Demand-System Asset Pricing: Theoretical Foundations*

Preliminary – comments welcome.

William Fuchs† Satoshi Fukuda‡ Daniel Neuhann§

March 18, 2024

Abstract

Recent approaches to asset pricing involve the estimation of demand systems for financial securities in which investors are permitted to have non-pecuniary *tastes* over cash flow-irrelevant asset characteristics. We investigate theoretical foundations of demand-system asset pricing using multiple approaches to integrating tastes with portfolio choice. Our analysis raises several conceptual issues, including the definition of no arbitrage, the pricing of "redundant" assets, and the cardinal interpretation of taste parameters. These issues imply multiple barriers to identifying demand systems for financial securities from observational data, and raise questions about the structural interpretation of financial demand elasticities. We discuss how these issues affect counterfactuals constructed from estimated demand systems.

*Special thanks to Andres Almazan and Andrey Ordin for detailed discussions. We thank Philippe van der Beck, Nicolae Garleanu (discussant), Amy Huber, Travis Johnson, Deborah Lucas, Aaron Pancost, Lukas Schmid, Yannis Stamatopoulos, Jonathan Wallen and seminar attendees at ASU Sonoran Winter Finance Conference, Finance Theory Webinar and UT Austin for useful feedback.

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

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

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

1 Introduction

Recent approaches to asset pricing following Koijen and Yogo (2019) involve the estimation of *demand systems* for financial assets in which investors are permitted to have *tastes* (or dogmatic beliefs) over both pecuniary and non-pecuniary attributes of financial assets, such as environmental and social scores or the identity of the issuer. According to this approach, data on portfolio holdings can be used to identify investor tastes and beliefs, and encodes information absent from prices that can inform researchers about the equilibrium response to a variety of counterfactual shocks, such as evolving tastes or changes in the wealth distribution across investor types (Koijen, Richmond, and Yogo, Forthcoming).

This focus on non-pecuniary tastes and portfolios represents a sharp break from neoclassical asset pricing, which emphasizes the fungibility of securities up to cash flows and focuses on price data disciplined by no arbitrage (Ross, 2004). Given the unsatisfactory empirical record of neoclassical approaches, this stark dichotomy may well be an advantage. However, introducing preferences over the *provenance* of cash flows also means jettisoning much of the well-established theoretical foundations that underlie the neoclassical approach. Thus, as we move forward in this direction, it is important to establish solid foundations for this approach and understand their implications for empirical work.

In this paper, we examine the theoretical foundations of demand-based asset pricing by way of synthesis with canonical models of portfolio choice. In particular, we enrich an endowment economy as in Lucas (1978) with potentially payoff-irrelevant characteristics that affect investors' *tastes* over assets, and derive portfolio demand functions that are sensitive to risk, return and tastes.[1] This allows us to distinguish "hedonic characteristics," such as environmental scores or issuer identify, from "cash flow characteristics" that can be used to summarize the statistical properties of the cash flow distribution.

The incorporation of hedonic characteristics forces us to revisit some fundamental conceptual issues in theoretical asset pricing, such as the appropriate definition of no arbitrage and the use of cardinal versus ordinal utility indices for tastes. With respect to arbitrage, we find that, depending on the security menu, there may not exist price systems that preclude arbitrage opportunities for all investors simultaneously. This

[1] As will become clear, there is close correspondence between preference-based tastes and dogmatic beliefs about asset returns. We focus on tastes to ease the exposition, but our results apply to dogmatic beliefs as well.

concern is particularly salient when there are assets with similar cash flows over which investors have different tastes, such as stocks and "replicating portfolios" of derivatives that do not inherit the same tastes. Non-pecuniary tastes can thus invalidate the organizing principles underlying stochastic discount factors and state prices, making it difficult to construct demand systems that are stable under counterfactuals. With respect to the modeling of non-pecuniary tastes, we find that integrating risk and return with tastes *requires* a cardinal interpretation of taste parameters. In empirical contexts, this implies that estimated demand systems may be highly sensitive to the assumed asset span and (arbitrary) changes in the unit of measurement for tastes, and particularly so if short sales are either prohibited or unobserved. As such, demand systems estimated using state-of-the-art methods may produce valid counterfactuals only for a narrow range of scenarios.

Why do these challenges rise? Demand estimation is perhaps the central problem in industrial organization. Accordingly, researchers have developed a rich set of tools to estimate preferences parameters from observational data (Berry and Haile, 2021). Yet, financial assets present a number of unique challenges that differ sharply from the demand for nonfinancial goods typically studied in industrial organization. Thus, as we discuss, the application of existing techniques to financial securities is far from straightforward.

The first challenge pertains to no arbitrage and the pricing of redundant securities. In contrast to most consumer good settings, investors who find a security too expensive are not forced to exit the market; they can (short) sell the asset. Importantly, they may do so either directly or by trading an alternative portfolio that replicates the asset's cash flows. To discipline equilibrium prices and trading behavior given these considerations, neoclassical asset pricing uses the notions of the law of one price and no arbitrage, which is the idea that investors should not be able to receive "something for nothing."

When investors have tastes over non-pecuniary characteristics, two assets with identical cash flows need no longer have the same price. Thus, we need to modify the definition of the law of one price to account for the provenance of the cash flows. Yet, even with this broader understanding of the value of an asset, we find that with sufficiently dissimilar tastes (given a set of allowable trades) in general there would not exist a pricing function (stochastic discount factor) that leaves no arbitrage opportunities. This has direct implications for equilibrium existence, incentives for short sales, and for the use of demand systems for pricing untraded securities. In applied contexts, estimated demand systems may be quite sensitive to misspecified investment universes (i.e., the set of assets investors can trade, and the extent to which investors can engage in short sales).

The second pertains to investor preferences. Models of portfolio choice are based on theories of choice under uncertainty, such as expected utility theory. As is well known, expected utility theory imposes a *cardinal* interpretation of utility, which is a stronger requirement than the ordinal rankings typically assumed in consumer good demand systems. Under a cardinal interpretation of tastes, identifying financial demand systems requires identification of the *intensity* of tastes *relative* to risk-return considerations, not just their ordinal ranking. In order to integrate non-pecuniary tastes with risk and return, one must therefore admit a cardinal interpretation of taste parameters as well.

This fact has practical relevance. For counterfactuals, preference parameters must be estimated for inframarginal investors as well. Moreover, portfolio choice is generically sensitive to rank-preserving transformations of tastes, such as changes in the units of measurement. Yet, investors may not agree on how to evaluate the "greenness" of an asset. They may also find little consensus on how to *aggregate* multiple hedonic characteristics, such as a firm's environmental social scores, into a single asset-level score. Empirical demand systems may be sensitive to such choices.

Having discussed these conceptual issues, we use a fully specified model of portfolio choice under tastes to examine questions of identification and counterfactuals based on observational data. The model is a variant of the Lucas (1978) endowment economy in which assets ("trees") may be endowed with payoff-irrelevant hedonic characteristics, and investors can differ in their tastes for these characteristics. Addressing the question of cardinal preferences, we work within the expected utility framework. This leads to a framework where investors trade off tastes against canonical risk-return considerations, and asset-level taste parameters can be interpreted in marginal utils.

Our model nests two important benchmarks. First, if two assets offer identical cash flows but differ in their hedonic characteristics, then an investor may prefer to buy only the one that aligns with his tastes. In this case, tastes lead to equilibrium sorting. Second, if all investors have the same tastes, then the model is equivalent to the canonical Lucas tree framework that considers only risk and return. When both channels are active, prices are affected by both fundamental cash flows and taste distributions.

This framework has several useful features. First, since we fully specify the model, we can explicitly solve for the "correct" counterfactual response to various shocks. Second, we can explicitly model common identification strategies that rely on mandates (i.e., exogenous restrictions on the type of assets an investor can hold) to obtain variation in prices. Hence, it is a useful laboratory for evaluating identification strategies and their challenges.

In general, demand system estimation requires price instruments whose variation is sufficient to identify tastes for specific assets. For a general portfolio choice problem, this is exceedingly difficult because the principle of diversification implies that any two assets may exhibit non-linear patterns of complementarity and substitutability that are shaped by the investor's overall portfolio holdings. Hence, in general, even "clean" variation in a single asset price is not enough to identify demand systems because endogenous portfolio changes create demand shocks in other assets (Berry and Haile, 2021).[2]

Even if price instruments are available, they may not suffice for identifying all taste parameters required to construct counterfactuals based on observational data. Given taste-based demand, an investor may only purchase her most preferred option in equilibrium, and may short sell less-preferred options. If short sales are prohibited, observed portfolio holdings may allow for inference on the *ordinal* ranking of a choice set, but there may not be enough information to estimate cardinal tastes for assets not purchased in equilibrium. When short sales are allowed but unobservable to the econometrician, as they often are in practice, demand systems may be misspecified. As we illustrate using our general equilibrium model, in either case it may be difficult to estimate taste parameters with sufficient accuracy to perform counterfactuals. More broadly, our analysis points to a critical question that must be addressed empirically, which is to what extent, if any, derivative securities inherit some of the taste properties of the underlying asset.

