Demand Elasticity is Not Well-Defined for Long-Short Portfolios*

Yu An

March 19, 2025

Abstract

Demand elasticity measures the percentage change in a product's demand when its price increases by 1%. This paper shows that demand elasticity is not well-defined for long-short portfolios. Prior studies aggregate portfolio demand by taking a weighted average of each asset's percentage demand change (with weights equal to the asset's portfolio weight). This heuristic approach lacks a formal theoretical foundation. This paper proposes a theoretical principle for aggregating demand and demonstrates that it yields an alternative procedure consistent with "portfolio betas" practice in risk management. Under this procedure, the notion of "percentage changes in demand for a long-short portfolio" in the elasticity definition is not economically meaningful, given that such portfolios have net-zero supply. Finally, the paper discusses alternative elasticity measures for long-short portfolios, and highlights the importance of incorporating risk exposure into elasticity measurement.

Keywords: demand elasticity, long-short portfolios, aggregation, quantity, risk

JEL Codes: G11, G12

*Carey Business School, Johns Hopkins University; yua@jhu.edu. For helpful comments, I thank Jiacui Li and Zeyu Zheng. All errors are my own.

1 Introduction

Demand elasticity is defined as $(\Delta Q/Q)/(\Delta P/P)$, measuring the percentage change in a product's demanded quantity when its price increases by 1%. Thanks to the pioneering works of Koijen and Yogo (2019) and Gabaix and Koijen (2022), researchers are now interested in empirically estimating and theoretically explaining demand elasticity for financial assets. In asset pricing, long-short portfolios (a.k.a., factors) are fundamental objects, as they characterize the arbitrage relationship among assets (Markowitz, 1952; Ross, 1976). One famous example is the Fama and French (1993) small-minus-big portfolio, which goes long \$1 in small companies and short \$1 in big companies. Given the importance of long-short portfolios in asset pricing, a series of works has attempted to estimate and interpret demand elasticity for them.

This paper shows that demand elasticity is not well-defined for long-short portfolios. Prior studies aggregate portfolio demand by taking a weighted average of each asset's percentage demand change (with weights equal to the asset's portfolio weight). This heuristic approach lacks a formal theoretical foundation. This paper proposes a theoretical principle for aggregating demand and demonstrates that it yields an alternative procedure consistent with "portfolio betas" practice in risk management. Under this procedure, the notion of "percentage changes in demand for a long-short portfolio" in the elasticity definition is not economically meaningful, given that such portfolios have net-zero supply. Finally, the paper discusses alternative elasticity measures for long-short portfolios, and highlights the importance of incorporating risk exposure into elasticity measurement.

Consider a setting where researchers have access to price and quantity data for individual assets and want to estimate the demand elasticity for some long-short portfolio. To do this, they need to aggregate both the quantity change $(\Delta Q/Q)$ and the price change $(\Delta P/P)$ at the portfolio level. (Demand instruments are also needed for the estimation. This issue is orthogonal to the portfolio aggregation issue considered in this paper.) The price aggregation

is theoretically grounded in the Law of One Price (LOOP)—researchers sum each asset's $(\Delta P/P)$ according to its portfolio weight. However, there is no clear theoretical guidance on aggregating quantities. In practice, researchers have analogously summed each asset's $(\Delta Q/Q)$ weighted by its portfolio weight. Yet, no analogous "law" exists to theoretically justify this aggregation.

Hence, to have a theory-grounded discussion on demand elasticity for long-short portfolios, one must first develop a theory for aggregating portfolio quantities. This paper proposes such a theoretical principle—much like LOOP underlies portfolio-price aggregation—and then derives an alternative formula for portfolio-quantity aggregation. This formula is consistent with "portfolio betas" practice in risk management.

