do not identify structural parameters.

To address this issue, the literature on demand-system asset pricing focuses on variation in *net supply*, defined as aggregate supply minus demand of a subset of market participants. [Figure 2](#) illustrates this approach. The left panel shows the canonical supply and demand diagram in an endowment economy for financial assets where the supply curve S is vertical. In the panel, D^1 and D^2 are demand curves for individual market participants, and D^A is aggregate demand. The right panel shows a potential solution: if exogenous shocks to aggregate supply S are not available, researchers can still estimate the structural parameters of demand function D^1 by finding exogenous variation in *residual supply* $S - D^2$. That is, we construct exogenous shocks to the residual supply curve of a given investor by finding exogenous shocks to the demand functions of other investors.

This approach places stringent constraints on the variation that can be used to identify demand systems. In particular, researchers must find settings in which there are demand shocks to a subset of investors that is uncorrelated with the demand of other investors. In the context of financial markets, this implies that one must find changes to market prices that are not driven by correlated shocks to discount rates and/or expected

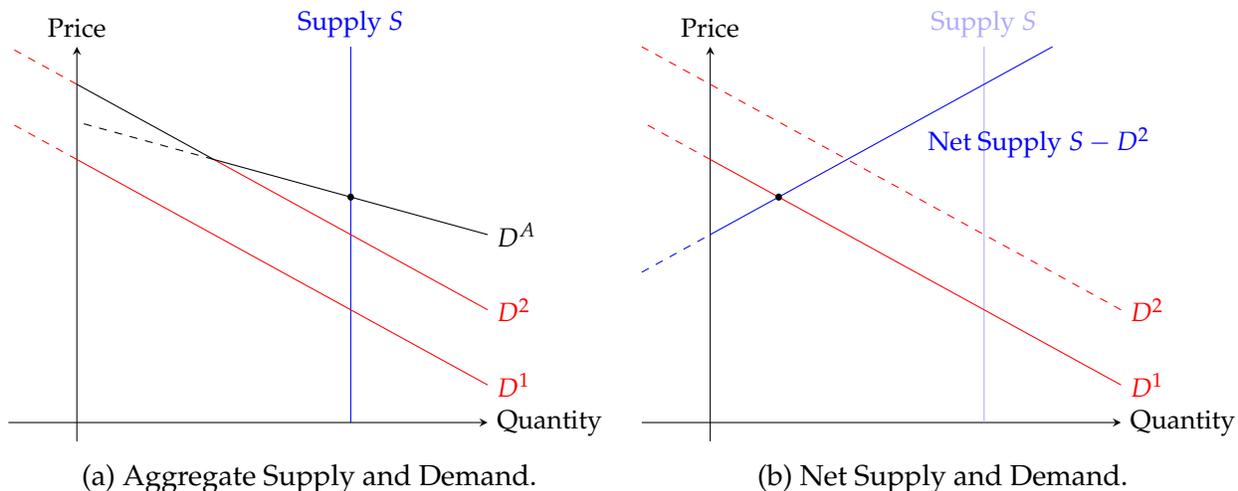


Figure 2: The basic identification issue in an endowment economy.

payoffs. Since financial assets are investment goods whose current value generically depends on their resale value, these requirements extend to expected future prices.

Two microfoundations for such shocks have been proposed. The first is cross-investor heterogeneity in *tastes* for particular assets, holding fixed a certain notion of expected cash flows. These could arise from differences in investor preferences over the provenance of cash flows, such as when some investors prefer to invest in environmentally-friendly firms. Or, investors may have dogmatic beliefs about returns that are orthogonal to the beliefs of other investors. While useful for identification, Appendix B shows that tastes can invalidate the principle of no arbitrage. The second is constraints or investment mandates that prevent some market participants from investing in a particular asset for *exogenous* reasons. We therefore use a framework that allows for both heterogeneous tastes (or dogmatic beliefs) and flexible constraints.

2.2 Formal model

We consider a one-shot portfolio choice problem in which an investor can choose to consume at date 0 and/or at date 1. Section 4.1 considers a simple dynamic extension.

There is a unit continuum of investors indexed by i . Investor i has a von Neumann-Morgenstern utility function defined over lotteries which determine the investor's consumption. A random state of the world $z \in \mathcal{Z} \equiv \{1, \dots, Z\}$ is realized at date 1, and the probability of state z is $\pi_z \in (0, 1)$. The set of assets is $\mathcal{J} \equiv \{1, \dots, J\}$. Asset $j \in \mathcal{J}$ offers

state-contingent cash flows $y_j(z)$ in state z . Investor i is endowed with e_j^i units of asset j and additional non-asset endowment w_0^i and $w_1^i(z)$ at dates 0 and 1, respectively. The aggregate endowment of asset j is $E_j = \int_i e_j^i di$.

Within this framework, we introduce payoff-augmenting tastes, defined as additional “consumption-equivalent” value that is generated by a particular asset. These taste parameters are designed to capture the heterogeneity in asset valuations that is required for a meaningful notion of exogenous shocks to residual demand. Formally, we say that investor i evaluates her payoffs from holding portfolio $a^i \equiv (a_j^i)_{j \in \mathcal{J}}$ by both the cash flows it generates and her *tastes* $\theta^i \equiv (\theta_j^i)_{j \in \mathcal{J}}$ over assets, where $\theta_j^i > 0$. Preferences are then defined over *effective units of consumption* delivered by portfolio a^i , and these are

$$\tilde{c}_1^i(z) \equiv \sum_{j \in \mathcal{J}} \theta_j^i y_j(z) a_j^i + w_1^i(z).$$

Investors may also be subject to portfolio constraints such as short-sale constraints, investment mandates, or restrictions on portfolio weights on particular asset classes. To capture these considerations in a flexible manner, we say that investor i faces $N \geq 0$ investment constraints on portfolio choices. The n -th investment constraint is defined as

$$M_n^i(a^i, p) \geq 0,$$

where a^i is investor i 's portfolio, $p \equiv (p_j)_{j \in \mathcal{J}}$ is the price vector, and the function $M_n^i(\cdot)$ is twice continuously differentiable in a_j^i for all j . We assume that the set of feasible portfolios induced by these constraints is convex.

Given these assumptions, investor i 's portfolio choice problem is:

$$\begin{aligned} \max_{a^i} \quad & (1 - \delta)u^i(c_0^i) + \delta \sum_{z \in \mathcal{Z}} \pi_z u^i(\tilde{c}_1^i(z)) & (1) \\ \text{s.t.} \quad & c_0^i = w_0^i - \sum_{j \in \mathcal{J}} p_j (a_j^i - e_j^i) \\ & \tilde{c}_1^i(z) = \sum_{j \in \mathcal{J}} \theta_j^i y_j(z) a_j^i + w_1^i(z) \text{ for all } z \\ & M_n^i(a^i, p) \geq 0 \text{ for all } n, \end{aligned}$$

where $\delta \in (0, 1]$ is a discount factor and the first two constraints are budget constraints at time 0 and time 1 in state z . We summarize investor i 's marginal valuations by the

taste-augmented marginal rate of substitution between state z and date 0, defined as

$$\tilde{\Lambda}^i(z) \equiv \frac{\delta \pi_z u^i(c_1^i(z))}{(1 - \delta) u^i(c_0^i)}$$

where u^i is marginal utility. Our equilibrium concept is competitive equilibrium.

Definition 1 (Competitive Equilibrium) *A competitive equilibrium consists of asset prices $(p_j)_{j \in \mathcal{J}}$ and investor portfolios $(a_j^i)_{j \in \mathcal{J}}$ for each i such that:*

1. *Given asset prices, investor portfolios solve decision problem (1) for each i .*
2. *The consumption goods market clears in every state.*
3. *Financial markets clear for every asset: $\int_i a_j^i di = E_j$ for each $j \in \mathcal{J}$.*

2.3 Model Discussion

Our model is designed to allow for a transparent and tractable analysis of demand estimation within a canonical asset pricing framework. The main departure is the introduction of tastes parameters which generate heterogeneous asset valuations across investors. We now briefly discuss some broader implications of tastes in asset pricing.

In our approach, tastes multiplicatively augment consumption. As we will show, they operate like “latent demand” shifters in [Kojien and Yogo \(2019\)](#). An alternative approach is to specify additive separable tastes, whereby investor obtains some additional value (or disutility) from holding certain assets that is separable from risk-return considerations. Both formulations deliver essentially identical conclusions.

Since portfolio choice requires a cardinal interpretation of utility, the intensity of tastes influences portfolio choice. This is in contrast to many settings in industrial organization where an ordinal ranking is sufficient. Asset demand systems thus require accurate identification of the specific values of a taste parameter. Moreover, the aggregation of assets into portfolios requires appropriately weighting tastes by marginal utility.

More broadly, there is a correspondence between tastes and heterogeneous beliefs. Specifically, the state space over which payoffs are defined can be enriched to include “taste-based payoffs.” Heterogeneous tastes then map into heterogeneous beliefs if investors differ in their subjective probabilities over this augmented state space. Impor-

tantly, such taste-related beliefs are dogmatic: investors agree to disagree, and in particular they may disagree on whether a particular state of the world can be realized. Such strong disagreement is *desirable* when trying to construct supply shocks because it allows for the possibility of orthogonal demand shocks. However, heterogeneous valuations are in tension with the organizing principle of no arbitrage. Appendix B shows that no arbitrage might fail given heterogeneous tastes, so that equilibrium prices may fail to exist.

3 Measuring and Interpreting Asset Demand Elasticities

In the previous section, we discussed the first main challenge of demand estimation, which is to develop conceptual frameworks for asset pricing which permit suitably exogenous shocks to residual supply. We now turn to the second main challenge, which is to develop a demand system which accurately identifies structural parameters from the data given this variation. We emphasize two critical factors that make this challenge particularly difficult in financial markets: (i) portfolio choice naturally exhibits *demand complementarities* whereby the marginal valuation of an asset depends on the rest of the investor's portfolio, and (ii) in market equilibrium, such demand complementarities generate *price spillovers* to other assets. Given these considerations, even clean exogenous variation in asset prices is generally not sufficient to identify structural parameters.

Throughout, we focus on estimating the elasticity of demand for asset j with respect to variation in some asset price p_s . This could be the asset's own price ($s = j$), or the price of another asset ($s \neq j$). Let $a_j^i(p, a_{-j}^i)$ denote investor i 's demand function for asset j , where p is the vector of asset prices and a_{-j}^i is the vector of investor i 's remaining asset positions.⁵ The elasticity of demand is the percentage change in i 's demand for asset j given a percentage change in the price of asset s , *holding all other prices fixed*:

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

We refer to \mathcal{E}_{jj} as the own-price elasticity and to \mathcal{E}_{js} ($s \neq j$) as the cross-price elasticity.

⁵Some implementations of the demand system approach define elasticities over portfolio shares rather than asset holdings, or consider demand relative to a benchmark asset. However, the appropriate units are sensitive to assumptions on investor utility and payoffs (Haddad, He, Huebner, Kondor, and Loualiche, 2025). Hence we use asset positions for now, and return to portfolio shares and relative demand later on.

3.1 The Identification Challenge

The first step is to clearly describe the identification challenge for demand systems in financial markets. To do so, we begin by deriving optimal portfolio choices (i.e., an investor's asset demand functions) from decision problem (1). Denote by λ_n^i the Lagrange multiplier associated with the n -th investment constraint, and by $m_{n,j}^i(a^i, p)$ the partial derivative of $M_n^i(a^i, p)$ with respect to a_j^i . Recall that $\tilde{\Lambda}^i(z)$ is the taste-adjusted marginal rate of substitution between date 0 and state z . Then investor i 's first-order necessary condition for a_j^i , her holdings of asset j , is

$$F_j^i(a^i, p) \equiv \theta_j^i \sum_{z \in \mathcal{Z}} y_j(z) \tilde{\Lambda}^i(z) + \sum_n \lambda_n^i \frac{m_{n,j}^i(a^i, p)}{(1-\delta)u'(c_0^i)} - p_j = 0. \quad (2)$$

Function $F_j^i(\cdot)$ has a natural interpretation as the marginal value of asset j net of the asset price and the shadow cost of constraints. Consequently, the optimal portfolio choice is determined by the condition that the net marginal value of every asset is equal to zero,

$$F^i(a^i, p) \equiv \begin{bmatrix} F_1^i(a^i, p) \\ F_2^i(a^i, p) \\ \vdots \\ F_j^i(a^i, p) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Importantly, this system exhibits *demand complementarities* whereby the *marginal* value of asset j generically depends on the quantities held of all other assets. Hence the willingness to substitute across assets is an endogenous object that depends on the entire vector of portfolio holdings. There are two natural sources of such complementarities. The first is the canonical notion of diversification, whereby the marginal value of an asset depends on its covariance with the rest of the investor's portfolio. The second is through constraints. If an investor faces a mandate to invest at least 50% of its assets in technology stocks, buying more of any given technology stock relaxes the constraint for all non-technology stocks. In either case, optimal asset positions, and the investor's willingness to substitute between assets, are inherently intertwined with each other.