To evaluate these concerns, we use our framework to model common identification approaches used in the demand-system asset-pricing literature. For example, Koijen and Yogo (2019) use the sparsity of observed portfolios to construct instruments for asset prices. As the argument goes, sparse portfolios indicate that the investor may have tightly prescribed mandates that make it costly to hold other stocks. If mandates and

[2] One way around this problem is to restrict attention to settings where optimal portfolio weights are linear functions of own prices and characteristics (Koijen, Richmond, and Yogo, Forthcoming). In a static setting, this is possible with, e.g., CARA preferences and normally distributed shocks to payoffs. However, even these assumptions would not suffice in a dynamic setting. Given the importance of taste for portfolio choice, prices would be a function of endogenous changes in the wealth distribution.

fund flows are sufficiently exogenous to current investment opportunities, variation in the extent to which a particular stock is present in observed portfolios may be used to construct a demand shifter that is independent of prices. We use our model to assess this identification strategy, *taking as given* the exogeneity of mandates and fund flows. An important concern is that mandates and tastes may be observationally equivalent given equilibrium play, even as they have potentially very different implications for counterfactuals. The reason is that mandate investors are insensitive to price changes, whereas taste-based investors are not. As such, misjudging the share of mandate investors can lead to *qualitatively* different counterfactual prices in response to shocks. In addition, even if a given mutual fund has a particular mandate, investors in said fund can still reallocate their holdings across funds. As we discuss in Section 4.1, the identification strategy fails if this is the case.

These concerns are naturally entangled with wealth effects. While price changes always induce income and substitution effects, in most consumer good settings wealth effects are likely to be negligible and are thus frequently ignored when modeling demand. In contrast, financial assets are investment goods, and so wealth changes may have first-order effects in portfolio choice. In particular, an investor's measured elasticity for two otherwise identical assets may be very different depending on if they already hold the asset in their portfolio or not.[3] Hence, it is important to control for the evolution of portfolios when estimating demand elasticities.

We end by discussing the structural interpretation of demand elasticities. First, we discuss how the structural interpretation of demand elasticities depends critically on the investment universe. In particular, when there are close substitutes available, *asset-level* demand elasticities may be very high even when *consumption-level* elasticities are low. Second we show that even in canonical general equilibrium frameworks such as Lucas (1978), demand elasticities can range from near zero to infinite, depending on whether the driving shocks are common to all investors or merely introduce reallocative tades between investors. For these reasons, it is difficult to discriminate between models of asset pricing based on estimated demand elasticities alone.

The rest of the paper is structured as follows. The rest of this section discusses related literature. Section 2 discusses the fundamental identification problem and its impli-

[3] This concern is amplified when interpreting investors as financial intermediaries, since capital flows from households can induce wealth changes even when intermediary constraints and preferences are fixed.

cations with regards to incorporating non-pecuniary tastes in an asset pricing framework. Section 3 presents our framework. Section 4 studies the implications on identification and counterfactuals. Section 5 discusses the structural interpretability of asset demand elasticities. Section 6 provides concluding remarks.

Related Literature

This paper studies the theoretical foundations of two closely-related literatures: demand-based asset pricing that tries to model equilibrium returns using estimated portfolio choice models as in Koijen and Yogo (2019), and models in which investors may hold certain financial securities because of non-pecuniary values associated with them (Starks, 2023).

Demand-based approaches have been used to address a number of substantive questions. These include computing counterfactuals for price informativeness and sustainable investing in response to changing non-pecuniary tastes or changes in the size distribution of institutional investors (Koijen, Richmond, and Yogo, Forthcoming), global imbalances and currency returns (Jiang, Richmond, and Zhang, 2023) or corporate bond returns (Bretscher, Schmid, Sen, and Sharma, 2022).

The fact that estimated demand systems appear to reveal that financial institutions exhibit rather 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). We show that such elasticities may not always be interpretable as deep parameters, and may differ by the level of aggregation and the nature of the shock. Inelastic trading patterns are also often attributed to mandates, and these are used as instruments to identify demand parameters. Our analysis suggests that it is difficult to empirically distinguish between mandates and "tastes," and this matters for counterfactuals.

Valued-based approaches have been used to study investment in so-called "green assets," such as stocks or bonds associated with sustainable, environmentally-friendly firms or government expenditures. Pastor, Stambaugh, and Taylor (2021) provide an equilibrium model of such sustainable investment, whereby firms differ in their "green scores," but the set of marketable securities consists only of firm shares and a risk-free asset. We address implications and micro-foundations of such sustainable investing for financial market equilibrium with redundant securities under various forms of tastes, and ask how such tastes might be identified in equilibrium.

Pastor, Stambaugh, and Taylor (2022) provide early evidence that such tastes may be reflected in yield differences between otherwise identical German government bonds, while D'Amico, Klausmann, and Pancost (2023) show that the underlying "greenium" appears to have shrunk over time. Our analysis shows that, in the presence of short selling by at least some investors, premiums extracted from marginal prices (i.e., prevailing market prices) may not be informative about infra-marginal preferences, and thus may not be sufficient to infer counterfactuals. More generally, we argue that counterfactuals are sensitive to the precise modeling of environmental concerns, including to simple monotone transformations of taste parameters. As such, one conclusion of our paper is that there is research in modeling the precise foundations of tasted-based investment.

2 The Identification Problem and its Implications

We begin by reviewing the fundamental identification problem as it relates to estimating demand systems for financial assets from observational data. Since demand estimation is a classic issue, particularly in industrial organization, we focus mainly on the novel features introduced by modeling demand for financial assets. We then examine the theoretical implications of models that permit features thought to be useful for identification.

In general, the goal of demand estimation is to measure market participants' willingness to pay for different assets. The main difficulty is that, since quantities and prices are generally jointly determined, simple regressions of quantities on prices do not identify structural parameters. Figure 1 illustrates this basic problem, as well as a potential solution. 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. Quantities may become negative because financial assets can be sold short.

A common empirical strategy is to trace out demand respones to exogenous changes in supply. However, this is not feasible in an endowment economy with fixed supply, which is the basic framework used in demand-system asset pricing. Even outside of endowment economies, it is difficult to find changes in the supply of a given financial asset that are not influenced by prevailing prices. To circumvent this issue, the literature on demand-system asset pricing focuses on variation in *net supply*, defined as aggregate sup-

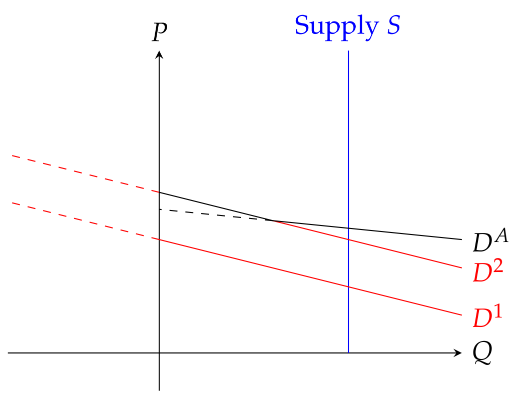

(a) Aggregate Supply and Demand.

(b) Net Supply and Demand.

Figure 1: The basic identification issue in an endowment economy.

ply minus demand of a subset of market participants. The basic idea is illustrated in the right panel: rather than finding exogenous shocks to aggregate supply S, researchers aim to estimate the structural parameters of demand function D^1 by finding exogenous variation in residual supply $S - D^2$. Put differently, the empirical strategy is to construct exogenous shocks to the residual supply curve faced by a particular investor by finding exogenous shocks to the demand functions of other investors.

This approach places stringent constraints on the type of variation that can be used to identify demand systems. In particular, researchers must find settings in which there are changes in residual supply (which is itself a type of demand shock) that are uncorrelated with the demand of remaining 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 payoffs. Since financial assets are investment goods whose current value generically depends on their resale value, these requirements extend not only to preferences over current cash flows, but also expected future prices. This presents additional identifiation challenges which we discuss in more detail in Section 2.5.

Two broad approaches for such shocks have been proposed. The first relies on cross-investor heterogeneity in *tastes* for particular assets, holding fixed a certain notion of expected cash flows. Such taste differences could be due to differences in investor preferences over the *provenance* of cash flows, such as when some investors prefer to invest in environmentally-friendly firms, while another does not care. Or, investors may have dogmatic heterogeneous beliefs about future cash flows that are orthogonal to the beliefs of other investors, so that investors do not update their own beliefs in response to changes in residual demand.

The second approach relies on the idea is that, if some funds are unable to invest in a particular asset for *exogenous* reasons, then it appears plausible that this asset will be cheaper than a similar asset that can be held by these funds. This approach is implemented empirically by trying to infer (unobservable) *constraints* that affect investor's investment opportunities, such as short-sale constraints or differences in investment mandates across investment funds.

Both approaches are relevant in practice: a growing literature documents that investors may care about non-pecuniary asset characteristics such as ESG scores (Starks, 2023), and models with belief differences have a long history in asset pricing. Finally, a wide array of funds are subject to a variety of investment mandates that constrain their portfolio choices.

At the same time, there are a number of open questions regarding the theoretical foundations of asset pricing with heterogeneous tastes, and about the feasibility of identifying demand systems in the presence of tastes and (unobserved) constraints. Moreover, well-established tools from industrial organization may have only limited applicability to financial markets because financial assets differ from consumer goods in at least three important ways: (i) assets can be flexibly bundled and unbundled using portfolios, (ii) preferences must admit a cardinal (rather than merely ordinal) interpretation, and (iii) there are resale considerations that affect current valuations.

The rest of this section addresses implications of these issues in more detail. Section 2.1 provides two ways of incorporating tastes in a canonical asset-pricing framework. Section 2.2 argues that asset pricing with tastes requires cardinal interpretations of taste parameters. Hence identifying these parameters presents greater challenges than in many classical settings from industrial organization that aim to identify ordinal rankings only. Section 2.3 demonstrates that tastes may invalidate standard notions of no arbitrage. This means that portfolio choice will generally be sensitive to assumptions on the security menu, and also that introducing tastes will generally require supplemental assumptions on the strategy space. Section 2.4 shows that typical (combinations of) utility functions, asset menus, and investment constraints generate cross-asset complementarities in portfolio choice, and that this renders it difficult to estimate asset demand systems. Finally, Section 2.5 considers implications of dynamic trading in settings where investors differ in terms of their tastes.