To understand the new principle, recall that canonical asset pricing builds on risks, and long-short portfolios capture distinct dimensions of risk. Hence, if one wants to include quantities (demand) in the asset pricing framework, one must account for how changes in quantity affect risk. This accounting should be the same whether one views risk from individual assets' perspective or from that of aggregated portfolios. For example, suppose some noise traders demand \$100 of NVIDIA stock and \$200 of META stock from some marginal investors. Because these stocks are risky, the extra demand shifts the marginal investors' risk exposures. One can measure these risk exposures either from the viewpoint of individual stocks or from the viewpoint of the risk-factor portfolios (such as market, company size, growth potential, and AI) to which these stocks are exposed. The risk-matching principle requires these two perspectives to match.

This paper proves that this principle, together with the LOOP, implies that one should aggregate a portfolio's quantity change ΔQ by summing each asset's ΔQ according to its *return beta with respect to the portfolio*. The same formula applies to the aggregation of total quantity Q. This result holds for any portfolio, not just long-short ones. Compared to the literature's approach, there are two key differences. First, rather than aggregating $(\Delta Q/Q)$ in one step, the new formula aggregates ΔQ and Q separately. The intuition is

that a 1% demand shock to a big company (e.g., NVIDIA) poses far greater risk than a 1% shock to a small company. Hence, when summing across assets, one should use the absolute dollar amounts ΔQ and Q rather than the percentage term $(\Delta Q/Q)$. Second, instead of summing by portfolio weights, the formula sums by each asset's beta exposure to the portfolio. This approach aligns with "portfolio betas" practice in risk management. For example, buying \$100 of a stock whose market beta is 2 is effectively buying \$200 of the market portfolio in terms of market risk exposure. Quantitatively, the difference from the literature's formula is most pronounced when the portfolio's long and short legs have vastly different market capitalization, or when the assets within each leg vary significantly in their betas—conditions especially relevant for the Fama-French small-minus-big portfolio.

With the new portfolio theory, one may ask: "Can demand elasticity $(\Delta Q/Q)/(\Delta P/P)$ be defined using the corrected aggregation formula?" The answer is yes for the market or any long-only portfolio, but no for long-short portfolios in general. The reason is that long-short portfolios are generally constructed to be uncorrelated with the market portfolio, whose weights are proportional to each asset's total quantity Q. As a result, if one sums up the total quantities Q_n of each asset weighted by their beta to the long-short portfolio, the sum is zero. In other words, *the total quantity Q of such a long-short portfolio is zero, meaning that this portfolio has net-zero supply.* Mathematically, division by zero is undefined in $\Delta Q/Q$; economically, the concept of "percentage change in demand for a long-short portfolio" in the elasticity definition does not apply.

Given that $(\Delta Q/Q)/(\Delta P/P)$ is not well-defined for long-short portfolios, one might ask whether alternative elasticity measures exist. One option is the semi-elasticity $\Delta Q/(\Delta P/P)$, often used in industrial organization to measure the absolute dollar change in a portfolio's demand when its price rises by 1%. Although intuitive, this measure is not scale-invariant and thus not directly comparable across different portfolios. For example, if a portfolio is levered by a factor of 2 (i.e., each asset's weight is doubled), the percentage price change $(\Delta P/P)$ doubles, but the quantity change ΔQ is halved (because each asset's beta to the levered

portfolio is halved due to the portfolio's increased volatility). Consequently, $\Delta Q/(\Delta P/P)$ drops to one-fourth of its original value. Because the semi-elasticity can be arbitrarily altered by scaling, it does not permit an apples-to-apples comparison across portfolios.

An (2024) and An and Huber (2025) propose a risk-based elasticity $\Delta Q \times \text{var}(r)/(\Delta P/P)$, where the portfolio's return variance $\text{var}(r)$ measures its risk. This elasticity thus quantifies the change in the portfolio's *demanded risk*—rather than just its dollar quantity—when its price increases by 1%. Importantly, the risk-based elasticity is theoretically equal to the reciprocal of the marginal investors' per-capita risk-aversion parameter, and captures the risk-return tradeoff conditional on quantity change. It is also scale-invariant: if a portfolio is levered by a factor of 2, its return variance quadruples, keeping the elasticity unchanged. Hence, comparing this measure across different portfolios provides a consistent, apples-to-apples view of marginal risk-bearing capacity. Section 5 illustrates this elasticity with a simple model.