To see how these considerations complicate inference of structural parameters, assume that we have an ideal instrument in hand: a purely exogenous supply shock χ_s that directly affects only asset s . For example, in line with instrumental variable strategy of

Koijen and Yogo (2019), we might imagine that an outside investor has decided to adjust her supply of asset s for purely exogenous reasons. Using our model, we can precisely describe how the investor responds to this shock. We totally differentiate the system of the first-order conditions $F^i(a^i, p) = \mathbf{0}$ with respect to shock χ_s .

By the implicit function theorem, the optimal change of i 's portfolio in response to an exogenous shock χ_s to asset s is then given by the system of equations

$$\begin{bmatrix} \frac{da_1^i}{d\chi_s} \\ \vdots \\ \frac{da_s^i}{d\chi_s} \\ \vdots \\ \frac{da_J^i}{d\chi_s} \end{bmatrix} = \begin{bmatrix} \frac{\partial a_1^i}{\partial p_1} & \cdots & \frac{\partial a_1^i}{\partial p_s} & \cdots & \frac{\partial a_1^i}{\partial p_J} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial a_s^i}{\partial p_1} & \cdots & \frac{\partial a_s^i}{\partial p_s} & \cdots & \frac{\partial a_s^i}{\partial p_J} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial a_J^i}{\partial p_1} & \cdots & \frac{\partial a_J^i}{\partial p_s} & \cdots & \frac{\partial a_J^i}{\partial p_J} \end{bmatrix} \begin{bmatrix} \frac{dp_1}{d\chi_s} \\ \vdots \\ \frac{dp_s}{d\chi_s} \\ \vdots \\ \frac{dp_J}{d\chi_s} \end{bmatrix} + \begin{bmatrix} \frac{\partial a_1^i}{\partial \chi_s} \\ \vdots \\ \frac{\partial a_s^i}{\partial \chi_s} \\ \vdots \\ \frac{\partial a_J^i}{\partial \chi_s} \end{bmatrix}. \quad (3)$$

This system formalizes the notion of cross-asset demand complementarities: while asset s is the only security that is directly affected by the shock, the observed response to the shock is the sum of endogenous quantity adjustments for *all* assets in the choice set. These adjustments reflect two channels: (i) the degree of substitutability among assets (shown in blue), which is endogenously determined from the marginal valuation of each asset given the vector of asset positions, and (ii) endogenous price spillovers to other assets (shown in red). The last term on the right-hand side captures income effects from the revaluation of endowments. Going forward, we formally denote price spillovers by

$$\mathcal{S}_{js} \equiv \frac{dp_j}{d\chi_s}.$$

The key empirical challenge is that the data does not directly reveal the structural elasticity of interest. Instead, even with exogenous price variation in hand, observational data on portfolio holdings show only the *total* derivative with respect to all margins of adjustment. We call this object the *observed elasticity* $\hat{\mathcal{E}}_{ss}^i$ and note that it is equal to

$$\hat{\mathcal{E}}_{ss}^i \equiv -\frac{\frac{da_s^i}{d\chi_s}}{\frac{dp_s}{d\chi_s}} \cdot \frac{p_s}{a_s^i}.$$

In the next result, we decompose the observed elasticity into its structural components: the sum over all assets of structural cross-elasticities multiplied by price spillovers,

and the set of income effects that occur due to price changes. We focus mainly on the first component since it poses a bigger identification challenge and can lead to large biases.⁶

Proposition 1 (Decomposition of the observed elasticity) *The difference between the structural and observed own-price elasticities can be decomposed as follows:*

$$\hat{\mathcal{E}}_{ss}^i = \mathcal{E}_{ss}^i - \underbrace{\sum_{j \neq s} \frac{\mathcal{S}_{js} \frac{1}{p_j}}{\mathcal{S}_{ss} \frac{1}{p_s}} \mathcal{E}_{sj}^i}_{\text{Complementarities and spillovers}} - \underbrace{\frac{\frac{\partial a_s^i}{\partial \chi_s} \frac{1}{a_s^i}}{\mathcal{S}_{ss} \frac{1}{p_s}}}_{\text{Income effects}}. \quad (4)$$

Proof. See Appendix A. ■

To understand this result, consider again our thought experiment from the introduction, in which an investor can allocate funds between a bond and two stocks, say Microsoft and Apple. If the two stocks are highly substitutable, the cross-elasticity \mathcal{E}_{MA}^i is high and investors would rapidly switch to Apple in response to a price increase for Microsoft. In equilibrium, a supply shock to Microsoft must therefore trigger a price spillover to Apple. This price spillover makes it less appealing to substitute, driving a wedge between the observed elasticity (which incorporates all adjustments) and the structural elasticity (which presumes that all other prices remain fixed). Using the observed elasticity as a measure of the structural elasticity can therefore lead to large biases.

Under heterogeneous tastes, moreover, the nature of cross-asset complementarities is investor-specific. Hence asset demand systems typically estimate substitution patterns at the level of the investor. However, there are also cross-investor interactions because price spillovers depend on the *marginal* investor's willingness to substitute across assets.

In principle, one can overcome the identification challenge by using multiple observed data moments. In particular, exogenous variation in prices allows researchers to measure observed elasticities and cross-elasticities for multiple assets. These observed elasticities can then be used to construct estimators of the true elasticity. The approach in [Kojen and Yogo \(2019\)](#) is to define a logit demand system in which demand is measured *relative* to an “outside asset.” Under the assumption of homogeneous substitution between all assets (the standard *Independence of Irrelevant Alternatives* (IIA) property of logit demand), observed elasticities of relative demand identify the associated structural

⁶Income effects can be eliminated by redefining the units of demand. Under iso-elastic utility, for example, there are no income effects if one defines demand in terms of portfolio shares. See Section 3.2.

elasticity. More recently, [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#) show that difference-in-difference estimators can be used to difference out cross-asset spillovers if the portfolio choice problem admits a symmetric substitution matrix.

The limitation of these approaches is that they are valid only under strong symmetry assumptions that are unlikely to hold in standard portfolio choice settings, at least without further conditioning information. In particular, in standard settings, substitution patterns between inside assets (i) are generically heterogeneous across different assets, both with respect to each other and the outside asset, and (ii) depend on interactions with the rest of the investors' portfolio. As such, the substitution matrix is endogenously determined alongside the asset portfolio itself. As we show next, violations of the assumption of symmetric substitution can lead to large biases in the estimated elasticities.

3.2 Biased Elasticities in Logit Demand Systems for Financial Assets

We have shown that asset demand systems can accurately recover structural elasticities only if they appropriately account for demand complementarities and price spillovers. We now show that approaches based on the logit framework from [Kojien and Yogo \(2019\)](#) do not satisfy this requirement for canonical portfolio choice models. As such, estimated elasticities are subject to large biases. To make this argument, we must compare estimated and structural elasticities. Our approach is to use our model to generate “data” and then ask whether current methodologies accurately identify structural model parameters.

Asset menu. Solving for demand functions in arbitrary asset menus is complicated and not necessary to make our points. Hence we consider a relatively sparse setting in which there are only three assets: two *inside assets* which are the focus of the demand estimation, and an *outside asset* towards which investors can substitute in response to shocks. This setting is rich enough to capture many cross-asset substitution patterns of practical interest, but sufficiently simple to derive closed-form expressions for many key objects of interest. Furthermore, under specific parameter restrictions, logit demand can accurately capture true model-derived demand functions.

Definition 2 (Three-asset economy) *There are two aggregate states, $z \in \{1, 2\}$, and a distributional shock $\iota \in \{r, g\}$ which further affects asset payoffs. The probability of state z is π_z , and the probability of distributional shock ι satisfies $\Pr(\iota = r) = \rho$. There are three assets:*

- (i) Tree 2, which pays $y(2)$ if and only if state 2 is realized. We will refer to this asset as the outside asset, and normalize its price to $p_2 = 1$. The aggregate endowment of this asset is 1.
- (ii) A green tree with price p_g which pays only in aggregate state 1, and pays more when $\iota = g$. The aggregate endowment of this asset is $\frac{1}{2} + \psi$, where ψ is an exogenous supply shifter.
- (iii) A red tree with price p_r which pays only in aggregate state 1, and pays more when $\iota = r$. The aggregate endowment of this asset is $\frac{1}{2}$.

The specific state-contingent payoffs of all three assets are summarized in Table 1.

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

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

We also make the following assumptions on investor endowments and constraints:

- (i) Investors have the same initial endowments: $e_j^i = E_j$ and $w_0^i = 0 = w_1^i(z)$ for all i, j and z .
- (ii) Investors care only about consumption at date 1: the discount factor is $\delta = 1$.

Variation in the parameter ϵ allows us to capture a number of different scenarios. If $\epsilon = 0$, then green and red trees are perfect substitutes with respect to their cash flows. Hence these parameter values capture investors who face a security menu with similar assets, such as when they choose among similar stocks or derivative assets as well as stock. The assets become more complementary as ϵ increases. Hence intermediate values of ϵ capture when assets are complementary because they allow the investor to diversify distributional risk, but investors are still willing to substitute between assets to some degree because distributional risk is not too large. Finally, the limit $\epsilon \rightarrow y(1)$ leads to three distinct states of the world, each associated with a single tree that cannot be substituted for each other. This maps into scenarios in which there are no diversification benefits between green and red trees at all, such as when we consider an investor choosing between well-diversified portfolios each exposed to certain aggregate risk factors. Lastly, substitutability is also modulated by latent taste parameters θ^i . Formally, these are unobserved demand shifters that cannot be controlled for using data on asset cash flows alone.

Logit specification. The logit demand system in [Koijen and Yogo \(2019\)](#) describes asset demand in terms of portfolio shares $\omega_j^i(p) \equiv \frac{p_j a_j^i}{W^i}$, where $W^i \equiv \sum_{j \in \mathcal{J}} p_j a_j^i$ is the market value of investor i 's portfolio. The *relative portfolio share* of asset j is the portfolio share of asset j divided by the portfolio share of the outside asset, $\omega_j^i(p) / \omega_2^i(p)$. Demand is specified in this manner because the logit demand system presumes that cross-asset spillovers can be appropriately controlled for measuring demand relative to the outside good.⁷ The underlying logic stems from the Independence of Irrelevant Alternatives (IIA) property of logit demand, whereby substitution patterns are assumed to be homogeneous across assets. However, this assumption is at odds with the heterogeneous substitution patterns that naturally arise in standard portfolio choice settings with demand complementarities.

Going forward, we adapt our definitions of elasticities to units of relative portfolio shares as well. In particular, structural and observed elasticities are now given by

$$\mathcal{E}_{jj}^i \equiv -\frac{\partial(\omega_j^i(p) / \omega_2^i(p))}{\partial p_j} \frac{p_j}{\omega_j^i(p) / \omega_2^i(p)} \quad \text{and} \quad \hat{\mathcal{E}}_{jj}^i \equiv -\frac{d(\omega_j^i(p) / \omega_2^i(p))}{dp_j} \frac{p_j}{\omega_j^i(p) / \omega_2^i(p)}.$$

Cross-elasticities of relative portfolio shares are defined in the analogous way.

Identification of logit demand. We first review the precise specification and identification of logit demand systems for financial assets. In this approach, relative portfolio shares are specified to be log-linear in log prices and a set of factor loadings $(x_k(j))_k$ on asset characteristics.⁸ Characteristics are used to summarize key properties of expected returns and the variance-covariance matrix of returns. Specifically, in [Koijen and Yogo \(2019\)](#), investor-level relative portfolio shares are assumed to satisfy the demand function

$$\frac{\omega_j(p)}{\omega_2(p)} = \frac{\omega_j}{\omega_2}(p_j) = \exp \left\{ \beta_0 \log p_j + \sum_{k=1}^{K-1} \beta_k x_k(j) + \beta_K \right\} \zeta(j), \quad (5)$$

where β_0 is the coefficient on the log price of asset j , $x_k(j)$ is asset j 's loading on the k -th characteristics-based factor, $(\beta_k)_{k=1}^{K-1}$ are the associated demand coefficients, and β_K and

⁷Given iso-elastic utility, working with portfolio shares has the additional benefit that supply shocks do not create direct income effects. [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#) provide a more general taxonomy of the "natural units" of analysis for different utility functions.