2.1 Heterogenous Tastes for Financial Assets

Demand-based asset pricing allows investors to have preferences over asset characteristics that need not be directly related to cash flows. Incorporating tastes within canonical asset pricing frameworks requires making a number of conceptual decisions that can alter some of the key theoretical underpinnings of asset pricing theory. This is because theories of choice under uncertainty place relatively stringent constraints on preferences.

The basic framework is standard. In line with Koijen and Yogo (2019), we consider a one-shot portfolio choice problem in which an investor can choose to consume at date 0 and/or at date 1. A random state of the world $z \in \mathcal{Z} \equiv \{1, \ldots, Z\}$ is realized at date 1, and the probability of state z is $\pi_z \in (0, 1)$. The set of assets is $\mathcal{J} \equiv \{1, \ldots, J\}$. Asset $j \in \mathcal{J}$ offers state-contingent cash flows $y_j(z)$ in state z. There is a set of investors indexed by i. Investor i has a von Neumann-Morgenstern utility function defined over lotteries.

Within this framework, we consider two main taste specifications. The first is consumption-augmenting tastes, defined as additional "consumption-equivalent" value that is generated by an asset of particular *provenance*. The second is additive-separable tastes, by which we mean that the investor obtains some additional value (or disutility) from holding certain assets that is separable from risk-return considerations. We show that both formulations deliver essentially identical conclusions.

Consumption-Augmenting Tastes. Under consumption-augmenting tastes, investor i evaluates her payoffs from holding portfolio $(a^i_j)_{j \in \mathcal{J}}$ by both the cash flows it generates and her *tastes* $(\theta^i_j)_{j \in \mathcal{J}}$ over assets, where $\theta^i_j > 0$. In particular, we assume that preferences are defined over the *effective units of consumption* delivered by a portfolio $(a^i_j)_{j \in \mathcal{J}}$ for investor i in state z, and define these as

$$\tilde{c}^i_1(z) \equiv \sum_{j \in \mathcal{J}} \theta^i_j y_j(z) a^i_j + w^i_1(z),$$

where $w_1^i(z) \geq 0$ is a non-marketable endowment.[4]

Investor i's portfolio maximization problem is to maximize expected utility over effective consumption subject to budget balance:

$$\begin{aligned} \max_{(a_j^i)_{j\in\mathcal{J}}} \quad & (1-\beta)u^i(c_0^i) + \beta \sum_{z\in\mathcal{Z}} \pi_z u^i(\tilde{c}_1^i(z)) \\ \text{s.t.} \quad & c_0^i = w_0^i - \sum_{j\in\mathcal{J}} p_j(a_j^i - e_j^i) \\ & \tilde{c}_1^i(z) = \sum_{j\in\mathcal{J}} \theta_j^i y_j(z) a_j^i + w_1^i(z), \end{aligned} \tag{P-CA}$$

where β is the discount factor, u^i is the utility function, p_j is the price of asset j, w_0^i is initial wealth, and e_j^i is investor i's endowment of asset j.

Effective consumption is useful for capturing the notion that an investor may, for example, value cash flows produced by environmentally-friendly firms more than an identical cash flows stream produced by other firms. As such, tastes differentiate effective consumption from pure consumption $c_1^i(z) = \sum_{j\in\mathcal{J}} y_j(z)a_j^i + w_1^i(z)$.

Remark 1 *There is a close correspondence between consumption-augmenting tastes and heterogeneous beliefs. In particular, it is generally possible to enrich the state space over which payoffs are defined to include "taste-based payoffs." Heterogeneous tastes can then be mapped into heterogeneous beliefs if we let investors differ in their probability assessments over this augmented state space. An important consideration in this regard is that such taste-related beliefs are dogmatic: investors must agree to disagree, and in particular they may disagree on whether a particular state of the world can be realized. Interestingly, such strong disagreement is desirable when trying to construct instruments for residual demand because it allows for the possibility of orthogonal demand shocks. However, we show below that it also comes with more undesirable consequences.*

[4] One can also define consumption-augmenting tastes in an additive manner: $\tilde{c}^i(z) = \sum_j(\theta_j^i + y_j(z))a_j^i + w_1^i(z)$. The main difference is that tastes operate like a "risk-free" component of returns for every asset, with obvious implications for portfolio choice. Overall, however, the main conclusions are unchanged.

Additive-Separable Tastes. Next we consider additive-separable tastes. Fix a function G^i that maps portfolio $(a^i_j)_{j\in\mathcal{J}}$ into utils, and define the investor's decision problem as:

$$\begin{aligned}
\max_{(a^i_j)_{j\in\mathcal{J}}} \quad & (1-\beta)u^i(c^i_0) + \beta\sum_{z\in\mathcal{Z}}\pi_z u^i(c^i_1(z)) + G^i\left((a^i_j)_{j\in\mathcal{J}}\right) && \text{(P-AS)}\\
\text{s.t.} \quad & c^i_0 = w^i_0 - \sum_{j\in\mathcal{J}} p_j(a^i_j - e^i_j)\\
& c^i_1(z) = \sum_{j\in\mathcal{J}} y_j(z)a^i_j + w^i_1(z).
\end{aligned}$$

In this problem, the utility index u^i is defined in the standard way over pure consumption $c^i_1(z) = \sum_{j\in\mathcal{J}} y_j(z)a^i_j + w^i_1(z)$, but the overall objective is augmented by an additive value to holding a portfolio. This specification captures the idea that an investor may earn a "warm glow" from holding some stocks, or a disutility from holding others. As such, additive tastes are closely related to the way in which some people model "convenience yields," an idea that goes all the way back to Sidrauski (1967)'s model of money demand.

Additive-separable tastes differ from consumption-augmenting tastes in that non-pecuniary benefits of holding certain assets do not directly depend on the properties of the utility function. For example, tastes do not necessarily induce wealth or substitution effects in portfolio choice. While this may be an advantage for particular applications, it also has the drawback that it is generally difficult to discipline the particular functional form of G^i, even though the functional form *will* generally determine the trade-off between pecuniary and non-pecuniary aspects of portfolio choice.[5]

2.2 Cardinal Interpretation of Tastes

Given the nature of the questions studied, standard methods in industrial organization are mainly developed for settings where it is sufficient to identify *ordinal* preferences over alternatives. This is because standard consumer preference rankings can be fully represented by ordinal utility functions. This simplification does not apply in the context of portfolio choice, where expected utility theory requires admitting a *cardinal* interpretation of utility functions. As we now establish, this has the consequence that even linear

[5] A further consideration is that if we have taste shocks (both aggregate and idiosyncratic) this would be another risk factor to contemplate when forming optimal portfolios. We abstract in the current paper from the additional considerations for identification this would entail.

transformations of taste parameters which leave ordinal rankings unchanged may lead to different portfolio choices.

We start with formulating the optimality conditions under heterogeneous taste specifications. To that end, under consumption-augmenting tastes, we define investor i's marginal rate of substitution over effective consumption in Program (P-CA) as:

$$\tilde{\Lambda}^i(z) \equiv \frac{\pi_z \beta u^{i'}(\tilde{c}_1^i(z))}{(1-\beta)u^{i'}(c_0^i)},$$

where $u^{i'}$ is marginal utility. Under additive-separable tastes, investor i's marginal rate of substitution over pure consumption in Program (P-AS) is analogous:

$$\Lambda^i(z) \equiv \frac{\pi_z \beta u^{i'}(c_1^i(z))}{(1-\beta)u^{i'}(c_0^i)}.$$

Then, the optimality condition for a_j^i under consumption-augmenting tastes is:

$$\theta_j^i \sum_{z \in \mathcal{Z}} y_j(z) \tilde{\Lambda}^i(z) = p_j.$$

The optimality condition for a_j^i under additive-separable tastes is:

$$\sum_{z \in \mathcal{Z}} y_j(z) \Lambda^i(z) + g_j^i \left((a_j^i)_{j \in \mathcal{J}} \right) = p_j,$$

where g_j^i is the partial derivative of G^i with respect to a_j^i.

These conditions relate the standard risk-return tradeoff, measured by the distribution of marginal utility across states, to investor tastes. The solutions to these programs depend on the functional form of tastes and are not invariant to simple monotone transformations. This means that identifying demand systems requires measuring the intensity of tastes, not just their ordinal ranking.

Proposition 1 (Sensitivity to Rank-preserving Transformations)

1. *The solution to Program (P-CA) is sensitive to rank-preserving transformations of $\theta^i = (\theta_j^i)_{j \in \mathcal{J}}$, including positive linear transformations and mean-preserving spreads of taste parameters.*

2. *The solution to Program (P-AS) is sensitive to monotone transformations of G^i, including linear positive transformations such as the re-scaling of the units in which tastes are*

measured.

The intuition for this result is straightforward. In the case of consumption-augmenting tastes, increasing the nominal value of θ_j increases the consumption-equivalent value of holding asset j, leading the investor to allocate more funds to this asset and distorting the overall portfolio. Increasing tastes for all assets simultaneously raises the value of holding all assets, which leaves portfolio weights unchanged but alters the consumption-saving decision between dates 0 and 1.

In the case of additive-separable tastes, portfolio choices trade off the marginal increase in non-pecuniary values against the risk-return trade-off as measured by marginal utility over consumption. Since expected utility is cardinal, any rank-preserving transformation (such as a change in units) will alter the optimal portfolio.

2.3 No Arbitrage with Tastes

So far, we have discussed two methods for integrating tastes with portfolio choices and clarified their link to models of heterogeneous beliefs. We also argued that tastes may be useful for identification because they leave open the possibility of orthogonal demand shocks. At the root of this identification argument is that investors may have dogmatic differences in asset valuations. We now argue that this same feature may invalidate a key organizing principle of asset pricing and portfolio choice, namely no arbitrage.