Related literature. This paper speaks directly to the empirical studies that estimate demand elasticity for long-short portfolios (e.g., Fama-French size and value factors) or for long-only portfolios (e.g., Fama-French 5×5 double-sorted portfolios on size and value) using mutual fund or retail investor demand shocks (Ben-David, Li, Rossi, and Song, 2022; Li, 2022; Peng and Wang, 2024; Li and Lin, 2025). This paper does not claim those studies are incorrect. On the contrary, they provide valuable data and novel demand instruments for the burgeoning literature on demand-based asset pricing. The only caution is that their estimates should not be interpreted as the demand elasticity of the long-short portfolio itself. Rather, those estimates capture a composite demand response of the underlying assets in the portfolio (and are perfectly valid as empirical moments), rather than a single elasticity parameter for the portfolio.

This paper contributes to the literature that clarifies the theoretical underpinnings of empirical estimates in demand-based asset pricing, including Fuchs, Fukuda, and Neuhann (2024) and Haddad, He, Huebner, Kondor, and Loualiche (2025). Both studies note that

because investors hold diversified portfolios, a demand shock in one asset can lead to rebalancing in others—a cross-substitution pattern unique to financial markets. This phenomenon complicates the theoretical interpretation of existing demand elasticity estimates. This paper emphasizes that long-short portfolios are specifically designed to capture substitution across different risk dimensions, making their demand effects critical for asset pricing. However, unlike the market portfolio or individual assets—where the conventional measure $(\Delta Q/Q)/(\Delta P/P)$ can still apply—long-short portfolios have net-zero supply, rendering the standard elasticity ill-defined. Consequently, as discussed, one must use alternative measures (such as semi-elasticities or risk-based elasticities) to obtain sound empirical estimates and theoretical interpretations.

In the remainder of this paper, Section 2 discusses how the literature defines demand elasticity for long-short portfolios. Section 3 provides the portfolio theory for aggregating quantity, and Section 4 applies this theory to show why the percentage-demand measure does not apply to long-short portfolios. Section 5 presents alternative measures of elasticity for long-short portfolios. Finally, Section 6 concludes.

2 How the Literature Defines Demand Elasticity for Long-Short Portfolios

This section introduces the model and explains how the literature defines demand elasticity for long-short portfolios.

Consider a two-period model with times $t = 0$ and $t = 1$, and a gross risk-free rate R_F. There are N assets, each with a total quantity of Q_n dollars for $n = 1, 2, \ldots, N$. For a stock, Q_n typically represents its market capitalization. The time-0 price of asset n is P_n, and at time 1, it liquidates for X_n. Suppose some noise traders demand ΔQ_n dollars of asset n from some marginal investors at time 0, causing the price of asset n to change from P_n to $P_n + \Delta P_n$. Here Q_n and ΔQ_n are measured in dollar terms at the pre-shock price P_n, consistent with standard empirical practices.

Next, consider a portfolio where the weight of asset n is w_n. That is, investing \$1 in the

portfolio allocates $\$w_n$ to asset n. If $\sum_{n=1}^{N} w_n \neq 1$, any leftover funds are invested in (or borrowed from) the risk-free asset. The goal is to measure portfolio-level demand elasticity, defined as:

$$\frac{\Delta Q^{\text{port}}/Q^{\text{port}}}{\Delta P^{\text{port}}/P^{\text{port}}}, \tag{1}$$

which represents the percentage change in the portfolio's demanded quantity for a 1% price increase.