⁸Technically, β_0 is the coefficient of log market equity, where market equity is the product of price and number of shares. This ensures neutrality to variation in the level of prices that is not economically meaningful, such as different choices for the initial number of shares in a firm. However, if the number of shares is constant, this term is another additive constant. Hence we focus simply on log prices.

$\zeta(j)$ are demand shifters for all inside assets and asset j , respectively. The latent asset-specific demand parameter $\zeta(j)$ corresponds to tastes θ_j in the context of our model.

A useful geometric interpretation of (5) is that β_0 measures the slope with respect to the own price, and all other demand shifters jointly determine the intercept. The key identifying assumption is that all demand shifters other than p_j are invariant to exogenous variation in p_j . Under this assumption, one can identify slope coefficient β_0 from the observed relative elasticity. In particular, totally differentiating (5) with respect to price p_j yields $\beta_0 = \hat{\mathcal{E}}_{jj}^i$, and this parameter recovers the structural elasticity as well, $\beta_0 = \mathcal{E}_{jj}^i$. This logic is shown in the left panel of Figure 1, where the observed change in relative portfolio weights of a given asset is treated as a move *along* a fixed demand curve. Appendix C characterizes how β_0 determines the *absolute* elasticity of demand for asset j , not just the relative demand elasticity with respect to the outside good. In the context of our model, the absolute elasticity is also precisely equal to β_0 .

Biased Measurement. Of course, logit demand systems fail to recover the structural elasticity if there are cross-asset spillovers and demand complementarities. This logic is depicted in the right panel of Figure 1, which incorporates the fact that a shock to one asset affects demand for other assets via complementarities and spillovers, triggering a shift in the demand curve. This effect is not captured in the logit demand system because factor loadings and characteristics, which proxy for expected returns and covariances, are assumed to be invariant to instrumented price shocks to a given asset. Given this restriction, observed elasticities are interpreted as determining the slope of demand.⁹

When complementarities exist, the observed elasticity is a biased estimator of the structural elasticity. The next result shows that the bias is directly proportional to complementarities and spillovers. Our formal definition of demand complementarities is that relative demand for asset j is affected by the prices of other assets, i.e., $\frac{\partial(\omega_j^i(p)/\omega_2^i(p))}{\partial p_{-j}} \neq 0$ in our three-asset economy, where p_{-j} is the price of the other inside asset.

Proposition 2 (Biased measurement of structural elasticities) *Consider the three-asset econ-*

⁹This restriction is derived under the assumption that *returns* are well-described by factor loadings. However, returns are endogenous objects that should respond to price changes.

omy. For investor i , the bias $\mathcal{B}_{jj}^i \equiv \mathcal{E}_{jj}^i - \hat{\mathcal{E}}_{jj}^i$ between structural and observed elasticities is

$$\mathcal{B}_{jj}^i = - \underbrace{\frac{\partial \left(\omega_j^i(p) / \omega_2^i(p) \right)}{\partial p_{-j}}}_{\text{Complementarity}} \underbrace{\left(\omega_j^i(p) / \omega_2^i(p) \right)}_{\text{Scaling terms}} \frac{p_{-j}}{p_j} \underbrace{\frac{dp_{-j}}{dp_j}}_{\text{Price Spillover}}. \quad (6)$$

Proof. See Appendix A. ■

The scaling terms ensure that the bias is reported in units of elasticities. We recover the logit identification result if and only if there are no complementarities or spillovers. Conversely, the bias from using logit demand is large when inside assets are very close substitutes (and thus much better substitutes for each other than for the outside asset). Indeed, we show formally that the bias diverges to infinity if inside assets are perfect substitutes, but zero if all assets are equally complementary.

To derive these results, we use our models to characterize structural demand elasticities under the optimal demand functions for our setting. To make these demand functions easier to interpret, we assume that there is no aggregate risk, $y(z) = 1$, although this is not necessary for any of the results to come. We have the following characterization.

Lemma 1 (Optimal portfolio shares) *Let $y(1) = y(2) = 1$. Given relative prices (p_g, p_r) and taste parameters (θ_g^i, θ_r^i) , the optimal relative portfolio shares for investor i are given by*

$$\frac{\omega_g^i(p_g, p_r)}{\omega_2^i(p_g, p_r)} = \theta_r^i \frac{\pi_1}{\pi_2} p_g \cdot \frac{(\theta_r^i p_g + \theta_g^i p_r) \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r) + 2\theta_g^i p_r \epsilon (1 - 2\rho)}{(\theta_r^i p_g + \theta_g^i p_r)^2 \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r)^2}; \quad (7)$$

$$\frac{\omega_r^i(p_g, p_r)}{\omega_2^i(p_g, p_r)} = \theta_g^i \frac{\pi_1}{\pi_2} p_r \cdot \frac{(\theta_r^i p_g + \theta_g^i p_r) \epsilon^2 + (\theta_r^i p_g - \theta_g^i p_r) - 2\theta_r^i p_g \epsilon (1 - 2\rho)}{(\theta_r^i p_g + \theta_g^i p_r)^2 \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r)^2}. \quad (8)$$

Proof. See Appendix A. ■

The critical feature of these demand functions is that they exhibit cross-asset complementarities, whereby the demand for green and/or red assets depends on the prices and tastes of both inside assets. As such, supply shocks to one asset alter relative portfolio shares of both assets, with the degree of complementarity influenced by parameter ϵ .

The next result formally characterizes the estimation bias that obtains when using the observed elasticity as an estimator for the structural elasticity. To derive this result, we

assume that we have access to a perfect asset-level instrument, namely a purely exogenous shock ψ to the supply of the green asset. While not necessary for the results, to obtain easily interpretable conditions for the bias we assume that all investors have the same tastes, $\theta_j^i = 1$, and that both aggregate states are symmetric, $\pi_1 = \frac{1}{2}$ and $y(1) = y(2) = 1$.

Proposition 3 (Measured vs Structural Elasticities) *Let $y(z) = 1 = \theta_j^i$ and $\pi_z = \frac{1}{2}$. Given an exogenous supply shock ψ around $\psi = 0$, for any i , observed and structural elasticities are:*

$$\hat{\mathcal{E}}_{gg}^i = \frac{(1 - \epsilon^2)}{(1 + \epsilon)^2 - 4\epsilon\rho} \quad \text{and} \quad \mathcal{E}_{gg}^i = \frac{(1 - \epsilon^2)(1 - \epsilon(1 - 2\rho))}{8\rho(1 - \rho)\epsilon^2}. \quad (9)$$

In the limit as green and red assets become perfect substitutes, we have:

$$\lim_{\epsilon \rightarrow 0} \hat{\mathcal{E}}_{gg}^i = 1 \quad \text{and} \quad \lim_{\epsilon \rightarrow 0} \mathcal{E}_{gg}^i = \infty. \quad (10)$$

For any $\epsilon > 0$, the bias, i.e., the difference between structural and observed elasticities, is

$$\mathcal{B}_{gg}^i = \frac{(1 - \epsilon^2)^2(1 + (1 - 2\rho)\epsilon)}{8\epsilon^2\rho(1 - \rho)((1 + \epsilon)^2 - 4\epsilon\rho)}. \quad (11)$$

The bias is positive, goes to infinity as $\epsilon \rightarrow 0$, is strictly decreasing in ϵ , and is zero iff $\epsilon = 1$.

Proof. See Appendix A. ■

Figure 3 illustrates the result. Measured elasticities are small throughout and of the same order of magnitude as leading estimates in the demand-system literature. In contrast, the structural elasticity approaches infinity as ϵ goes to zero (such as when inside assets are closely substitutable), but is zero when $\epsilon = 1$. This means that the measured elasticity is not closely related to the underlying structural elasticity.

To understand the intuition for this disconnect, consider the case where the two inside assets are close substitutes, $\epsilon \approx 0$. (Note that, in the extreme case where $\epsilon = 0$, any change in the green asset's price leads to an immediate arbitrage opportunity relative to the red asset.) Precisely because investors would like to respond to the green supply shock by buying more of the red asset, the red price must adjust to clear the market. Given this price spillover, it becomes less attractive to substitute across assets, leading to low *measured* elasticities in response to the shock. However, the reason that prices adjust in this manner is precisely that each individual investor would have responded very elastically had prices not adjusted.

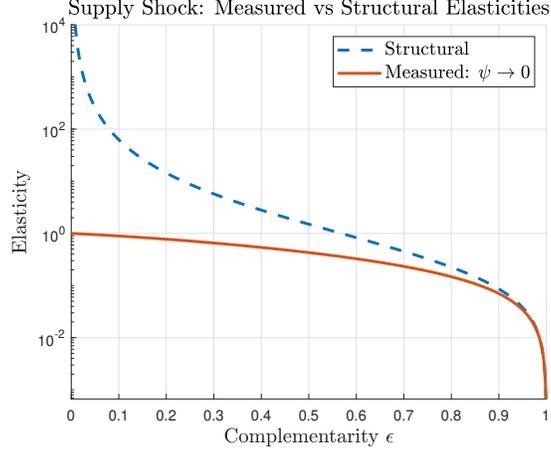


Figure 3: Measured vs structural elasticity given an exogenous supply shock ψ as $\psi \rightarrow 0$. The measured elasticity is $\hat{\mathcal{E}}_{gg}^i$. The structural elasticity is \mathcal{E}_{gg}^i . Parameters: $\pi_1 = \frac{1}{2}$, $y(1) = y(2) = 1$, $\theta_j^i = 1$, and $\rho = \frac{1}{4}$.

Our findings thus align directly with the basic mechanism of neoclassical finance: if financial markets rapidly reprice substitute assets in response to shocks, equilibrium portfolio responses may suggest very low elasticities even when the structural elasticity is high. The reason is that individual demand responses are strategic substitutes: when other investors adjust their portfolios, the resulting price response implies that any individual investor will rebalance her portfolio less than she otherwise would. Hence low observed elasticities are *not* dispositive evidence that financial markets are indeed slow to respond to profitable trading opportunities.

While the bias is strictly positive, it is decreasing in ϵ and converges to zero in the limit as $\epsilon \rightarrow 1$. In this limit, inside assets are not substitutable: the payoff structure features three states, each of which are associated with a single asset. Hence substitution between inside assets is symmetric to substitution with the outside asset, and demand complementarities and cross-asset spillovers are immaterial. As a result, logit demand accurately captures the underlying substitution patterns.

As such, our analysis also admits a more positive interpretation: while the bias is severe when close substitutes are available, it may be more muted when investors face choice sets with limited substitutability, as may be the case in international finance or foreign exchange (Jiang, Richmond, and Zhang, 2023; An and Huber, 2024; Koijen and Yogo, 2020).¹⁰ Nevertheless, substitutability is ultimately a latent variable that must be esti-

¹⁰Limited substitutability can also stem from tastes. Investors with a taste for one asset may not reallocate to another asset in response to a price shock to their preferred asset. See Section 4.2 for an example.

mated from data. Our results suggest that logit demand cannot accurately discriminate between varying degrees of substitutability: the measured elasticity ranges from zero to one while the structural elasticity ranges from zero to infinity.

More generally, an ideal estimation procedure would control for spillovers by estimating a fixed point between individual demands and the matrix of spillovers. While theoretically appropriate, the implementation of such a fixed-point approach may be challenging. Price spillovers depend on the market-wide degree of substitutability between assets, which in turn depends on the unobserved cross-sectional distribution of taste parameters. Hence researchers would have to jointly estimate the spillover matrix alongside individual demand functions, checking for consistency using market clearing. This means that demand systems may be simpler to estimate for more aggregated portfolios, such as stocks versus bonds, where cross-asset spillovers are likely to be relatively small. However, it may also be more difficult to find instruments for such settings.

Remark 1 (Aggregation from “Micro-logit” to Asset-level Elasticities) *We have derived estimation biases at the level of the individual investor (“micro logit”). One could also derive stock-level elasticities by averaging across investors. This does not solve the identification problem: since investor-level elasticities are always underestimated, so are asset-level elasticities. In the specific setting of Proposition 3, asset-level elasticities are identical to the asset-investor-level elasticity.*

More generally, one can consider variations on our economy with heterogeneous investors. For example, some investors may have a strong taste for green assets, while the remainder have no specific taste for either asset (as in the baseline). In this case, investors with a taste-based preference for a given asset will exhibit low (or even zero) structural elasticities, which the logit demand system will identify relatively accurately. However, it will fail to accurately identify the elasticity of the investors without specific tastes. Furthermore, precisely because taste-based investors are less willing to substitute, non-taste investors will be the marginal investors whose preferences determine the relative prices of green and red assets (and thus the asset-level price response to shocks). Hence estimated asset-level demand elasticities will again be severely biased.