No arbitrage is the notion that an equilibrium price system should not admit trades in which an investor receives something for nothing. Under no arbitrage, one can turn the problem of asset pricing into the problem of pricing state-contingent payoffs, which is much simpler than pricing individual assets when there is a large number of assets and potential portfolios, some of which may have (partially) redundant cash flows.[6]

Operationalizing no arbitrage requires a theory of value. In the neoclassical approach, the appropriate notion of value is cash flows, and an arbitrage is "something for nothing," or, more formally, "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). Hence a payoff space X (the set of attainable payoffs) and a pricing function p are sufficient to define no arbitrage.

[6]Such concerns do not arise in, say, consumer good settings, where a consumer is unable to combine parts of multiple cars to arrive at a more desirable bundle of characteristics.

Definition 1 (No Arbitrage (Cochrane, 2005)) *A payoff space X and pricing function p leave no arbitrage opportunities if, every payoff $x \in X$ that is weakly positive (i.e., $x \geq 0$) almost surely and strictly positive (i.e., $x > 0$) with some positive probability has positive price: $p(x) > 0$.*

This definition is critical for deriving basic properties of price systems, including the existence of positive stochastic discount factors. It has also been used to underpin the Arbitrage Pricing Theory which forms the backbone of factor approaches to modeling returns that is used in Koijen and Yogo (2019) to reduce the dimensionality of the asset space.

When investors differ in their tastes, Definition 1 is not sufficient to define no arbitrage because investors do not evaluate investment opportunities on the basis of cash flows alone. In particular, they may have subjective views on what constitutes "something for nothing" and thus a payoff space alone is not sufficient for defining valuations. We illustrate this concern using consumption-augmenting tastes (analogous results obtain with additive-separable tastes).

To define an appropriate notion of no arbitrage with tastes, we provide two preliminary definitions. First, letting a subset $\mathcal{A}$ of $\mathbf{R}^J$ denote the set of feasible portfolios and denoting by p_j the price of asset j, pricing function $P : \mathcal{A} \to \mathbf{R}$ 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 taste function $v^i : \mathcal{A} \to \mathbf{R}^Z$ that maps a portfolio a into a $1 \times Z$ vector $v^i(a)$ of state-contingent taste-augmented payoffs for investor i. In the absence of tastes, denoting by $Y = (y_j(z))_{j,z}$ the $J \times Z$ matrix of cash flows, all investors care only about cash flows as in standard asset pricing: $v^i(a) = a'Y$. With these notations, below we define no arbitrage with tastes.

Definition 2 (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 (i.e., $v^i(a) \geq 0$) almost surely and strictly positive (i.e., $v^i(a) > 0$) with strictly positive probability, the associated price is positive: $P(a) > 0$.*

We then have the following result regarding no arbitrage with tastes.

Proposition 2 (Generic Arbitrage Opportunities with Tastes) *Fix taste functions v^i for all investors. There does not exist pricing function P that leaves no arbitrage opportunities if and only if:*

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

A sufficient but not necessary 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:

$$v_j^i \geq v_{j'}^i \text{ and } v_j^{i'} \leq v_{j'}^{i'} \text{ with at least one inequality strict,}$$

where v_j^i is the marginal taste with respect to asset a_j.

Hence, no arbitrage fails if tastes are sufficiently heterogeneous and the asset menu is sufficiently rich. This can be most transparently seen in the following example.

Example 1 (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 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 1's taste-augmented payoffs for green and red assets satisfy $\theta_g^1 > \theta_r^1$, type 2 has $\theta_g^2 < \theta_r^2$.*

We consider a long-short portfolio consisting of selling one unit of the green asset and buying one unit of the red asset: $a = (-1, 1)$. The price of this portfolio is $P(a) = p_r - p_g$. Hence, investor i's taste-augmented payoff is $v^i(a) = \theta_r^i - \theta_g^i$. If there exists a pricing function P^* that leaves no arbitrage opportunities, then the absence of arbitrage opportunities for type 2 requires that $P^*(a) < 0$. Since type 1 can conduct the trade in reverse, no arbitrage for that type requires $P^*(a) > 0$, which is a contradiction. Hence, there does not exist pricing functions that leaves no arbitrage opportunities.

The result is a direct implication of the facts (i) that investors care about non-pecuniary factors, (ii) that taste differences are invariant to quantities (i.e., marginal valuation differences are invariant to portfolio holdings), and (iii) that investors are free to short either asset. Price changes are then not sufficient to equilibriate asset markets.

Two remarks are in order. First, it is well-known that problems of arbitrage and mispricing may also arise in models of dogmatic differences in beliefs.[7] To address these difficulties, theoretical models with heterogeneous beliefs generally impose strong restrictions on feasible strategies that preclude the existence of *risk-free* arbitrages. Prominent

[7]See, for instance, Hong and Stein (2007) for a survey on implications of heterogeneous beliefs on asset pricing.

examples include short sale constraints or sparse asset menus, such as restricting attention to one risk-free asset and one risky asset, with disagreement only about the dividends of the risky asset. It follows that models with tastes must generally rely on similar restrictions to ensure the existence of equilibria with desirable pricing properties. However, this gives rise to an additional concern that equilibrium outcomes are highly sensitive to the precise form of these restrictions. This presents challenges when trying to empirically estimate demand systems without a-priori knowledge of the precise investment opportunities and constraints faced by investors.[8]

Second, violations of no arbitrage may exist even if the law of one price (LOOP) holds conditional on the asset "color." In particular, in the presence of tastes, LOOP can be defined as requiring that two assets which deliver identical taste-augmented payoffs must have the same price. As the example shows, even if LOOP holds for individual assets, one can still construct portfolios over which investors have strict disagreements. This suggests a deeper question, which is whether and to what extent assets that are redundant in terms of their cash flows inherit the non-pecuniary benefits of underlying assets. For example, does a portfolio of an option and a bond generate similar tastes as a stock? Answers to questions such as this appear critical for developing a full-fledged theory of taste-based asset pricing and a necessary first step for proper measurement.

2.4 Mandates and Demand Complementarities

Koijen and Yogo (2019) emphasize the role of (unobserved) portfolio constraints and investment mandates in creating non-fundamental variation in asset prices that can serve to identify demand systems. To evaluate this strategy, we enrich decision problem (P-CA) by assuming that the investor faces (unobserved) constraints on portfolio choices. We then show that demand complementarities emerge whenever such constraints admit some degree of substitutability of assets.

Formally, we assume that investor i faces (unobserved) $K \geq 0$ constraints on port-

[8] An alternative approach is to restore no arbitrage using a notion of decreasing marginal tastes, whereby the marginal taste function converges to zero for sufficiently large asset holdings. This allows investors to potentially agree on marginal valuations even if they differ in their infra-marginal tastes. The empirical downside of this approach is that one must now identify an entire taste function.

folio choices. The k-th constraint is defined as

$$F^i_k(a^i, p) \leq 0,$$

where $a^i = (a^i_j)_{j\in\mathcal{J}}$ is investor i's portfolio, $p = (p_j)_{j\in\mathcal{J}}$ is the price vector, and the function $F^i_k(\cdot)$ is twice continuously differentiable in a^i_j for all j. For ease of exposition, we assume that the set of feasible portfolios induced by constraints $F^i_k \leq 0$ is convex.

A variety of constraints can be modeled in this way. First, a short-sale constraint on asset j says that $a^i_j \geq -\underline{a}_j$ for some constant $\underline{a}_j$. Second, an investor cannot invest in an asset when $a^i_j \leq 0$ and $a^i_j \geq 0$. Finally, an investment mandate which constrains portfolio weights can be modeled as follows. Let χ denote a vector of asset characteristics, where $\mathcal{C}(\chi)$ is the set of assets that share these characteristics. Given (a, p), the portfolio weight of assets with characteristic χ in investor i's portfolio is

$$\omega^i(\chi) = \frac{\sum_{j\in\mathcal{C}(\chi)} p_j a^i_j}{\sum_{j\in\mathcal{J}} p_j a^i_j}.$$

An *investment mandate over* χ is then a restriction that $\omega^i(\chi) \geq \underline{\omega}(\chi)$ and $\omega^i(\chi) \leq \overline{\omega}(\chi)$ for some constants $\underline{\omega}(\chi)$ and $\overline{\omega}(\chi)$.

Given these assumptions, investor i's portfolio maximization problem is

$$\begin{aligned}
\max_{a^i} \quad & (1-\beta)u^i(c^i_0) + \beta \sum_{z\in\mathcal{Z}} \pi_z u^i(\tilde{c}^i_1(z)) \\
\text{s.t.} \quad & c^i_0 = w^i_0 - \sum_{j\in\mathcal{J}} p_j(a^i_j - e^i_j) \\
& \tilde{c}^i_1(z) = \sum_{j\in\mathcal{J}} \theta^i_j y_j(z) a^i_j + w^i_1(z) \\
& F^i_k(a^i, p) \leq 0 \qquad \text{for all } k.
\end{aligned}$$

Define λ^i_k to be the Lagrange multiplier associated with constraint k, and $f^i_{k,j}(a^i, p)$ to be the partial derivative of $F^i_k(a^i, p)$ with respect to a^i_j. Then optimal portfolios are determined by the following system of first-order optimality conditions:

$$p_j = \theta^i_j \sum_{z\in\mathcal{Z}} y_j(z)\tilde{\Lambda}^i(z) + \sum_k \lambda^i_k \frac{f^i_{k,j}(a^i, p)}{(1-\beta)u^{i'}(c^i_0)} \quad \text{for all } j \in \mathcal{J}. \tag{1}$$

The system of equations defined by (1) features two forms of cross-asset restrictions that are generally of first-order importance in financial markets. First, due to the benefits of diversification, an investor's desired position in a particular asset depends on the holdings of all other securities (in particular, its covariance with the rest of the portfolio). As such, marginal rates of substitution across time and states generically depend on the entire vector of portfolio holdings.