To evaluate (1), one must aggregate both the price change and the quantity change at the portfolio level. For the price change, the literature adopts

$$\frac{\Delta P^{\text{port}}}{P^{\text{port}}} := w_1 \frac{\Delta P_1}{P_1} + w_2 \frac{\Delta P_2}{P_2} + \cdots + w_N \frac{\Delta P_N}{P_N}, \tag{2}$$

which is supported by the Law of One Price (LOOP).[1]

The main challenge is that there is no clear theoretical basis for aggregating quantity changes. In practice, prior studies use equation (3), computing a weighted average of each asset's percentage demand change $\Delta Q_n/Q_n$ using the portfolio weights w_n:

$$\frac{\Delta Q^{\text{port}}}{Q^{\text{port}}} := w_1 \frac{\Delta Q_1}{Q_1} + w_2 \frac{\Delta Q_2}{Q_2} + \cdots + w_N \frac{\Delta Q_N}{Q_N}. \tag{3}$$

However, unlike the price aggregation (justified by LOOP), there is no analogous principle to support this quantity aggregation.

Under the literature's aggregation, the portfolio-level demand elasticity reflects a composite demand response of the underlying assets $(\Delta Q_n/Q_n)/(\Delta P_n/P_n)$, with both numerator and denominator weighted by the portfolio weights w_n:

$$\frac{\Delta Q^{\text{port}}/Q^{\text{port}}}{\Delta P^{\text{port}}/P^{\text{port}}} := \frac{w_1 \frac{\Delta Q_1}{Q_1} + w_2 \frac{\Delta Q_2}{Q_2} + \cdots + w_N \frac{\Delta Q_N}{Q_N}}{w_1 \frac{\Delta P_1}{P_1} + w_2 \frac{\Delta P_2}{P_2} + \cdots + w_N \frac{\Delta P_N}{P_N}}. \tag{4}$$

[1] Specifically, the portfolio invests $\$w_n$ in asset n, which changes in price from P_n to $P_n + \Delta P_n$. The remaining balance, $1 - \sum_{n=1}^{N} w_n$, is placed in the risk-free account, whose price remains unchanged. Thus, $P^{\text{port}} = 1$ and $\Delta P^{\text{port}} = \sum_{n=1}^{N} w_n \Delta P_n / P_n$.

As the paper shows below, such an approach lacks a formal theoretical basis and may misinterpret the portfolio's true elasticity, especially for long-short portfolios. To establish a theoretically grounded discussion of demand elasticity for long-short portfolios, the next Section 3 develops a portfolio theory for aggregating ΔQ^{port} and Q^{port}. Section 4 then shows that, once this aggregation formula is in place, $\Delta Q^{\text{port}}/Q^{\text{port}}$ is no longer meaningful for long-short portfolios, since $Q^{\text{port}} = 0$ (the portfolio has net-zero supply).

3 Portfolio Theory for Aggregating Quantity

This section develops a general theory for aggregating the portfolio quantities ΔQ^{port} and Q^{port}. First, it introduces a theoretical principle for quantity aggregation; then, it applies this principle to derive the corresponding formula. For simplicity, the focus is on ΔQ^{port}, but the same approach applies to Q^{port} by replacing ΔQ^{port} with Q^{port}.

Define the excess return of asset n from $t = 0$ to $t = 1$ as $r_n = X_n/P_n - R_F$. When marginal investors accommodate an additional demand ΔQ_n for asset n, they assume an extra return of $\Delta Q_n r_n$. Summing across all N assets, the extra return from the quantity change is

$$\Delta Q_1 r_1 + \Delta Q_2 r_2 + \cdots + \Delta Q_N r_N. \tag{5}$$

Now suppose one forms $K \leq N$ portfolios, allowing for potential dimension reduction if some assets are redundant. Let $w_{n,k}$ be the weight of asset n in portfolio k. By the Law of One Price (LOOP), the excess return of portfolio k is given by

$$r_k^{\text{port}} = \sum_{n=1}^{N} w_{n,k} r_n. \tag{6}$$

Denote the quantity change of portfolio k by ΔQ_k^{port}, which must be determined. Summing across all K portfolios, the extra return from the quantity change is

$$\Delta Q_1^{\text{port}} r_1^{\text{port}} + \Delta Q_2^{\text{port}} r_2^{\text{port}} + \cdots + \Delta Q_K^{\text{port}} r_K^{\text{port}}. \tag{7}$$