3.3 Control Variables

One suggested approach to the problem of heterogeneous substitution and spillovers is the use of control variables. For instance, a supply shock to a given asset may trigger

spillovers primarily to other assets with similar factor exposures. In this case, controlling for common factor exposures mitigates the resulting asymmetry in substitution.

While useful for estimating certain parameters of interest, controls directly alter the object of analysis by changing the degree of substitutability between choices. In particular, two assets may be highly substitutable precisely *because* they have common exposures. In this case, controlling for common exposures yields demand elasticities defined over the *residual* cash flows unaccounted for by controls, rather than asset-level elasticities. Moreover, if residual cash flows are less substitutable than the asset, residual elasticities are lower and may carry little information about asset elasticities. Hence control-variables approaches are informative about asset demand only under additional assumptions.

We next illustrate this effect in the context of risk-based portfolio choice when asset payoffs obey a factor structure.¹¹ Substitutability is determined by covariances: assets are substitutable if cash flows covary positively, and complementary when they covary negatively. Hence conditional and unconditional substitutability differ whenever the conditional and unconditional covariances differ. If these differences are large, conditional demand functions are not informative about asset-level demand functions.

Example 1 (Controlling for exposures) *Let asset cash flows Y_j follow as single-factor model,*

$$Y_j = \beta_j F + \eta_j,$$

where $F \sim \mathcal{N}(\mu, \sigma^2)$ is the single factor, β_j asset j 's loading on the factor and η_j is mean-zero noise that is uncorrelated across assets. Then the covariance of cash flows for two assets a and b is

$$\text{Cov}(Y_a, Y_b) = \beta_a \beta_b \sigma^2, \tag{12}$$

while the covariance conditional on the factor is zero,

$$\text{Cov}(Y_a, Y_b \mid F) = 0.$$

Hence the residual cash flows are more complementary than the asset as a whole if

$$\text{Cov}(Y_a, Y_b \mid F) < \text{Cov}(Y_a, Y_b), \quad \text{that is, } \beta_a \beta_b > 0.$$

¹¹Factors are common control variables because they allow researchers to hold fixed certain risk exposures. The underlying logic is based on no arbitrage pricing of risk exposures. However, allowing for heterogeneous tastes may invalidate no arbitrage: see Appendix B.

Under relatively weak assumptions, knowledge of (taste-adjusted) factor loadings may be sufficient to determine the *sign* of the difference between residual and asset-level demand elasticities. For example the residual elasticity is likely to be lower than the asset-level elasticity if $\beta_a\beta_b > 0$ and higher if $\beta_a\beta_b < 0$. However, the precise *quantitative* mapping between the two depends on a number of (unobserved) variables, including the perceived contribution to total variance of the controls, the investor’s demand elasticities over the subset of asset cash flows correlated with the controls, and the interaction of each asset with the rest of the investor’s portfolio. Moving from residual elasticities to asset-level elasticities therefore requires additional, potentially strong, assumptions.

In Appendix D, we formally characterize the informativeness of residual and factor-level demand elasticities in the context of our model. We decompose payoffs into factors and introduce small idiosyncratic noise. This allows us to explicitly model demand functions for factors and for residual cash flows. Under standard assumptions, factor demand elasticities are zero and residual elasticities are equal to one *independently* of the underlying asset-level elasticity. Since the asset-level elasticity ranges from zero to infinity, factor and residual demand elasticities carry little information about asset-level demand.

These findings relate to [Haddad, He, Huebner, Kondor, and Loualiche \(2025\)](#), who argue that conditioning information can be used to obtain settings with specific symmetry properties. Specifically, if the asset menu does not satisfy symmetry unconditionally, judicious controls may yield a decision problem that is symmetric conditional on controls. Given such symmetry, they show that a difference-in-difference estimator identifies the *relative elasticity* between two assets subject to symmetric spillovers (namely, the change in the demand difference between the two assets given a change in the price difference). However, this procedure does not recover the absolute elasticity, and it works only if substitution patterns are symmetric across *all* assets within the asset menu, as well as with outside assets. In practice, this assumption is unlikely to hold without conditioning information.¹² Using controls leads to an estimate of the *residual* relative elasticity.

¹²While their symmetry assumption is satisfied in our example economy, small perturbations of the payoff structure which break symmetry can lead to very large biases. These results are available upon request.

4 Are Elasticities Structural Parameters?

The previous section discussed the measurement and interpretation of demand elasticities, but stopped of establishing whether asset demand elasticities are structural parameters that are invariant to counterfactuals, even if they are well-measured. This is a critical concern if asset demand systems are to be used to inform policy.

We discuss two main challenges. The first is that investors care about the resale value of their assets. Hence demand elasticities reflect not only the investor’s individual tastes for an asset, but also her expectations about *other* investors’ future valuations. Second, it is difficult to separately identify unobserved tastes and constraints. The standard logit asset demand system cannot separately identify these different sources of demand, since all unobserved variation is summarized using a single “latent demand” parameter. However, many counterfactuals are sensitive to the precise microfoundation.

4.1 Dynamic trading: whose preferences are being measured?

We begin by discussing the role of resale considerations. Forward-looking demand makes it difficult to separately identify individual and “market-wide” tastes because investors may buy an asset they personally dislike if they expect other investors will pay a high price for it (Keynes, 1936; Harrison and Kreps, 1978). However, counterfactuals involving shifts in the wealth distribution require knowledge of individual preferences.

In Appendix E we formally construct a simple two-period variant of our baseline framework for the special case where green and red trees have identical cash flows ($\epsilon = 0$), but variation in tastes creates variation in prices. Here, we only report the key equation determining asset demand in the first period.

Let $p_{j,1}$ denote the price of asset j at date 1, R_j the gross return of asset j between dates 1 and 2, and W_2^i the investor’s wealth at date 2. Given discount factor δ , an interior choice of $a_{j,1}^i$, the investor i ’s holdings of asset j at date 1, must satisfy

$$\frac{\pi_1 y(1) \theta_{j,1}^i}{p_{j,1}} u'(\tilde{c}_1^i(1)) = \pi_2 y(2) u'(\tilde{c}_1^i(2)) + \frac{\delta}{1-\delta} \mathbb{E}^i \left[\frac{R_2 - R_j}{W_2^i} \right]. \quad (13)$$

The left-hand side is the date-1 marginal benefit of buying asset j at current price $p_{j,1}$ in state 1. The parameter $\theta_{j,1}^i$ is investor i ’s taste for asset j at date 1. The right-hand

side consists of two components: the marginal loss from consuming less in state 2 at date 1, and the expected return reduction from carrying wealth forward in the form of tree j rather than tree 2. Hence demand is increasing in both private tastes and expected market returns. Since market returns are determined by the tastes of tomorrow’s marginal investor, observed demand elasticities do not reveal whether an investor is buying based on her tastes or her beliefs over market-wide tastes.

4.2 Tastes versus constraints

Next we consider the interpretation of observed elasticities when investors differ in both tastes and unobservable investment mandates. To this end, we study our baseline economy but assume that there are two types of investors: some prefer green to red, and the rest prefer red to green. “Green investor” owns a share γ of the aggregate endowment of every tree. This allows us to model counterfactual wealth distributions.

To ensure equilibrium existence, we assume that each investor faces short-sale constraints. Moreover, we assume that a share m of each type faces a strict mandate to only invest in their preferred trees (i.e., a green investor with a mandate cannot buy red trees). This mandate is unobserved to the econometrician. All derivations are in Appendix F.

Since investors have different tastes, the equilibrium may feature sorting. In particular, if diversification motives are sufficiently small ($\epsilon \approx 0$) and green investors are not too wealthy ($\gamma \leq \bar{\gamma}$ for some $\bar{\gamma}$), all green investors buy only green assets. In a sorting equilibrium, green investors without a mandate are observationally equivalent to green investors with a mandate. However, sorting breaks down when assets are complementary ($\epsilon \gg 0$) or when green investors are wealthy ($\gamma > \bar{\gamma}$), pushing up the green price.

Figure 4 shows the equilibrium green price p_g as a function of the wealth share γ and complementarity ϵ for two economies: one in which very few investors face a mandate (left panel), and the other with many constrained investors (right panel).

The two economies are observationally equivalent near the origin where unconstrained investors choose to specialize in their preferred color to the same extent as mandate investors. However, they differ sharply under counterfactual wealth distributions or payoff structures. For unconstrained investors, a shock to ϵ creates more demand for diversification. In the left panel, the price of green trees is thus decreasing in ϵ when green investors choose to hold inside assets. In contrast, mandate investors do not buy red trees

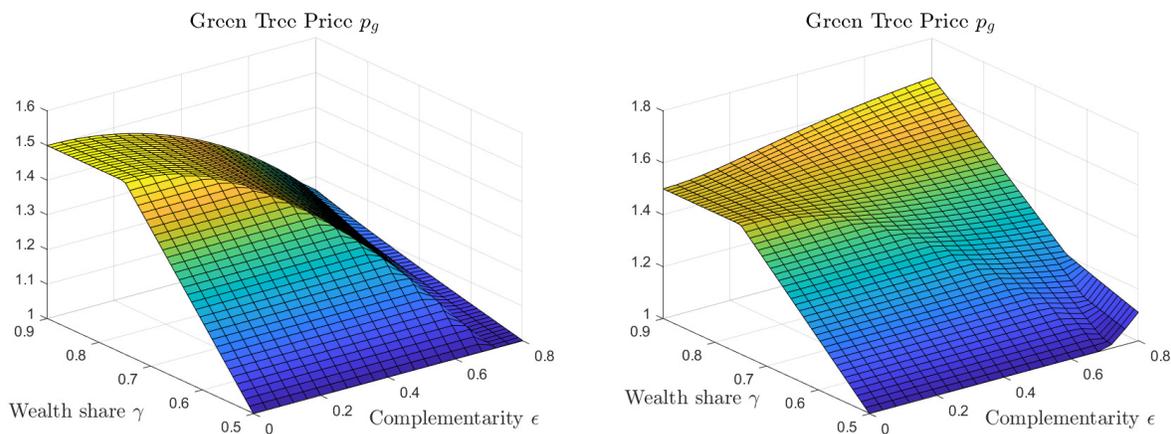


Figure 4: Green Price. Left: Mandate Share $m = 0.01$. Right: Mandate Share $m = 0.85$.

at any price. In the right panel, the price of green trees thus *increases* in ϵ . Hence, the two economies are observationally equivalent for some parameters, but qualitatively different under counterfactuals. Hence researchers should be cautious when assessing counterfactuals using demand systems that cannot separately identify tastes and constraints.

5 Conclusion

We present an analysis of demand systems for financial assets. Our results highlight the critical role of heterogeneous demand complementarities and equilibrium price spillovers. Specifically, we show that a failure to properly account for these effects can lead to severe biases in measured elasticities. This offers a simple reconciliation of the striking difference between demand elasticities in demand-system approaches and canonical benchmarks.

We see two main paths for future work. The first is to develop new structural frameworks that can account for richer cross-asset interactions. This could involve finding applications with additional data or structure, as in [Allen, Kastl, and Wittwer \(2025\)](#). The second is to develop empirical tests to bound the estimation biases that we discuss.

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 4.]
- ALLEN, J., J. KASTL, AND M. WITTEW (2025): “Estimating Demand Systems with Bidding Data,” *Working Paper*. [Cited on pages 5 and 27.]
- AN, Y., AND A. W. HUBER (2024): “Demand Propagation Through Traded Risk Factors,” Discussion paper. [Cited on pages 6 and 21.]
- BEN-DAVID, I., J. LI, A. ROSSI, AND Y. SONG (2022): “What Do Mutual Fund Investors Really Care About?,” *Review of Financial Studies*, 35(4), 1723–1774. [Cited on page 5.]
- BERRY, S. T., AND P. A. HAILE (2021): “Foundations of Demand Estimation,” *Handbook of Industrial Organization*, 4(1), 1–62. [Cited on page 5.]
- BRECKENFELDER, J., AND V. DE FALCO (2023): “Investor Heterogeneity and Large-Scale Asset Purchases,” *Working paper, ECB and Imperial College*. [Cited on page 5.]
- BRETSCHER, L., L. SCHMID, I. SEN, AND V. SHARMA (2022): “Institutional Corporate Bond Pricing,” *Working Paper*. [Cited on page 5.]
- CHANG, Y.-C., H. HONG, AND I. LISKOVICH (2015): “Regression Discontinuity and the Price Effects of Stock Market Indexing,” *Review of Financial Studies*, 28(1), 212–246. [Cited on page 5.]
- CHAUDARY, M., Z. FU, AND J. LI (2023): “Corporate Bond Multipliers: Substitutes Matter,” *Columbia University Working Paper*. [Cited on page 6.]
- COVAL, J., AND E. STAFFORD (2007): “Asset Fire Sales (and Purchases) in Equity Markets,” *Journal of Financial Economics*, 86(2), 479–512. [Cited on page 5.]
- D’AMICO, S., J. KLAUSMANN, AND A. A. PANCOST (2023): “The Benchmark Greenium,” *Working Paper*. [Cited on page 6.]
- DARMOUNI, O., K. Y. SIANI, AND K. XIAO (2023): “Nonbank Fragility in Credit Markets: Evidence from a Two-Layer Asset Demand System,” *Columbia University and MIT Working paper*. [Cited on page 5.]