Second, another form of cross-asset restrictions stems from portfolio constraints and operates whenever these constraints admit some degree of substitutability of assets. We provide three examples. The first example is a bond fund which may face a requirement to invest a certain proportion of its wealth in high-yield bonds but has flexibility over which particular bonds to invest in. The second example is an index fund that is designed to track a particular index but is permitted to have some degree of tracking error. The third example is a constraint on market-weighted portfolio shares held in different asset classes, for example, x% in stocks of firms with a high ESG score or a value-weighted ESG score.

In demand estimation, such interdependence between multiple goods (or assets) is referred to as *demand complementarities* and is known to have sharp implications for identification. In particular, given demand complementarities, even exogenous variation in a single price is generally not sufficient to identify specific demand parameters such as own-price elasticities (Berry and Haile, 2021). This is because complementarities induce endogenous movements in the prices of related goods, thereby contaminating the demand response to price changes. In financial markets, these issues are likely to be of first-order importance precisely because of the possibility of diversification and the prevalence of price-weighted mandates with tracking error.

Koijen and Yogo (2019) rely on instruments defined at the asset level, and impose restrictions on the decision problem to circumvent demand complementarities. In particular, they mute portfolio diversification considerations by modeling mean-variance investors with linear marginal utility who invest in characteristics-based portfolios that are suitably orthogonal to each other.[9] They address cross-asset restrictions in mandates by considering only *extensive margin quantity* restrictions on individual assets, whereby investors can hold positive quantities of only some assets (the so-called investment uni-

[9]Rationales for this approach must be statistical, since, as we demonstrated, no arbitrage does not necessarily hold under tastes.

verse), and no short positions at all. Since these constraints are purely asset-specific, they rule out the empirically relevant case where mandates permit some degree of substitution across assets. In addition, they rule out an important class of constraints where mandates depend on prevailing asset prices. Hence, the validity of the identification strategy depends on the realism of these restrictions in applied contexts.

Even taking this approach as given, it is important to note that asset-by-asset price instruments are generally not sufficient to *separately* identify tastes and constraints on portfolio holdings. In particular, different combinations of tastes and constraints can induce the same observed choices *given* a particular shock to market prices, but they may lead to different outcomes in counterfactuals. Hence, evaluating the validity of counterfactuals requires even more information on the underlying economy. Lastly, an additional issue that can contaminate the analysis is that even if individual funds are restricted by mandates, the individual investors are not.

2.5 Dynamic Trading

Another major difference between financial assets and consumer goods is that assets are investment goods whose value is at least partially determined by expected future resale considerations. In particular, irrespective of their personal valuation of the underlying cash flows, investors are willing to pay more for an asset today when the asset is expected to fetch a high price tomorrow. This means that investor valuations for long-lived assets will generally depend on other investors' valuations.

This can be seen in the classical framework of Harrison and Kreps (1978). In this model, different types of investors dynamically trade a single stock subject to short sale constraints and heterogeneous beliefs about future dividends. These assumptions are similar to those made in demand-system asset pricing: investors differ in their tastes for an asset, and they are subject to investment restrictions that preclude actionable arbitrage opportunities. The main difference is that Harrison and Kreps (1978) consider dynamic trading, which is a first-order concern in practice.

The main result in Harrison and Kreps (1978) is that asset prices are shown to admit a *speculative component*, whereby investors are willing to pay more for a stock than they would be if obliged to hold it forever. The reason is that short-sale constraints prevent pessimistic beliefs from affecting prices today, and investors are willing to pay a high price today if they get to resell to a future optimist.

This has direct implications for identifying asset demand systems in practice: since current willingness to pay depends on both private valuations and expected future market valuations, it is not clear *which* demand parameters (individual or market) can be measured by shifts in current quantities. Moreover, the fixed point problem between individual and market demand renders it infeasible to separately identify individual demands investor by investor.

3 Analytical Framework: Lucas (1978) with Tastes

The previous section discussed theoretical issues related to non-pecuniary tastes in asset pricing and requirements for identification of demand systems. This discussion left open the precise biases generated by these concerns and did not discuss broader meaning and counterfactual implications of estimated demand systems. We now construct a fully-specified model economy based on Lucas (1978) to address these issues.[10]

Environment. There is a single period of trading and all information is public. Asset payoffs depend on the realization of an aggregate state $z \in \{1, 2\}$ that is realized at date 1. The probability of state z is given by $\pi_z \in (0, 1)$. Associated with each state z is a Lucas tree that pays off $y(z)$ if the state is z. Trees are perfectly divisible, and the aggregate supply of each tree is equal to one.

There are two equally-sized trees associated with aggregate state 1 (but only with state 1): *red* and *green*. Conditional on aggregate state 1, the green tree pays $y_g(\iota)$ and the red tree pays $y_r(\iota)$, where $\iota \in \{r, g\}$ is a distributional shock that determines which of the two trees offers more cash flows. In particular, let

$$y_g(\iota) = \begin{cases} y(1) - \epsilon & \text{if } \iota = r \\ y(1) + \epsilon & \text{if } \iota = g \end{cases} \quad \text{and} \quad y_r(\iota) = \begin{cases} y(1) + \epsilon & \text{if } \iota = r \\ y(1) - \epsilon & \text{if } \iota = g \end{cases}.$$

[10] We use the consumption-augmenting approach to modeling tastes because this allows us to use much of the theoretical scaffolding of expected utility theory. It also it aligns closely with existing approaches in the literature (e.g., Koijen, Richmond, and Yogo, Forthcoming). However, the main results carry over to the case of additive-separable tastes.

Given these payoffs, it is clear that the distributional shock is fully diversifiable because

$$y(1) = \frac{1}{2}y_g(\iota) + \frac{1}{2}y_r(\iota) \text{ for all } \iota \in \{r, g\}.$$

Parameter $\epsilon \in [0, y(1))$ determines the substitutability of red and green trees. If $\epsilon = 0$, then red and green trees are perfect substitutes with respect to their cash flows. If $\epsilon > 0$, they are complements because holding both serves to diversify distributional risk. The probability of the red tree doing better is denoted by $\Pr(\iota = r) = \rho$.

Given this structure, one can think of assets as having two "characteristics:" the aggregate state of the world in which their payoffs accrue (i.e., 1 or 2), and their color. These characteristics determine in which states cash flows accrue (and thus serve as useful statistical summaries of the overall cash flow distribution), and they can also be used to define non-pecuniary tastes. When type 1 trees are perfect substitutes ($\epsilon = 0$), the color characteristic is irrelevant for cash flows and *only* matters through its link with tastes. When type 1 trees are imperfect substitutes ($\epsilon > 0$), even investors without tastes ($\theta^i_j = 1$) care about the color characteristic because it summarizes cash flow risk.

To focus on variation in the price of red and green trees, we use the following assumption in all of our numerical examples.

Assumption 1 (Aggregate symmetry) *Aggregate payoffs are $y(1) = y(2) = 1$, and the probability of each aggregate state is equal to one half, $\pi_1 = \frac{1}{2}$.*

Investors. There are two types of investors indexed by i.[11] Types determine an investors' endowment, tastes, and mandates. Specifically, let investor i be endowed with e^i_g, e^i_r and e^i_2 units of green, red, and state 2 trees, where aggregate feasibility dictates that

$$\sum_i e^i_j = \frac{1}{2} \text{ for } j \in \{g, r\} \quad \text{and} \quad \sum_i e^i_2 = 1.$$

Although this is not necessary, it is helpful to work with a relatively symmetric setting. Hence we will typically assume that type 1 owns share $\omega \geq \frac{1}{2}$ of the aggregate endowment of each tree, $e^1_g = e^1_r = \frac{\omega}{2}$ and $e^1_2 = \omega$.

Investor i takes positions a^i_j in asset $j \in \mathcal{J} \equiv \{g, r, 2\}$, and may be subject to short sale constraints: $a^i_j \geq 0$. The investor evaluates the payoffs of his portfolio using *effective*

[11] Our results readily generalize to many types, or to a continuum of types.

units of consumption. Similarly to Section 2.1, this object is defined as

$$\tilde{c}^i(\iota) \equiv \sum_{j \in \mathcal{J}} \theta^i_j y_j(\iota) a^i_j \text{ for each } \iota \in \{g, r, 2\}.$$

Note that $\tilde{c}^i(r)$ and $\tilde{c}^i(g)$ represent effective consumption depending on whether the red or green tree offers relatively higher cash flows in state 1, respectively, and that $\tilde{c}^i(2)$ is effective consumption in state 2. In the definition of effective consumption $\tilde{c}^i(\iota)$, the taste parameters $(\theta^i_j)_{j \in \mathcal{J}}$ represent agent i's private *tastes* over assets. Taste parameters allow us to nest characteristics-based demand as distinct from cash-flowed based risk-return considerations.

For simplicity, we assume that tastes are irrelevant for tree 2: $\theta^i_2 = 1$ for all i. Hence, tastes only affect relative preferences for red and green trees. Since Section 2.2 has shown that rank-preserving variation in the taste distribution can have independent effects on portfolio choice, for transparency we will mainly focus on the sparse specification $\theta^1_g = 1 + t$ and $\theta^1_r = 1 - t$, while $\theta^2_g = 1 - t$ and $\theta^2_r = 1 + t$. This means that type 1 prefers green while type 2 prefers red. However, this choice is not necessary.