- DAVIS, C., M. KARGAR, AND J. LI (2025): “Why Do Portfolio Choice Models Predict Inelastic Demand?,” *Available at SSRN: <https://ssrn.com/abstract=4195089> or <http://dx.doi.org/10.2139/ssrn.4195089>*. [Cited on page 5.]
- FUCHS, W., S. FUKUDA, AND D. NEUHANN (2025): “Kantian versus Consequentialist Portfolio Choice,” *Working paper, University of Texas at Austin and Bocconi University*. [Cited on page 6.]
- 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 5.]
- GENTZKOW, M. (2007): “Valuing New Goods in a Model with Complementarity: Online Newspapers,” *American Economic Review*, 97(3), 713–744. [Cited on page 5.]
- GOMPERS, P. A., AND A. METRICK (2001): “Institutional Investors and Equity Prices,” *Quarterly Journal of Economics*, 116(1), 229–259. [Cited on page 5.]
- GREENWOOD, R. M., AND M. SAMMON (2025): “The Disappearing Index Effect,” *Journal of Finance*, 80(2), 657–698. [Cited on page 5.]
- HADDAD, V., Z. HE, P. HUEBNER, P. KONDOR, AND E. LOUALICHE (2025): “Causal Inference for Asset Pricing,” *Discussion paper*. [Cited on pages 6, 11, 15, 17, and 24.]
- HADDAD, V., P. HUEBNER, AND E. LOUALICHE (2025): “How Competitive is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing,” *American Economic Review*, 115(3), 975–1018. [Cited on page 5.]
- 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 4.]
- HARRISON, J. M., AND D. M. KREPS (1978): “Speculative Investor Behavior in a Stock Market with Heterogenous Expectations,” *Quarterly Journal of Economics*, 92(2), 323–336. [Cited on page 25.]
- HARTZMARK, S. M., AND D. H. SOLOMON (2024): “Marketwide Predictable Price Pressure,” *Working Paper, Boston College*. [Cited on page 5.]

- HE, Z., AND A. KRISHNAMURTHY (2013): "Intermediary Asset Pricing," *American Economic Review*, 103(2), 732–70. [Cited on page 4.]
- JIANG, Z., R. RICHMOND, AND T. ZHANG (2023): "A Portfolio Approach to Global Imbalances," *NYU Stern School of Business Working Paper*. [Cited on pages 5 and 21.]
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest, and Money*. Macmillan, London. [Cited on pages 4 and 25.]
- KOIJEN, R. S., AND M. YOGO (2020): "Exchange Rates and Asset Prices in a Global Demand System," *NBER Working Paper 27342*. [Cited on pages 6 and 21.]
- KOIJEN, R. S. J., R. J. RICHMOND, AND M. YOGO (2024): "Which Investors Matter for Equity Valuations and Expected Returns?," *Review of Economic Studies*, 91(4), 2387–2424. [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 1, 2, 3, 10, 13, 14, 15, 17, 36, and 41.]
- (2021): "Asset Demand Systems in Macro-Finance," *NBER Reporter*, 4, 21–24. [Cited on page 1.]
- KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity*, 42(2), 215–287. [Cited on page 5.]
- (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120(2), 233–267. [Cited on page 4.]
- LI, Z. (2025): "Long Rates, Life Insurers, and Credit Spreads," *Working paper, Imperial College*. [Cited on page 5.]
- LOU, D. (2012): "A Flow-Based Explanation for Return Predictability," *Review of Financial Studies*, 25(12), 3457–3489. [Cited on page 5.]
- LUCAS, R. E. J. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–1445. [Cited on pages 1 and 7.]

- NENOVA, T. (2025): “Global or Regional Safe Assets: Evidence from Bond Substitution Patterns,” *London Business School Working Paper*. [Cited on pages 5 and 6.]
- PASTOR, L., R. F. STAMBAUGH, AND L. A. TAYLOR (2021): “Sustainable investing in equilibrium,” *Journal of Financial Economics*, 142(2), 550–571. [Cited on page 6.]
- PAVLOVA, A., AND T. SIKORSKAYA (2023): “Benchmarking Intensity,” *Review of Financial Studies*, 36(3), 859–903. [Cited on page 5.]
- ROSS, S. A. (2004): *Neoclassical Finance*. Princeton University Press, Princeton, NJ. [Cited on pages 4 and 35.]
- SELGRAD, J. (2023): “Testing the Portfolio Rebalancing Channel of Quantitative Easing,” Working paper. [Cited on page 5.]
- SHLEIFER, A. (1986): “Do Demand Curves for Stocks Slope Down?,” *Journal of Finance*, 41(3), 579–590. [Cited on page 4.]
- STARKS, L. T. (2023): “Presidential Address: Sustainable Finance and ESG Issues - Value versus Values,” *Journal of Finance*, 78(4), 1837–1872. [Cited on page 6.]
- TOBIN, J. (1969): “A General Equilibrium Approach to Monetary Theory,” *Journal of Money, Credit and Banking*, 1(1), 15–29. [Cited on page 4.]
- VAN DER BECK, P. (2022): “On the Estimation of Demand-Based Asset Pricing Models,” *Swiss Finance Institute Research Paper Series*, 22(67). [Cited on page 6.]

A Proofs

Proof of Proposition 1. First, it follows from (the s -th row of) equation (3) that

$$\frac{da_s^i}{d\chi_s} = \frac{\partial a_s^i}{\partial p_s} \frac{dp_s}{d\chi_s} + \sum_{j \neq s} \frac{\partial a_s^i}{\partial p_j} \frac{dp_j}{d\chi_s} + \frac{\partial a_s^i}{\partial \chi_s}.$$

Second, multiplying both sides by $-\frac{1}{\frac{dp_s}{d\chi_s} a_s^i} p_s$ and noting the definitions of $\hat{\mathcal{E}}_{ss}^i$, \mathcal{E}_{ss}^i , \mathcal{S}_{js} , and \mathcal{S}_{ss} , we have:

$$\hat{\mathcal{E}}_{ss}^i = \mathcal{E}_{ss}^i - \left(\sum_{j \neq s} \frac{\partial a_s^i}{\partial p_j} \frac{\mathcal{S}_{js} p_s}{\mathcal{S}_{ss} a_s^i} + \frac{\frac{\partial a_s^i}{\partial \chi_s} p_s}{\mathcal{S}_{ss} a_s^i} \right).$$

Third, then, rearranging, we obtain equation (4). ■

Proof of Proposition 2. Observe that the structural elasticity can be rewritten as:

$$\mathcal{E}_{jj}^i = -\frac{\partial \log \left(\frac{\omega_j^i(p)}{\omega_2^i(p)} \right)}{\partial p_j} p_j = - \left(\frac{\partial \omega_j^i(p)}{\partial p_j} \frac{p_j}{\omega_j^i} - \frac{\partial \omega_2^i(p)}{\partial p_j} \frac{p_j}{\omega_2^i} \right).$$

In contrast, the measured elasticity is written as:

$$\begin{aligned} \hat{\mathcal{E}}_{jj}^i &= -\frac{d \log \left(\frac{\omega_j^i(p)}{\omega_2^i(p)} \right)}{dp_j} p_j = - \left(\frac{d\omega_j^i(p)}{dp_j} \frac{p_j}{\omega_j^i} - \frac{d\omega_2^i(p)}{dp_j} \frac{p_j}{\omega_2^i} \right) \\ &= - \left(\frac{\partial \omega_j^i(p)}{\partial p_j} \frac{p_j}{\omega_j^i} - \frac{\partial \omega_2^i(p)}{\partial p_{-j}} \frac{p_j}{\omega_2^i} \right) - \left(\frac{\partial \omega_j^i(p)}{\partial p_{-j}} \frac{p_j}{\omega_j^i} - \frac{\partial \omega_2^i(p)}{\partial p_{-j}} \frac{p_j}{\omega_2^i} \right) \frac{dp_{-j}}{dp_j}. \end{aligned}$$

Hence,

$$\begin{aligned} \mathcal{B}_{jj}^i &= \mathcal{E}_{jj}^i - \hat{\mathcal{E}}_{jj}^i = - \left(\frac{\partial \omega_j^i(p)}{\partial p_{-j}} \frac{p_{-j}}{\omega_j^i} - \frac{\partial \omega_2^i(p)}{\partial p_{-j}} \frac{p_{-j}}{\omega_2^i} \right) \frac{p_j}{p_{-j}} \frac{dp_{-j}}{dp_j} \\ &= - \frac{\partial \left(\omega_j^i(p) / \omega_2^i(p) \right)}{\partial p_{-j}} \frac{p_{-j}}{\left(\omega_j^i(p) / \omega_2^i(p) \right)} \frac{p_j}{p_{-j}} \frac{dp_{-j}}{dp_j}, \end{aligned}$$

as desired. ■

Proof of Lemma 1. Observe that, under $y(1) = y(2) = 1$, the first-order conditions with

respect to a_g^i and a_r^i can be rewritten as

$$p_g = \frac{\pi_1(1-\rho)}{\pi_2} \frac{(1+\epsilon)\theta_g^i a_2^i}{(1+\epsilon)\theta_g^i a_g^i + (1-\epsilon)\theta_r^i a_r^i} + \frac{\pi_1\rho}{\pi_2} \frac{(1-\epsilon)\theta_g^i a_2^i}{(1-\epsilon)\theta_g^i a_g^i + (1+\epsilon)\theta_r^i a_r^i}; \quad (14)$$

$$p_r = \frac{\pi_1(1-\rho)}{\pi_2} \frac{(1-\epsilon)\theta_r^i a_2^i}{(1+\epsilon)\theta_g^i a_g^i + (1-\epsilon)\theta_r^i a_r^i} + \frac{\pi_1\rho}{\pi_2} \frac{(1+\epsilon)\theta_r^i a_2^i}{(1-\epsilon)\theta_g^i a_g^i + (1+\epsilon)\theta_r^i a_r^i}. \quad (15)$$

Substituting the budget constraint

$$a_2^i = p_g e_g^i + p_r e_r^i + e_2^i - p_g a_g^i - p_r a_r^i$$

and solving for (a_g^i, a_r^i) , we obtain:

$$a_g^i = \theta_r^i \pi_1 \cdot \frac{(p_g e_g^i + p_r e_r^i + e_2^i) \left((\theta_r^i p_g + \theta_g^i p_r) \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r) + 2\theta_g^i p_r \epsilon (1-2\rho) \right)}{(\theta_r^i p_g + \theta_g^i p_r)^2 \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r)^2}; \quad (16)$$

$$a_r^i = \theta_g^i \pi_1 \cdot \frac{(p_g e_g^i + p_r e_r^i + e_2^i) \left((\theta_r^i p_g + \theta_g^i p_r) \epsilon^2 + (\theta_r^i p_g - \theta_g^i p_r) - 2\theta_r^i p_g \epsilon (1-2\rho) \right)}{(\theta_r^i p_g + \theta_g^i p_r)^2 \epsilon^2 - (\theta_r^i p_g - \theta_g^i p_r)^2}. \quad (17)$$

We then obtain the portfolio weights (ω_g^i, ω_r^i) in equations (7) and (8), respectively. ■

Proof of Proposition 3. First, observe that the representative agent (thus we suppress the superscript i) consumes the aggregate endowment in equilibrium. Thus, under the assumptions that $y(1) = y(2) = 1$, $\pi_1 = \frac{1}{2}$, and no tastes, the first-order conditions yield:

$$p_g = (1-\rho) \frac{1+\epsilon}{1+(1+\epsilon)\psi} + \rho \frac{1-\epsilon}{1+(1-\epsilon)\psi};$$

$$p_r = (1-\rho) \frac{1-\epsilon}{1+(1+\epsilon)\psi} + \rho \frac{1+\epsilon}{1+(1-\epsilon)\psi}.$$

Then, the change in prices $\Delta p_g \equiv p_g - p_g^0$ satisfies:

$$\Delta p_g \equiv p_g - p_g^0 = -\frac{1+\epsilon^2 + (1-\epsilon^2)\psi + (1-2\rho)(2+(1-\epsilon^2)\psi)\epsilon}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)} \psi.$$

When $\epsilon = 0$, we have $p_g = \frac{1}{1+\psi}$, $p_g^0 = 1$, and $\Delta p_g = -\frac{\psi}{1+\psi}$.

Second, substituting the equilibrium prices into the portfolio weight function ω_g

yields the equilibrium portfolio weights of green tree:

$$\begin{aligned}\omega_g &= \frac{1+2\psi}{4} \cdot \frac{1+\psi+\epsilon(1-2\rho-\epsilon\psi)}{(1+(1+\epsilon)\psi)(1+(1-\epsilon)\psi)}; \\ \omega_g^0 &= \frac{1}{4}(1+(1-2\rho)\epsilon); \\ \Delta\omega_g \equiv \omega_g - \omega_g^0 &= \frac{1}{4} \cdot \frac{(1-\epsilon^2)\psi(1+\psi-\epsilon(1-2\rho)\psi)}{(1+(1-\epsilon)\psi)(1+(1+\epsilon)\psi)}.\end{aligned}$$

In the limit as $\epsilon = 0$, we have $\omega_g = \frac{1+2\psi}{4(1+\psi)}$, $\omega_g^0 = \frac{1}{4}$, and $\Delta\omega_g = \frac{\psi}{4(1+\psi)}$.

Third, the measured elasticity reduces to

$$-\frac{\Delta\omega_g}{\Delta p_g} \frac{p_g^0}{\omega_g^0} = \frac{(1-\epsilon^2)(1+\psi-\epsilon(1-2\rho)\psi)}{1+\epsilon^2+(1-\epsilon^2)\psi+(1-2\rho)(2+(1-\epsilon^2)\psi)\epsilon}. \quad (18)$$

In the limit as $\psi \rightarrow 0$, we obtain the first equation in (9). Also, in the limit as $\epsilon \rightarrow 0$, the measured elasticity tends to the first equation in (10).

Forth, to compute the structural elasticity, we can compute

$$-\frac{\partial\omega_g}{\partial p_g} \frac{p_g}{\omega_g}$$

as a function of (p_g, p_r) when $(e_g, e_r, e_2) = (\frac{1}{2} + \psi, \frac{1}{2}, 1)$. Note that, since ω_g does not depend on (e_g, e_r, e_2) , this term does not depend on (e_g, e_r, e_2) , i.e., ψ . Then, evaluating this at the initial equilibrium prices (p_g^0, p_r^0) , the structural elasticity is given by the second equation in (9), which goes to infinity as $\epsilon \rightarrow 0$, establishing the second equation in (10).

Fifth, as in Proposition 1, we can decompose the measured demand response (in the limit as $\psi \rightarrow 0$) as the true response and the bias term:

$$\frac{\frac{d\omega_g}{d\psi}}{\frac{d\log p_g}{d\psi}} = \frac{\frac{\partial\omega_g}{\partial \log p_g} \frac{d\log p_g}{d\psi} + \frac{\partial\omega_g}{\partial \log p_r} \frac{d\log p_r}{d\psi}}{\frac{d\log p_g}{d\psi}} = \frac{\partial\omega_g}{\partial \log p_g} + \frac{\partial\omega_g}{\partial p_r} \frac{\frac{dp_r}{d\psi}}{\frac{d\log p_g}{d\psi}}.$$

Then, we can decompose the measured elasticity (in the limit as $\psi \rightarrow 0$) into the structural elasticity and the bias term:

$$-\frac{\frac{d\omega_g}{d\psi}}{\frac{d\log p_g}{d\psi}} \frac{p_g}{\omega_g} = -\frac{\partial\omega_g}{\partial \log p_g} \frac{p_g}{\omega_g} - \frac{\partial\omega_g}{\partial p_r} \frac{\frac{dp_r}{d\psi}}{\frac{d\log p_g}{d\psi}} \frac{p_g}{\omega_g}.$$

Consequently, the measured elasticity in the limit as $\psi \rightarrow 0$ is:

$$-\frac{\frac{d\omega_g}{de_g} p_g}{\frac{dp_g}{de_g} w_g} = \frac{(1 - \epsilon^2)}{(1 + \epsilon)^2 - 4\epsilon\rho'}$$

which coincides with the limit of equation (18) as $\psi \rightarrow 0$. In the limit $\psi \rightarrow 0$, the bias is:

$$\frac{\partial\omega_g}{\partial p_r} \frac{\frac{dp_r}{d\psi}}{\frac{d\log p_g}{d\psi}} \frac{p_g}{\omega_g} = \frac{(1 - \epsilon^2)^2(1 + (1 - 2\rho)\epsilon)}{8\epsilon^2\rho(1 - \rho)((1 + \epsilon)^2 - 4\epsilon\rho)}. \quad (19)$$

It follows from equation (19) that the bias is positive (when $\epsilon < 1$) and goes to infinity as $\epsilon \rightarrow 0$. Also, the derivative of the bias term with respect to ϵ is

$$-\frac{(1 + \epsilon)^4 + 2\epsilon(1 - \epsilon)(1 + \epsilon)^2(7 + \epsilon)(1 + \epsilon^2)\rho + 16\epsilon^2(1 - \epsilon^4)\rho^2}{8\epsilon^3\rho(1 - \rho)((1 + \epsilon)^2 - 4\epsilon\rho)^2} < 0,$$

which establishes that the bias is strictly decreasing in ϵ . ■

B No Arbitrage with Tastes

In many applications, researchers aim to control for certain risk exposures or use a low-dimensional factor representation of the matrix of expected returns to model demand. No arbitrage is used to ensure that assets are priced according to their risk exposures. We now discuss the relation between taste heterogeneity and the principle of no arbitrage.

In the standard definition, investors care only about cash flows and an arbitrage is “an investment strategy that guarantees a positive payoff in some contingency with no possibility of a negative payoff and no initial net investment” (Ross, 2004). When investors differ in tastes, they have subjective views on the payoffs of a given trade.

To define no arbitrage with tastes, we provide two preliminary definitions. First, letting a vector subspace \mathcal{A} of \mathbf{R}^J denote the set of feasible portfolios and letting $(p_j)_{j \in \mathcal{J}}$ be a price vector of individual assets, the pricing function $P : \mathcal{A} \rightarrow \mathbf{R}$ maps a portfolio $a = (a_j)_{j \in \mathcal{J}}$ into its price according to $P(a) \equiv \sum_{j \in \mathcal{J}} p_j a_j$. Second, investor i has a linear taste function $v^i : \mathcal{A} \rightarrow \mathbf{R}^Z$ that maps a portfolio a into a vector $v^i(a)$ of state-contingent taste-augmented payoffs for investor i . Specifically, letting $Y^i \equiv (\theta_j^i y_j(z))_{z,j}$ be the $Z \times J$ matrix of investor i 's payoff-augmented cash flows, $v^i(a) \equiv Y^i a$. This generalizes the

neoclassical approach in which $\theta_j^i = 1$ for all j . We then have the following definition.

Definition 3 (No Arbitrage with Tastes) *Let taste functions v^i be given for all investors i . The pricing function P leaves no arbitrage opportunities if, for any investor i and any portfolio $a \in \mathcal{A}$ such that the effective payoff is weakly positive almost surely (i.e., $v^i(a) \geq 0$) and strictly positive with strictly positive probability (i.e., $v^i(a) > 0$), the associated price is positive: $P(a) > 0$.*

That is, pricing function P leaves no arbitrage opportunities given taste functions $(v^i)_i$ if and only if, for every i , the pricing function P leaves no arbitrage opportunities in the standard sense if the cash-flow matrix is replaced by the taste-augmented payoff matrix. The key difficulty is that this payoff matrix is investor-specific. The main restriction is that taste functions are linear, as they are in [Kojien and Yogo \(2019\)](#).

Theorem 1 (Generic Arbitrage Opportunities with Tastes) *Fix taste functions v^i for all i . There does not exist a pricing function P that leaves no arbitrage opportunities if:*

$$\text{there exist } a, i, \text{ and } i' \text{ such that } v^i(a) > 0 \text{ and } 0 \geq v^{i'}(a). \quad (\text{C})$$

A sufficient condition for (C) is that there exist assets j and j' such that

(i) both assets have identical cash flows:

$$y_j(z) = y_{j'}(z) \text{ for all } z \in \mathcal{Z};$$

(ii) there exist investors i and i' with sufficiently heterogeneous tastes with respect to these assets:

$$\theta_j^i > \theta_{j'}^i \text{ and } \theta_j^{i'} \leq \theta_{j'}^{i'}.$$

Proof. We first show that there does not exist a pricing function P that leaves no arbitrage opportunities if (C) holds. We then establish the stated sufficiency condition for (C).

Suppose first that (C) holds. Generically, we can assume that $v^i(a) > 0$ and $v^{i'}(a) < 0$. Now, suppose to the contrary that there exists a pricing function P that leaves no arbitrage opportunities. Applying Definition 3 to investor i yields $P(a) > 0$. But applying Definition 3 to investor i' yields $P(-a) > 0$, i.e., $P(a) < 0$. This is a contradiction.

Next, suppose that the stated conditions hold: there exist two assets j and j' and two investors i and i' satisfying the two conditions. Denoting by $v_j^i \equiv (\theta_j^i y_j(z))_z$ investor

i 's marginal taste with respect to asset j , the two conditions imply:

$$v_j^i > v_{j'}^i \text{ and } v_j^{i'} \leq v_{j'}^{i'}. \quad (\text{C}')$$

It is then sufficient to show that (C') implies condition (C). Let $a \in \mathbf{R}^J$ be such that

$$a_\ell = \begin{cases} 1 & \text{if } \ell = j \\ -1 & \text{if } \ell = j' \\ 0 & \text{otherwise} \end{cases}.$$

Then, we obtain $v^i(a) = v_j^i - v_{j'}^i > 0$ and $v^{i'}(a) = v_j^{i'} - v_{j'}^{i'} \leq 0$, as desired. ■

The following example provides a simple illustration consistent with our model.

Example 2 (Green and Red Assets) *There are a green asset and a red asset with prices denoted by p_g and p_r , respectively. Both assets deliver a unit payoff with certainty. There are two investor types, denoted by α and β , that differ in their relative taste for the two assets. For each investor type i , the taste function is given by $v^i(a_g, a_r) = \theta_g^i a_g + \theta_r^i a_r$ with the following properties: while type α 's taste-augmented payoffs for green and red assets satisfy $\theta_g^\alpha > \theta_r^\alpha$, type β has $\theta_g^\beta < \theta_r^\beta$. Then there are no prices such that both investors agree on the value of a long-short portfolio selling one unit of the green asset and buying one unit of the red asset.*

C Computing the absolute elasticity given β_0

In the baseline logit model, the key estimation equation is written in terms of portfolio shares relative to the outside good. We now show how β_0 can be used to back out the absolute elasticity of the portfolio share of asset j , $\partial \log(\omega_j) / \partial \log(p_j)$ given knowledge of the demand for the outside asset.