Investors care only about consumption at date 1. Relative to Section 2, this simplifies matters in that variation in the level of non-pecuniary tastes cannot distort any consumption savings decision. Hence, we can focus on identifying cross-sectional asset pricing and portfolio choice effects. We work within the expected utility framework. In particular, we assume that investor preferences over state-contingent effective units of consumption are given by a CRRA utility function u. Our numerical examples use log utility.

Investors (or delegated managers) may also face different *mandates*, which are exogenous restrictions on permissible portfolios. We incorporate mandates because they have been argued to be useful for identification of demand systems (Koijen and Yogo, 2019). Since tastes only affect red and green trees, we define mandates over these assets as well. Given market prices p_g and p_r, define the portfolio share of green trees among red and green trees as

$$w^i_g = \frac{p_g a^i_g}{p_g a^i_g + p_r a^i_r}.$$

A mandate is a restriction that imposes, for parameters $\overline{w}_g^i$ and $\underline{w}_g^i$,

$$w_g^i \in [\underline{w}_g^i, \overline{w}_g^i].$$

Decision Problem. We normalize the price of tree 2 to $p_2 = 1$. The budget constraint is:

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

Substituting consumption in state 2, the decision problem of investor i is:

$$\begin{aligned} \max_{a_g^i, a_r^i \geq 0} \quad & \pi_1 \left[\rho u \left(\theta_g^i y_g(r) a_g^i + \theta_r^i y_r(r) a_r^i \right) + (1-\rho) u \left(\theta_g^i y_g(g) a_g^i + \theta_r^i y_r(g) a_r^i \right) \right] \\ & + \pi_2 u \left(e_2^i + p_g(e_g^i - a_g^i) + p_r(e_r^i - a_r^i) \right) \\ \text{s.t.} \quad & w_g^i \in [\underline{w}_g^i, \overline{w}_g^i]. \end{aligned} \tag{2}$$

Our equilibrium concept is a competitive equilibrium.

Definition 3 (Competitive Equilibrium) *A competitive equilibrium consists of asset prices* (p_g, p_r) *and portfolios* (a_g^i, a_r^i, a_2^i) *for each i such that:*

1. *Given asset prices, portfolios solve decision problem* (2) *for each i.*

2. *Markets clear for every asset:*

$$\sum_i a_j^i = \frac{1}{2} \text{ for } j \in \{r, g\} \quad \text{and} \quad \sum_i a_2^i = 1.$$

3. *The goods market clears.*

The rest of this section solves for equilibrium demand systems and prices when mandates do not bind for any investor. Section 3.1 provides the first-order conditions for optimal portfolios. As a benchmark, Section 3.2 solves for an equilibrium demand system without tastes. Finally, Section 3.3 solves for an equilibrium. We study implications of mandates and short-sales constraints in Section 4.

3.1 Optimal Portfolio Choice

It is instructive to begin with the case where mandates do not bind for any investor. In this case, demand functions are determined by the following first-order conditions:

$$\begin{aligned} a_g^i: &\quad \pi_1 \rho \theta_g^i (y(1) - \epsilon) u' \left(\frac{\tilde{c}^i(r)}{\tilde{c}^i(2)} \right) + \pi_1 (1-\rho) \theta_g^i (y(1) + \epsilon) u' \left(\frac{\tilde{c}^i(g)}{\tilde{c}^i(2)} \right) \leq (1 - \pi_1) p_g; \\ a_r^i: &\quad \pi_1 \rho \theta_r^i (y(1) + \epsilon) u' \left(\frac{\tilde{c}^i(r)}{\tilde{c}^i(2)} \right) + \pi_1 (1-\rho) \theta_r^i (y(1) - \epsilon) u' \left(\frac{\tilde{c}^i(g)}{\tilde{c}^i(2)} \right) \leq (1 - \pi_1) p_r. \end{aligned}$$

These conditions hold with equality whenever the investor chooses a positive quantity of the associated asset. Whether this is the case in equilibrium depends on the distribution of tastes.

This demand system is non-linear and exhibits complementarities: a change in the price of one asset alters the demand for all other assets. This is because, when $\epsilon > 0$, there is a diversification benefit to holding both red and green trees. As Berry and Haile (2021) point out, in settings with demand complementarities it is generally not enough to have a valid instrument for a particular price.

Asset demand is also sensitive to the *intensity* of tastes. By this we mean that variation in θ_j^i will drive changes in portfolios even when the ordinal preference ranking is preserved. This is not necessarily true in some discrete choice models of durable good purchases, where the outcome of interest is a binary choice. An implication is that the estimation of counterfactual asset demands generally requires identifying *cardinal* values of tastes. In the language of asset pricing, defining a stochastic discount factor requires incorporating the exact value of the marginal investor's tastes. The difficulty in identifying such a SDF is that tastes may be latent given equilibrium play.

3.2 Representative Agent Benchmark: No Tastes

To build intuition, we first solve the model without tastes: $\theta_j^i = 1$. As in Lucas (1978), this leads to the representative agent framework in which the representative agent holds the aggregate endowment in equilibrium, and is thus well-diversified in aggregate state 1. Given that total output in state 1 is constant, we can define by

$$p_1 \equiv \frac{p_g + p_r}{2}$$

the price of a sure claim on one unit of consumption in state 1. Then prices are determined by

$$p_1 = \frac{\pi_1}{1-\pi_1} y(2) \quad \text{and} \quad p_r - p_g = 2\epsilon \frac{\pi_1}{1-\pi_1} (2\rho - 1) \frac{y(2)}{y(1)}.$$

The prices of claims on aggregate states reflect the relative scarcity of aggregate consumption across the two states, and price differences between red and green assets are driven by the distribution over the distributional shock ρ. With symmetric aggregate states (i.e., $y(1) = y(2) = 1$ and $\pi_1 = \frac{1}{2}$), this yields

$$p_1 = 1 \quad \text{and} \quad p_r = 1 + (2\rho - 1)\epsilon.$$

3.3 Equilibrium with Tastes: Endogenous Sorting

We now consider the case with tastes. Because different investors may disagree on the marginal value of investing in a particular asset, there may be endogenous sorting in equilibrium. By this, we mean that investors with a taste for green assets will hold only green assets, while those with a taste for red assets will hold only red assets. (Of course, both will hold tree 2 as well.)

We guess and verify that type 1 specializes in green assets, and vice versa. Then type 1's consumption in state 1 is $\theta_g^1 y_g(\iota) a_g^1$, and vice versa for type 2. By the first-order conditions for optimal portfolios, demand functions are

$$\begin{aligned} \text{Type 1:} \quad a_g^1 &= \frac{1}{p_g} \frac{\pi_1}{1-\pi_1} a_2^1; \\ \text{Type 2:} \quad a_r^2 &= \frac{1}{p_r} \frac{\pi_1}{1-\pi_1} a_2^2; \end{aligned}$$

and consider only the trade-off between tree 2 and one specific color. Observe that these demand functions are independent of the particular intensity of tastes: since the equilibrium features sorting, only the ordinal ranking of tastes matters.

While this is reminiscent of consumer good settings (for example, discrete choice over automobiles), sparse portfolio choices may be particularly problematic in financial markets. In particular, under the cardinal interpretation of preferences required for expected utility framework, the intensity of tastes will affect portfolio choice when investors hold both assets in equilibrium. Since taste intensities are latent on the equilibrium path when sorting occurs (in particular, they can only be identified up to their ordinal properties), counterfactuals are vulnerable to incorrect inference about taste intensities.

To illustrate this issue, we verify whether sorting can be sustained in equilibrium. Evaluating the first-order condition for a_r^1 at $a_r^1 = 0$, it is indeed optimal for type 1 to refrain from purchasing red trees if and only if

$$\frac{p_r}{p_g} \geq \frac{\theta_r^1}{\theta_g^1} \left[\rho \frac{y(1) + \epsilon}{y(1) - \epsilon} + (1 - \rho) \frac{y(1) - \epsilon}{y(1) + \epsilon} \right].$$

This inequality states that the relative price of red trees must be high enough relative to the relative taste for red trees, adjusted by the benefits of diversification within state 1. In the equilibrium with sorting, moreover, relative prices are driven by wealth shares:

$$\frac{p_r}{p_g} = \frac{1 - \omega}{\omega}.$$

Hence, shocks to the wealth distribution or the degree of complementarity of red and green trees (as determined by ϵ and ρ) can lead sorting to break down. Shocks to wealth are precisely the type of counterfactual entertained by Koijen, Richmond, and Yogo (Forthcoming).

We now demonstrate that equilibrium demand functions do indeed depend on the intensity of taste parameters once there is partial sorting. We obtain simple closed-form solutions for the case where trees are perfect substitutes, i.e., when $\epsilon = 0$.

Proposition 3 (Equilibrium with and without sorting) *Let $\theta_r^1 < \theta_g^1$ and assume $\epsilon = 0$. There exists a threshold $\overline{\omega}$ for type 1's wealth share ω such that type 1 buys only green trees if $\omega \leq \overline{\omega}$ and buys both red and green trees if $\omega > \overline{\omega}$. When this is the case, prices are determined by type 1's taste parameters, $p_g = \theta_g^1$ and $p_r = \theta_r^1$, and the quantity of red trees held by type 2 is*

$$a_r^2 = \frac{1 - \omega}{2} \frac{\frac{1}{2} E_2 + \theta_g^1 E_g + \theta_r^1 E_r}{\theta_r^1},$$

where $E_2 = 1$ and $E_g = E_r = \frac{1}{2}$ denote the aggregate endowments of each tree.

The result highlights that, conditional on a shock to the wealth distribution, prices and quantities are now determined by the *intensity* of type 1's tastes, not just their ordinal ranking. Since these are latent conditional on equilibrium play, it is difficult to conduct counterfactuals based on observational data that feature sparse portfolios. While we illustrate this in a setting with only two types, the insight naturally generalizes to many types. In this case, equilibrium portfolios do not reveal preferences of infra-marginal investors. Yet, identifying preference parameters of inframarginal investor is critical for any counterfactual in which investors may rebalance their portfolios on the extensive margin.

Figure 2 shows equilibrium prices for the entire range of wealth share ω and substitutability ϵ. The left panel shows the green price p_g, and the right panel shows the price of a sure claim on state 1, $p_1 = \frac{p_g + p_r}{2}$. In the left panel, sorting occurs in the linear region near the origin, but breaks down as either type 1 becomes too wealthy or the diversification benefits become too large. In response to shocks, the relative price of green trees (and thus the underlying demand system) is highly non-linear in fundamentals. In contrast, the aggregate price of state 1 is flat in the entire region, as in the representative agent benchmark. This is because, *conditional* on a well-specified stochastic discount factor that takes into account tastes, every investor remains well-diversified within state 1.

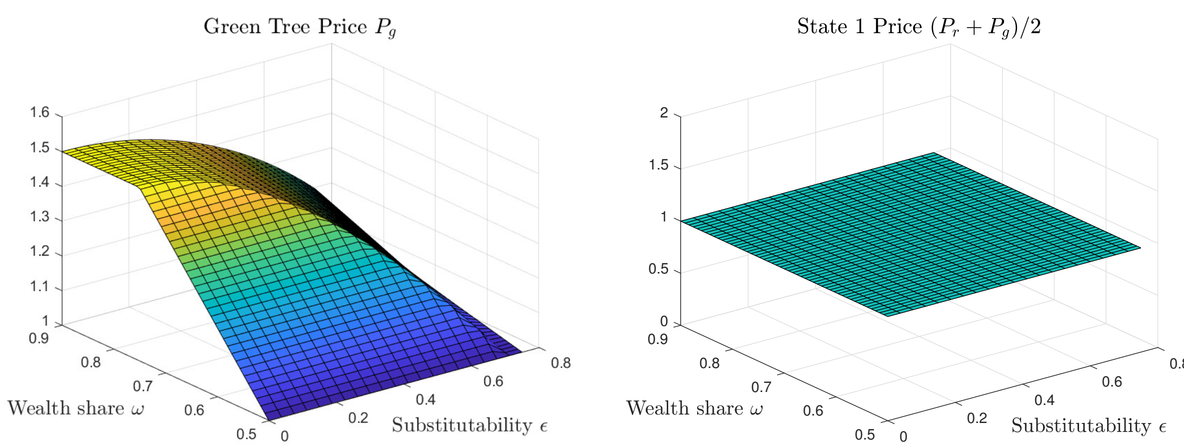

Figure 2: Green Price (Left) and State 1 Price $\frac{p_g+p_r}{2}$ (Right).

4 Identification and Counterfactuals

Having analyzed equilibrium demand systems without mandates binding in Section 3, we study the implications for identification and counterfactuals when mandates or short-sale constraints are present. We also discuss endogenous wealth effects.

4.1 On Identification and Counterfactuals with Mandates

In addition to tastes, investors may be distinguished by their mandates. In practice, this means that certain funds only invest in S&P 500 companies, or only in companies that have high Environmental, Social, and Governance (ESG) scores. While mandates may be a contributing factor to observed portfolio choices, they have also been put forth as helpful for *identifying* demand systems (Koijen and Yogo, 2019). As the argument goes, an asset that is inside the "investment universe" of many investors will be in higher demand, and thus see higher prices, than an otherwise similar asset that is not widely held in many investment universes. Putting aside for a minute the concern that mandates may be chosen in response to investment opportunities, we can consider the implications of this identification strategy in our model.

In particular, say there is a share m of type 1 investors that are not permitted to invest in red trees. While this is a very stark mandate, it is the type of mandate that would be ideal for the identification strategy in Koijen, Richmond, and Yogo (Forthcoming) because it is entirely inflexible with respect to changes in investment opportunities. Mandates are observed by investors, but not by the econometrician.

The demand function of *mandate investors* (superscript M) trades off green trees with tree 2. Under log utility, they spend share π_1 of their wealth on green trees:

$$a_g^M = \pi_1 \cdot \frac{e_2^M + p_g e_g^M + p_r e_r^M}{p_g}.$$

Mandates and tastes for green assets are observationally equivalent to the extent that the equilibrium features sorting; that is, tastes and mandates are both valid microfoundations for sparse portfolios. However, they differ when non-mandate investors choose to hold both red and green assets. This threatens the validity of counterfactuals. In particular, shocks to the wealth of type 1 will result in different counterfactual prices when many investors are subject to mandates versus when they are not.

Proposition 4 *The equilibrium response to shocks to wealth share ω or substitutability ϵ may be qualitatively different depending on the share of mandate investors m.*

Figure 3 illustrates this result. When there are almost no investors with mandates, a shock to ϵ creates more demand for diversification. Thus, the price of green trees is

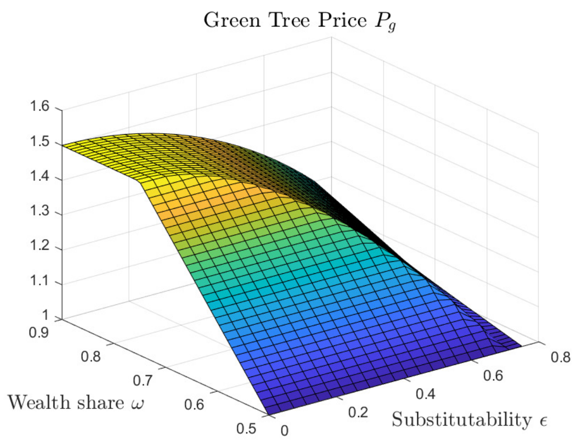

Figure 3: Green Price. Left: Low Mandate Share ($m \approx 0$). Right: High Mandate Share ($m = 0.85$).

decreasing in ϵ if type 1 investors choose to hold both types of trees. Mandate investors do not buy red trees at any price. Hence, shocks to ϵ do not reduce their demand for green trees even as type 2's demand increases. Thus, the price of green trees may be increasing in ϵ when there are sufficiently many mandate investors. The observational equivalence of tastes and mandates therefore creates the risk that counterfactuals are misspecified.

More broadly, identification based on mandates is threatened by the lack of a theory of delegation. In practice, mandates are typically imposed on funds (such as mutual funds), not end investors (such as households). This means that even very tight mandates are irrelevant as long as end investors can flexibly reallocate investments across funds.

Proposition 5 *Consider a two-layer structure where households invest through funds, and funds are subject to mandates. Suppose further that there exist at least one red and one green fund. Absent other frictions, equilibrium is invariant in mandates.*

In practice, researchers have pointed out that households may be slow to rebalance, or do so in predictable manners at regular intervals (i.e., quarter end). While it is possible that this may help with identification, the argument is incomplete: if some investors are known to rebalance intermittently, other investors may trade preemptively only to later sell. In this sense, intermittent rebalancers' tastes may be reflected in market demand even when they are not actively trading.

4.2 Equilibrium Consequences of Short Sales

As discussed in Section 2.3, the neoclassical arbitrage pricing may fail to hold when investors differ in terms of their tastes. We now illustrate consequences of this fact by introducing a set of investors who can freely sell short any asset. Such investors are likely to have outsize implications for equilibrium prices. We show this mechanism under the assumption that red and green trees are perfect substitutes in terms of their cash flows, $\epsilon = 0$. In this case, optimal portfolio choices are bang-bang, and the investor takes an infinite short position whenever the relative price of red and green trees is misaligned with her tastes.

Proposition 6 *Let $\epsilon = 0$. We enrich the model with a single investor type, indexed by S, who can freely short. Then the relative price of red and green trees is given by*

$$\frac{p_g}{p_r} = \frac{\theta_g^S}{\theta_r^S}$$

and is independent of any other parameters in the model.

The proposition states that an agent who can freely short trades until relative prices are aligned with her tastes *irrespective* of any other parameters. In practice, there may be short-sale constraints or other limits to arbitrage that prevent large short positions. However, this merely means that researchers have to measure *when* they might bind.

Two considerations make this difficult in practice: researchers may not observe short positions, nor do they have universal coverage of all investors in a given market. These data limitations make it difficult to infer taste parameters from equilibrium play.

4.3 Endogenous Wealth Effects

We now discuss another important feature of financial markets—portfolios are regularly marked to market. This means that, even holding preferences fixed, an individual who already owns a particular stock will exhibit different demand elasticities in response to a price change than an investor who does not. Hence standard instruments that may work well in consumer good settings (where purchases are one shot) will not be sufficient to identify asset demand systems.

This simple logic also has implications for the case of "index deletions."[12] In particular, assume that an investor is mandated to hold only assets that are in a particular index. Assume that green and red were initially the index, before red surprisingly drops out. Hence, the investor must divest upon deletion. We call the short-run demand curve the one that determines demand right upon deletion, and the long-run demand curve the one that obtains once short-run adjustments have occurred. Assume for illustrative purposes that the investor is not forward-looking, and that $\epsilon = 0$. Then demand functions satisfy

$$\text{Short-run demand:} \quad a_g^M = \pi_1 \cdot \frac{e_2^M + p_g e_g^M + p_r e_r^M}{p_g};$$

$$\text{Long-run demand:} \quad a_g^M = \pi_1 \cdot \frac{e_2^M + p_g e_g^M}{p_g}.$$

The difference is due to the valuation of endowments. In the short run, changes in the red price affect demand because wealth is marked to market. In the long run, demand is independent of the red price because the investor was forced to divest. In general, there are thus important dynamic considerations that differ from consumer markets, where most purchases are generally not resold or marked to market.