Consider $J - 1$ inside assets and one outside asset. Let o denote the outside asset. For any inside assets j and k with $j \neq k$, the logit demand system assumes that

$$0 = \frac{\partial}{\partial \log p_k} \log \left(\frac{\omega_j}{\omega_o} \right) (p) = \frac{\partial \log \omega_j(p)}{\partial \log p_k} - \frac{\partial \log \omega_o(p)}{\partial \log p_k}.$$

Thus, we obtain:

$$\frac{\partial \log \omega_j(p)}{\partial \log p_k} = \frac{\partial \log \omega_o(p)}{\partial \log p_k}.$$

This equation means that, for the j -th row of the elasticity matrix $\left(\frac{\partial \log \omega_j(p)}{\partial \log p_k}\right)_{j,k}$, the (j, k) -th element (with $k \neq j$) is

$$\frac{\partial \log \omega_o(p)}{\partial \log p_k},$$

which does not depend on j . For the (j, j) -element of the elasticity matrix, we have

$$\frac{\partial}{\partial \log p_j} \log \left(\frac{\omega_j}{\omega_o} \right) (p) = \frac{\partial \log \omega_j(p)}{\partial \log p_j} - \frac{\partial \log \omega_o(p)}{\partial \log p_j}.$$

Under the assumption of the logit demand system that

$$\frac{\partial}{\partial \log p_j} \log \left(\frac{\omega_j}{\omega_o} \right) (p) = \beta_0,$$

we have:

$$\frac{\partial \log \omega_j(p)}{\partial \log p_j} = \beta_0 + \frac{\partial \log \omega_o(p)}{\partial \log p_j}.$$

We then obtain the following elasticity matrix:

$$\begin{bmatrix} \beta_0 + \frac{\partial \log \omega_o(p)}{\partial \log p_1} & \frac{\partial \log \omega_o(p)}{\partial \log p_j} & \dots & \frac{\partial \log \omega_o(p)}{\partial \log p_{J-1}} \\ \frac{\partial \log \omega_o(p)}{\partial \log p_1} & \beta_0 + \frac{\partial \log \omega_o(p)}{\partial \log p_2} & \dots & \frac{\partial \log \omega_o(p)}{\partial \log p_{J-1}} \\ \frac{\partial \log \omega_o(p)}{\partial \log p_1} & \frac{\partial \log \omega_o(p)}{\partial \log p_2} & \dots & \frac{\partial \log \omega_o(p)}{\partial \log p_{J-1}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \log \omega_o(p)}{\partial \log p_1} & \frac{\partial \log \omega_o(p)}{\partial \log p_2} & \dots & \beta_0 + \frac{\partial \log \omega_o(p)}{\partial \log p_{J-1}} \end{bmatrix}.$$

Letting ω_j be such that $\frac{\partial \log \omega_o(p)}{\partial \log p_j} = -\omega_j \beta_0$, the elasticity matrix reduces to:

$$\begin{bmatrix} (1 - \omega_1)\beta_0 & -\omega_2\beta_0 & \dots & -\omega_{J-1}\beta_0 \\ -\omega_1\beta_0 & (1 - \omega_2)\beta_0 & \dots & -\omega_{J-1}\beta_0 \\ -\omega_1\beta_0 & -\omega_2\beta_0 & \dots & -\omega_{J-1}\beta_0 \\ \vdots & \vdots & \ddots & \vdots \\ -\omega_1\beta_0 & -\omega_2\beta_0 & \dots & (1 - \omega_{J-1})\beta_0 \end{bmatrix}.$$

Given information on demand elasticities for the outside asset, we can then compute the

elasticity in absolute terms. In the three-asset economy, the elasticity matrix is:

$$\begin{bmatrix} \beta_0 + \frac{\partial \log \omega_2(p)}{\partial \log p_g} & \frac{\partial \log \omega_2(p)}{\partial \log p_r} \\ \frac{\partial \log \omega_2(p)}{\partial \log p_g} & \beta_0 + \frac{\partial \log \omega_2(p)}{\partial \log p_r} \end{bmatrix} = \begin{bmatrix} \left(1 - \frac{\frac{\partial \log \omega_2(p)}{\partial \log p_g}}{\beta_0}\right) \beta_0 & -\frac{\frac{\partial \log \omega_2(p)}{\partial \log p_r}}{\beta_0} \\ -\frac{\frac{\partial \log \omega_2(p)}{\partial \log p_g}}{\beta_0} & \left(1 - \frac{\frac{\partial \log \omega_2(p)}{\partial \log p_r}}{\beta_0}\right) \beta_0 \end{bmatrix}.$$

In our model, we can also directly compute the optimal demand for the outside asset. Under log utility, it is trivial that $\omega_2 = \pi_2$. Hence the the elasticity matrix is

$$\begin{bmatrix} \beta_0 & 0 \\ 0 & \beta_0 \end{bmatrix}.$$

D Residual and Factor Demand Elasticities

In this appendix, we formally characterize the difference between asset-level elasticities and *residual* elasticities in the context of our model. We will show that neither factor-level demand functions nor residual demand functions accurately inform asset-level demand elasticities. To ensure that no arbitrage holds (as is required for control variables based on factor exposures), we assume that investors do not have heterogeneous tastes, $\theta_j^i = 1$.

We begin by decomposing the payoff structure from Table 1 into underlying factors. For pedagogical purposes, assume that the factors are Arrow securities on the states. Since there are two aggregate states and a distributional shock, the set of states is $\{g, r, 2\}$. By a slight abuse of notation, use z to index a generic state. Given the payoff structure from Table 1, asset j has factor loading $y_j(z)$ on the state- z Arrow security, and factor loadings perfectly capture asset payoffs. Since green and red trees can be combined to replicate Arrow securities, all factors are traded. By no arbitrage, the factors thus have well-defined prices $q(z)$ and they are related to asset prices by

$$p_j = \sum_z y_j(z) q(z).$$

Factor demands. We first show that *factor-level* demand elasticities are not useful for capturing asset-level elasticities. Since the factors are traded, we can define a decision problem directly over state-contingent consumption $c(z)$ with associated price $q(z)$. There are two notions of factor demand functions. In the first, we simultaneously choose demand

functions for all factors. This corresponds to the following decision problem:

$$\begin{aligned} \max_{c(r), c(g)} \quad & \pi_2 u(c(2)) + \pi_1 \rho u(c(r)) + \pi_1 (1 - \rho) u(c(g)) && \text{(Factor demand)} \\ \text{s.t.} \quad & c(2) = W - q(r)c(r) - q(g)c(g). \end{aligned}$$

Alternatively, we could derive the demand function for a single factor, holding other factor exposures constant. For example, consider the demand for the green factor holding fixed the exposure to the red asset. This corresponds to the conditional decision problem:

$$\begin{aligned} \max_{c(g)} \quad & \pi_2 u(c(2)) + \pi_1 \rho u(c(r)) + \pi_1 (1 - \rho) u(c(g)) && \text{(Cond'l factor demand)} \\ \text{s.t.} \quad & c(2) = W - q(r)c(r) - q(g)c(g). \end{aligned}$$

We next derive the solution to these problems in terms of relative portfolio shares.

Proposition 4 (Zero factor elasticities) *Let $\omega^F(z) = q(z)c(z)/W$ denote the factor-level portfolio share. The conditional and unconditional relative portfolio shares satisfy*

$$\frac{\omega^F(z)}{\omega^F(2)} = \frac{\pi_z}{\pi_2} \quad \text{for } z \in \{r, g\}.$$

Hence factor demand elasticities are equal to zero for any factor prices.

Proof. The solution is standard. The first-order condition for Arrow security $z \in \{r, g\}$ is

$$\pi_2 q(z) u'(c(2)) = \pi_z u'(c(z)), \quad \text{that is,} \quad q(z) = \frac{\pi_z u'(c(z))}{\pi_2 u'(c(2))}.$$

Under log utility, this yields $q(z) = \frac{\pi_z c(2)}{\pi_2 c(z)}$. Imposing the budget constraint yields

$$c(z)q(z) = \frac{\pi_z}{\pi_2} \pi_2 W = \pi_z W, \quad \text{that is,} \quad \frac{c(z)q(z)}{W} = \pi_z.$$

Since portfolio shares are invariant in prices, the elasticity is always zero. ■

The result shows that factor demand functions exhibit low elasticities *for any prices and any payoff structures*. In particular, factor elasticities are zero for any value of ϵ . Since asset-level elasticities are strictly decreasing in ϵ and diverge to infinity as $\epsilon \rightarrow 0$, knowledge of the factor elasticities is not informative about the asset-level elasticities.

It is difficult to infer asset-level elasticities from factor elasticities mainly because no arbitrage ensures that factor prices are identical no matter through which asset factor exposures are acquired. Hence there is no asset-level variation in factors that can be used to identify cross-asset substitution.

Residual Demands We next measure residual demands in the case where factors do not subsume all asset cash flows. Because factors fully capture cash flows in our baseline model, we add idiosyncratic noise to green and red assets. The augmented payoff \tilde{y}_j is

$$\tilde{y}_j = y_j + \eta_j,$$

where η_j is a random variable with mean μ_j , standard deviation σ_j that is uncorrelated across assets. Volatility σ_j is small in order to study a small perturbation of our model.

Next, suppose that the factors are directly traded, either because Arrow securities exist or because they are sufficiently many assets to form well-diversified factor portfolios. Let q_k denote the price of a unit of exposure to factor k and let each investor choose factor quantities α_k^i . An asset is a bundle of its factor exposures and the residual idiosyncratic component η_j . By no arbitrage, there exists a well-defined price for η_j , say \tilde{p}_j .

We can then model the portfolio choice in two steps. First, choose positions in the underlying assets. Second, adjust factor positions to achieve desired factor exposures. This leads to a natural *conditional* decision problem: controlling for factor exposures, choose positions \tilde{a}_j for the idiosyncratic components of each asset. The solution to this problem yields *residual demand functions* given fixed factor positions. To make this point as simply as possible, assume as in [Kojien and Yogo \(2019\)](#) that the investor can invest in some outside asset in elastic supply, and that \tilde{p}_j are again relative prices with respect to the outside good. Then the decision problem determining residual demand functions is

$$\max_{(\tilde{a}_j)} \mathbb{E}[\log(\bar{c} + \sum_j \eta_j \tilde{a}_j)],$$

where

$$c = \underbrace{\sum_k \alpha_k F_k}_{\equiv \bar{c}} + \sum_j \eta_j \tilde{a}_j$$

is the consumption process for the investor and \bar{c} is the (fixed) component of consumption

that is due to factor exposures. As is standard, this problem can also be stated as:

$$\max_{(\tilde{\omega}_j)} \mathbb{E}[\log(\bar{c} + W \sum_j \frac{\eta_j}{\tilde{p}_j} \tilde{\omega}_j)]$$

where $\tilde{\omega}_j = \tilde{p}_j \tilde{a}_j / W_j$ is the portfolio share of the residual component and we assume that changes in demand for the idiosyncratic demand is accommodated by a change in the demand for the outside asset. Since η_j is a payoff, η_j / \tilde{p}_j is a return.

To illustrate the determination of residual demand elasticities, fix $\bar{c} = 1$ (as is the case in our model if $y(z) = 1$ and the residual component is small). Then a standard approximation to this problem as in Campbell-Viceira yields the decision problem:

$$\max_{(\tilde{\omega}_j)} \mathbb{E} \left[\sum_j \frac{\eta_j}{\tilde{p}_j} \tilde{\omega}_j \right] - \frac{W}{2} \mathbb{V} \left[\sum_j \frac{\eta_j}{\tilde{p}_j} \omega_j \right].$$

Given that the idiosyncratic components are uncorrelated, optimal residual demand functions (relative to the outside good) have the simple form

$$\tilde{\omega}_j^* = \frac{\mu_j \tilde{p}_j}{\tilde{W} \sigma_j^2},$$

where σ_j is the variance of η_j . Observe that the residual elasticity is

$$\frac{\partial \tilde{\omega}_j^*}{\partial \tilde{p}_j} \frac{\tilde{p}_j}{\tilde{\omega}_j} = \frac{\mu_j}{\tilde{W} \sigma_j^2} = 1.$$

This elasticity is low and constant, and uninformative about asset elasticities.

E Dynamic Trading

Consider a two-period variant of our baseline model. Trees are durable assets which pay dividends in two periods. Investors trade at the beginning of each period, and consume the per-period payoffs generated by their trees at the end of each period. The payoff structure is the same as in our baseline framework, and the aggregate state processes $z \in \{1, 2\}$ and $\iota \in \{g, r\}$, which determine payoffs, are *i.i.d.* across periods. The discount factor is given by $\delta \in (0, 1)$.

We denote by \mathbf{S}_t the aggregate state variable sufficient for determining prices,

which naturally includes the aggregate distribution over investor wealth and tastes. We use this structure to introduce potential variations in market prices over time. In particular, investors know the realization of \mathbf{S}_1 when forming portfolio allocations at date 1, and their choices are also influenced by their expectations of \mathbf{S}_2 .

We write the price of asset $j \in \mathcal{J} \equiv \{g, r, 2\}$ as a function of the aggregate state: $p_{j,t} = P_j(\mathbf{S}_t)$ for some endogenous function P_j . The state-contingent gross return on asset $j \in \mathcal{J}$ is $R_j(\mathbf{S}_2) = \frac{P_j(\mathbf{S}_2)}{P_j(\mathbf{S}_1)}$. Given heterogeneity in tastes, a single investor will thus be concerned with the fact that changes in the preferences or wealth of other investors can induce changes in prices and, therefore, her perceptions of expected returns. It is sufficient for our purposes to consider the decision problem of a single investor who takes as given the stochastic process over the aggregate state. The investor's wealth at the beginning of period t is determined by the realized state and previous asset positions: $W_t^i(\mathbf{a}_{t-1}^i, \mathbf{S}_t) = \sum_{j \in \mathcal{J}} P_j(\mathbf{S}_t) a_{j,t-1}^i$.