5 On The Structural Interpretation of Demand Elasticities

Demand elasticities are one of the main objects of interest in industrial organization. The reason is that a well-identified demand elasticity which can be related to, e.g., preferences for automobiles may inform a policymaker of the quantity response to a tax policy that raises automobile prices. In line with this view, researchers in demand-system asset pricing often argue that demand elasticities are a useful diagnostic that might distinguish their method from more neoclassical approaches. Against this background, we now discuss the structural interpretability of demand elasticities in financial markets.

An important difference between consumer goods and financial markets is that portfolio choice is generally modeled, at least in part, using preferences over state-contingent payoffs rather than asset characteristics alone. We will therefore argue that *asset-level* de-

[12] See, for instance, Chang, Hong, and Liskovich (2015) and Pavlova and Sikorskaya (2023) for recent approaches to estimating demand elasticities from index deletions/additions.

mand elasticities may not be informative about preference parameters whenever multiple (portfolios) of assets can deliver the same state-contingent payoff stream. In particular, the structural interpretation of demand elasticities depends on the security menu, as well as on whether there are "outside options" for an investor to pursue in response to price changes. This in turn is linked to the general equilibrium consequences of price changes.

5.1 Elasticities and the Security Menu

To establish a clean benchmark, we first show an example where asset-level demand elasticities *are* informative about preferences. In particular, we assume that the security menu consists *only* of the full set of Arrow securities, and that investors do not exhibit tastes over assets. Since each asset is uniquely tied to a particular state, asset-demand elasticities are then informative about state-contingent valuations. In particular, consider a generic investor choosing Arrow security positions $(a(z))_{z\in\mathcal{Z}}$ to solve:

$$\begin{aligned} \max_{(a(z))_z} \quad & (1-\beta)u(c_0) + \beta \sum_{z\in\mathcal{Z}} \pi_z u(c_1(z)) \\ \text{s.t.} \quad & c_0 = w_0 - \sum_{z\in\mathcal{Z}} p(z)a(z) \\ & c_1(z) = a(z) + w_1(z), \end{aligned}$$

where $p(z)$ is the price of security z and $w_1(z)$ an exogenous state-contingent endowment. If we define the marginal rate of substitution (or the state price) associated with state z to be

$$\Lambda(z) = \frac{\beta\pi_z u'(c_1(z))}{(1-\beta)u'(c_0)},$$

then the first-order condition is

$$p(z) = \Lambda(z).$$

The implicit function theorem yields an equation linking demand elasticities to preference parameters,

$$\epsilon(z)\Lambda(z)\left[-\frac{\alpha(c_1(z))}{p(z)} - \alpha(c_0)\right] = \frac{1}{a(z)} + \Lambda(z)\alpha(c_0),$$

where $\alpha(c)$ is the coefficient of absolute risk aversion at c and $\epsilon(z) = \frac{\partial a(z)}{\partial p(z)}\frac{p(z)}{a(z)}$ is the price elasticity of demand for Arrow security z. Thus, when the security menu is the case of Arrow securities, suitable price instruments allow researchers to estimate preference parameters $\Lambda(z)\alpha(c_0)$ from portfolio data. With multiple price instruments, it may also be possible to disentangle both components under mild parametric assumptions.

Next, consider the effects of changes in the security menu. To remove direct effects on prices, we hold the set of marketable payoffs fixed. In particular, we begin with the Arrow security menu and, for some generic state of the world z^*, introduce a new Arrow security that also pays off only in state z^*. Denote the demand for the original security by $a_0(z^*)$, and the demand for the new security by $a_1(z^*)$. If the Law of One Price holds, then it must be the case that both assets have the same initial price, $p_0(z^*) = p_1(z^*)$, and the investor will be indifferent between holding both assets.

Proposition 7 (Elasticities with Redundant Assets) *Suppose that the investor holds a positive position in both securities referencing state z^*. Now consider an exogenous increase in price of the new security, holding all other prices fixed. Then the demand elasticity for the new security is $-\infty$, while the demand for consumption in state z^* is unchanged. Hence, estimated demand elasticities for the new security are uninformative about preference parameters.*

While this example is deliberately stark, it is sufficient to highlight the critical role of the security menu, and redundant assets, for the structural interpretability of demand elasticities. Outside the case of Arrow securities, moreover, investors may need to combine multiple assets in certain proportions to achieve a certain consumption stream, and asset-level demand elasticities may only have structural interpretations when multiple assets are considered jointly. Finally, when there are redundant assets, investor quantities may not be uniquely pinned down by optimality conditions, yet they will still affect measured elasticities. It is precisely because of this indeterminacy that neoclassical asset pricing has relied so extensively on no arbitrage relationships between prices.

5.2 Outside Goods

Next, we consider the role of outside options, and how they are shaped by general equilibrium forces. In industrial organization, it has long been recognized that the interpretation of demand elasticities depends critically on the assumed notion of an "outside good," which is the alternative use of money available to a consumer that, for instance, chooses not to buy a car (Berry and Haile, 2021). In consumer good settings, a common strategy is to model preferences as quasi-linear, with the implicit understanding that the outside good is a consumption bundle whose utility scales approximately linearly with wealth.

Such an approach may be less viable in financial markets, at least as long as one is interested in *general equilibrium* economies. To see this, return to our baseline Lucas Tree economy from Section 3 and assume that the equilibrium features perfect sorting, with only type 1 investors buying green trees and only type 2 investors buying red trees. In this case, as in Section 3.3, type 1's demand function for green trees satisfies:

$$a_g^1 = \frac{1}{p_g}\frac{\pi_1}{1-\pi_1}a_2^1.$$

Differentiating this expression with respect to the price p_g yields

$$\begin{aligned}\frac{\partial a_g^1}{\partial p_g} &= -\frac{1}{(p_g)^2}\frac{\pi_1}{1-\pi_1}a_2^1 + \frac{1}{p_g}\frac{\pi_1}{1-\pi_1}\frac{\partial a_2^1}{\partial p_g} \\ &= -\frac{a_g^1}{p_g} + \frac{a_g^1}{a_2^1}\frac{\partial a_2^1}{\partial p_g}.\end{aligned}$$

The first term is the own price effect that is negative. The second is the cross-demand elasticity of asset 2. As is well known, under natural assumptions, the overall response is negative, with investors substituting away from green trees as their price increases.

However, this cannot be true in equilibrium for any *fundamental* shock that might increase the price of green trees. Consider for example an increase in type 1's wealth share ω. Our preceding analysis shows that this must lead to an increase in p_g. However, since type 1 agents are symmetric within type and only type 1 investors buy green trees, market clearing ensures that $a_g^1 = \frac{1}{2}$ in equilibrium. Hence, the *equilibrium* elasticity is zero, not negative, because there is no outside good for investors to move to.

In applied contexts, it is therefore critical to assess the degree of substitutability between "inside" and "outside" assets. For example, corporate bonds may be a better substitute for Treasury bonds than equities, and the degree of substitutability may differ depending on the level of aggregation.

6 Conclusion

We present a synthesis between neoclassical asset pricing and recent demand-system approaches to asset pricing, using multiple methods of incorporating non-pecuniary tastes into equilibrium models of portfolio choice. Our analysis highlights important conceptual concerns, including the definition of no arbitrage, the pricing of redundant assets, and the cardinal interpretation of taste parameters. Based on these concerns, we highlight several barriers to identification of demand systems that may threaten the validity of counterfactuals based on estimated demand systems.

References

BERRY, S. T., AND P. A. HAILE (2021): "Foundations of Demand Estimation," *Handbook of Industrial Organization*, 4(1), 1–62. [Cited on pages 2, 4, 19, 25, and 34.]

BRETSCHER, L., L. SCHMID, I. SEN, AND V. SHARMA (2022): "Institutional Corporate Bond Pricing," *Working Paper*. [Cited on page 6.]

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 32.]

COCHRANE, J. H. (2005): *Asset Pricing: Revised Edition*. Princeton University Press. [Cited on page 15.]

D'AMICO, S., J. KLAUSMANN, AND A. A. PANCOST (2023): "The Benchmark Greenium," *Working Paper*. [Cited on page 7.]

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 6.]

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 20.]

HONG, H., AND J. C. STEIN (2007): "Disagreement and the Stock Market," *Journal of Economic Perspectives*, 21(2), 109–128. [Cited on page 16.]

JIANG, Z., R. RICHMOND, AND T. ZHANG (2023): "A Portfolio Approach to Global Imbalances," *NYU Stern School of Business Working Paper*. [Cited on page 6.]

KOIJEN, R. S. J., R. J. RICHMOND, AND M. YOGO (Forthcoming): "Which Investors Matter for Equity Valuations and Expected Returns?," *Review of Economic Stdies*. [Cited on pages 1, 4, 6, 21, 27, and 29.]

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, 4, 6, 10, 15, 17, 19, 23, and 29.]

LUCAS, R. E. J. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–1445. [Cited on pages 1, 3, 5, 21, and 25.]

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.]

——— (2022): "Dissecting green returns," *Journal of Financial Economics*, 146(2), 403–424. [Cited on page 6.]

PAVLOVA, A., AND T. SIKORSKAYA (2023): "Benchmarking Intensity," *Review of Financial Studies*, 36(3), 859–903. [Cited on page 32.]

ROSS, S. A. (2004): *Neoclassical Finance*. Princeton University Press, Princeton, NJ. [Cited on pages 1 and 14.]

SIDRAUSKI, M. (1967): "Rational Choice and Patterns of Growth in a Monetary Economy," *American Economic Review*, 57(2), 534–544. [Cited on page 12.]

STARKS, L. T. (2023): "Presidential Address: Sustainable Finance and ESG Issues - Value versus Values," *Journal of Finance*, 78(4), 1837–1872. [Cited on pages 6 and 9.]