We denote investor i 's purchases of trees at the beginning of period $t \in \{1, 2\}$ by $\mathbf{a}_t^i = (a_{g,t}^i, a_{r,t}^i, a_{2,t}^i)$, and we denote by \mathbf{a}_0^i investor i 's exogenous endowment. The investor i 's state variable at the beginning of period $t \in \{1, 2\}$ consists of the portfolio of asset positions purchased in the previous period, \mathbf{a}_{t-1}^i , and the current-period taste parameters $\Theta_t^i = (\theta_{j,t}^i)_{j \in \{g, r, 2\}}$, which are permitted to evolve stochastically over time. The individual state at the beginning of period t is therefore $\mathbf{s}_t^i = (\mathbf{a}_{t-1}^i, \Theta_t^i)$.

We solve the problem by backwards induction, assuming that investors face short-sale constraints. For ease of exposition, assume that green and red trees are perfect substitutes: $\epsilon = 0$. In this case, the second-period choice between green and red trees is bang-bang: the investor buys only green trees if $\theta_{g,2}^i / P_g(\mathbf{S}_2) > \theta_{r,2}^i / P_r(\mathbf{S}_2)$, and only red trees if the inequality is reversed. We then have the following characterization of the second-period value function.

Lemma 2 (Value function) *Let $\epsilon = 0$ and $u = \log$. The second-period value function satisfies:*

$$V_2^i(\mathbf{s}_2^i, \mathbf{S}_2) = H(\Theta_2^i, \mathbf{S}_2) + \log \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) \right),$$

where $H(\Theta_2^i, \mathbf{S}_2)$ depends on investor i 's tastes and market prices in period 2, but is independent of any investor choices at date 1.

Proof. Following the discussion of the static optimization problem, the second-period

value function can be written as:

$$V_2^i(\mathbf{s}_2^i, \mathbf{S}_2) = \max_{\mathbf{a}_2^i \geq 0} \pi_1 \left[\rho u \left(\theta_{g,2}^i y_g(r) a_{g,2}^i + \theta_{r,2}^i y_r(r) a_{r,2}^i \right) + (1 - \rho) u \left(\theta_{g,2}^i y_g(g) a_{g,2}^i + \theta_{r,2}^i y_r(g) a_{r,2}^i \right) \right] \\ + \pi_2 u \left(\frac{y(2)}{P_2(\mathbf{S}_2)} \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) - P_g(\mathbf{S}_2) a_{g,2}^i - P_r(\mathbf{S}_2) a_{r,2}^i \right) \right).$$

Since green and red assets are perfect substitutes (i.e., $\epsilon = 0$), the solution to the second-period decision problem is bang-bang, depending on whether $\theta_g^2/P_g(\mathbf{S}_2) \geq \theta_r^2/P_r(\mathbf{S}_2)$. Suppose that this inequality holds (i.e., green trees are cheap). Then, the value of the second-period problem is:

$$V_{g,2}^i(\mathbf{s}_2^i, \mathbf{S}_2) = \max_{a_{g,2}^i \geq 0} \pi_1 u \left(\theta_{g,2}^i y(1) a_{g,2}^i \right) + u \left(\frac{y(2)}{P_2(\mathbf{S}_2)} \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) - P_g(\mathbf{S}_2) a_{g,2}^i \right) \right).$$

The first-order condition together with log utility yields

$$a_{g,2}^i = \frac{\pi_1 W_2^i(\mathbf{a}_1^i, \mathbf{S}_2)}{P_g(\mathbf{S}_2)}.$$

This means that

$$V_{g,2}^i(\mathbf{s}_2^i, \mathbf{S}_2) = \pi_1 u \left(\theta_{g,2}^i \frac{\pi_1 y(1) W_2^i(\mathbf{a}_1^i, \mathbf{S}_2)}{P_g(\mathbf{S}_2)} \right) + \pi_2 u \left(\frac{y(2)}{P_2(\mathbf{S}_2)} \pi_2 W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) \right).$$

With log utility this can be written as

$$V_{g,2}^i(\mathbf{s}_2^i, \mathbf{S}_2) = \pi_1 \log \left(\frac{\theta_{g,2}^i}{P_g(\mathbf{S}_2)} \right) + \pi_2 \log \left(\frac{1}{P_2(\mathbf{S}_2)} \right) + \pi_1 \log \left(\pi_1 y(1) \right) + \pi_2 \log \left(\pi_2 y(2) \right) \\ + \log \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) \right).$$

Similarly, if $\theta_g^2/P_g(\mathbf{S}_2) \leq \theta_r^2/P_r(\mathbf{S}_2)$, the value of the second-period problem is:

$$V_{r,2}^i(\mathbf{s}_2^i, \mathbf{S}_2) = \pi_1 \log \left(\frac{\theta_{r,2}^i}{P_r(\mathbf{S}_2)} \right) + \pi_2 \log \left(\frac{1}{P_2(\mathbf{S}_2)} \right) + \pi_1 \log \left(\pi_1 y(1) \right) + \pi_2 \log \left(\pi_2 y(2) \right) \\ + \log \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) \right).$$

Hence, the second-period value function is written as

$$V_2^i(\mathbf{s}_2, \mathbf{S}_2) = \max \left\{ V_{g,2}^i(\mathbf{s}_2, \mathbf{S}_2), V_{r,2}^i(\mathbf{s}_2, \mathbf{S}_2) \right\},$$

and it satisfies

$$V_2^i(\mathbf{s}_2, \mathbf{S}_2) = H(\Theta_2^i, \mathbf{S}_2) + \log \left(W_2^i(\mathbf{a}_1^i, \mathbf{S}_2) \right),$$

where

$$\begin{aligned} H(\Theta_2^i, \mathbf{S}_2) = & \pi_1 \max \left\{ \log \left(\frac{\theta_{r,2}^i}{P_r(\mathbf{S}_2)} \right), \log \left(\frac{\theta_{g,2}^i}{P_g(\mathbf{S}_2)} \right) \right\} + \pi_2 \log \left(\frac{1}{P_2(\mathbf{S}_2)} \right) \\ & + \pi_1 \log \left(\pi_1 y(1) \right) + \pi_2 \log \left(\pi_2 y(2) \right). \end{aligned}$$

■

The separability of the value function into a taste component and a wealth component follows from log utility. However, this feature is not essential for our results below. What is important is that investors take into account future market prices when forming portfolios. Indeed, weakening separability would further complicate identification.

Now turn to the first-period decision problem. Maintaining the assumptions of log utility and normalizing $P_2(\mathbf{S}_1) = 1$, we can then write the period 1 decision problem under discount factor δ as:

$$\begin{aligned} V_1^i(\mathbf{s}_1, \mathbf{S}_1) = & \max_{\mathbf{a}_1^i \geq 0} (1 - \delta) \left[\pi_1 \log \left(y(1) \left(\theta_{g,1}^i a_{g,1}^i + \theta_{r,1}^i a_{r,1}^i \right) \right) \right. \\ & \left. + \pi_2 \log \left(y(2) \left(W_1^i(\mathbf{a}_0^i, \mathbf{S}_1) - P_g(\mathbf{S}_1) a_{g,1}^i - P_r(\mathbf{S}_1) a_{r,1}^i \right) \right) \right] \\ & + \delta \mathbb{E}^i \left[H(\Theta_2^i, \mathbf{S}_2) + \log \left(\sum_{j \in \mathcal{J}} P_j(\mathbf{S}_2) a_{j,1}^i \right) \right], \end{aligned}$$

where the expectation \mathbb{E}^i is taken with respect to private tastes and the aggregate state variable at date 2. Differentiating this object with respect to $a_{j,1}^i$ yields (13).

F Tastes versus constraints

We use the three-asset economy with the following simplifying assumptions:

- (i) There are two types of investors, $i \in \{h, \ell\}$, each of which faces short-sale constraints. Tastes satisfy: $\theta_g^h = 1 + t$ and $\theta_r^h = 1 - t$, while $\theta_g^\ell = 1 - t$ and $\theta_r^\ell = 1 + t$.
- (ii) Type h owns share $\gamma \geq \frac{1}{2}$ of the aggregate endowment of each tree.
- (iii) There is no aggregate risk, $y(1) = y(2) = 1$ and $\pi_1 = \frac{1}{2}$.

We begin with the baseline where investors face only short-sale constraints. In this case, taste differences can lead to endogenous sorting in equilibrium.

First, look for an equilibrium in which type h specializes in green trees while type ℓ investor specializes in red trees. It follows from the first-order conditions with respect to a_g^h and a_r^ℓ and the aggregate resource constraint on tree 2 that

$$\frac{p_r}{p_g} = \frac{1 - a_2^h}{a_2^h} \quad \text{and} \quad p_1 \equiv \frac{p_g + p_r}{2} = \frac{\pi_1}{1 - \pi_1} = 1,$$

where the last equality follows from our simplifying assumptions. Substituting p_1 and the first-order condition with respect to a_g^h into type h investor's budget constraint, together with the aggregate resource constraint, one can show:

$$(a_2^h, a_2^\ell) = (\gamma, 1 - \gamma).$$

Thus, if sorting occurs in equilibrium, then the prices satisfy

$$(p_g, p_r) = \left(2\gamma \frac{\pi_1}{1 - \pi_1}, 2(1 - \gamma) \frac{\pi_1}{1 - \pi_1} \right) = (2\gamma, 2(1 - \gamma))$$

and the investors' portfolio choices are

$$(a_g^h, a_r^h, a_2^h) = \left(\frac{1}{2}, 0, \gamma \right) \quad \text{and} \quad (a_g^\ell, a_r^\ell, a_2^\ell) = \left(0, \frac{1}{2}, 1 - \gamma \right).$$

To show that this constitutes an equilibrium, we need to show that the first-order conditions with respect to a_r^h and a_g^ℓ hold at the zero holding. It can be seen that these first-order conditions are met as long as type h 's incentive is satisfied:

$$\gamma \leq \bar{\gamma} \equiv \frac{\theta_g^h}{\theta_g^h + \frac{(y(1))^2 + \epsilon^2}{(y(1))^2 - \epsilon^2} \theta_r^h} = \frac{\theta_g^h}{\theta_g^h + \theta_r^h}.$$

Next, we consider an equilibrium in which one type of investor, denoted by i , holds both green and red trees while the other type specializes in one tree. Guessing that $\tilde{c}^i(g) = \tilde{c}^i(r) = \tilde{c}^i(2)$, first-order conditions with respect to a_g^i and a_r^i imply that

$$(p_g, p_r, p_1) = (\theta_g^i, \theta_r^i, 1).$$

Combining these conditions with the budget constraint implies that

$$a_2^i = (1 - \pi_1)e_2^i + \pi_1 \frac{y(1)}{y(2)} (\theta_g^i e_g^i + \theta_r^i e_r^i) = \frac{e_2^i + \theta_g^i e_g^i + \theta_r^i e_r^i}{2}.$$

We also guess and verify that $i = h$. Since type ℓ specializes in red, $a_g^\ell = 0$ and $a_r^\ell = \frac{1}{2}$. Then, by the aggregate resource constraint,

$$a_r^h = \frac{a_2^h - \theta_g^h E_g}{\theta_r^h}.$$

Since the first-order condition with respect to a_r^ℓ yields $a_r^\ell = \frac{a_2^\ell}{p_r}$, it follows from the budget constraint that

$$a_2^\ell = (1 - \pi_1)(e_2^\ell + p_g e_g^\ell + p_r e_r^\ell) = 1 - \gamma \quad \text{and} \quad a_r^\ell = (1 - \pi_1) \frac{e_2^\ell + p_g e_g^\ell + p_r e_r^\ell}{p_r} = \frac{1 - \gamma}{p_r}.$$

Thus, we obtain, as in the statement of the proposition,

$$a_r^\ell = \frac{1 - \gamma}{2} \frac{E_2 + \theta_g^h E_g + \theta_r^h E_r}{\theta_r^h}.$$

When $\gamma > \bar{\gamma}$, it can be seen that $a_r^h > 0$ and that the first-order condition with respect to a_g^ℓ at $a_g^\ell = 0$ is also met (i.e., $a_g^\ell = 0$).