Both (5) and (7) measure the extra return from the quantity change, but from two different perspectives—individual assets versus portfolios. Since these are simply alternative accounting views, they should yield the same result. However, if the number of portfolios K is smaller than the number of assets N, one cannot expect a perfect match—only the best possible approximation. Thus, the principle of quantity aggregation is to select $\Delta Q_1^{\text{port}}, \ldots, \Delta Q_K^{\text{port}}$ to minimize the variance of the difference between the extra return measured at the asset level and risk measured at the portfolio level:

$$\min_{\Delta Q_1^{\text{port}},\ldots,\Delta Q_K^{\text{port}}} \text{var}\left(\Delta Q_1 r_1 + \cdots + \Delta Q_N r_N - \Delta Q_1^{\text{port}} r_1^{\text{port}} - \cdots - \Delta Q_K^{\text{port}} r_K^{\text{port}}\right). \tag{8}$$

This ensures that portfolio-level quantity changes align as closely as possible with the risk implications of individual asset-level changes.[2]

Solving this minimization problem yields the aggregation formula:

$$\Delta Q_k^{\text{port}} = \Delta Q_1 \beta_{1,k} + \Delta Q_2 \beta_{2,k} + \cdots + \Delta Q_N \beta_{N,k}, \tag{9}$$

where $\beta_{n,k}$ represents the beta of asset n with respect to portfolio k.[3]

Compared to the literature's aggregation formula in (3), the formula (9) introduces two key differences. First, instead of aggregating $(\Delta Q/Q)$ in a single step, it separates the aggregation of ΔQ and Q. The intuition is that a 1% demand shock to a big company carries far greater risk than a 1% shock to a small company. Thus, when summing across assets, one should use the absolute dollar amounts ΔQ and Q rather than the percentage term $(\Delta Q/Q)$.

Second, instead of using portfolio weights, the formula weights each asset's contribution by its return beta with respect to the portfolio. This aligns with the "portfolio betas" practice in risk management. For example, buying \$100 of a stock with a market beta of 2

[2]Minimizing variance corresponds to a standard OLS regression that includes a constant. Economically, this allows individual assets to have nonzero alphas with respect to the K portfolios.

[3]Specifically, $\beta_{n,k}$ is obtained from the OLS regression: $\min_{\beta_{n,1},\ldots,\beta_{n,K}} \text{var}\left(r_n - \beta_{n,1} r_1^{\text{port}} - \cdots - \beta_{n,K} r_K^{\text{port}}\right)$.

is equivalent to buying \$200 of the market portfolio in terms of market risk exposure.

Having clarified the conceptual difference from the literature's formula (3), the next step is to assess the quantitative difference. To this end, consider a typical long-short portfolio that sorts stocks based on certain characteristics and then holds market-capitalization-weighted positions within its long and short legs, respectively.[4] For such a portfolio,

$$w_n = \begin{cases} \frac{Q_n}{\sum_{m\in\{\mathrm{L}\}} Q_m} & \text{if } n \in \{L\} \text{ (the long leg);} \\ -\frac{Q_n}{\sum_{m\in\{\mathrm{S}\}} Q_m} & \text{if } n \in \{S\} \text{ (the short leg);} \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

For this portfolio, the literature's formula (3) simplifies to

$$w_1 \frac{\Delta Q_1}{Q_1} + w_2 \frac{\Delta Q_2}{Q_2} + \cdots + w_N \frac{\Delta Q_N}{Q_N} = \frac{\sum_{n\in\{\mathrm{L}\}} \Delta Q_n}{\sum_{n\in\{\mathrm{L}\}} Q_n} - \frac{\sum_{n\in\{\mathrm{S}\}} \Delta Q_n}{\sum_{n\in\{\mathrm{S}\}} Q_n}. \tag{11}$$

Compared to the theory-founded formula (9), the literature's formula introduces two biases. First, instead of weighting the long and short legs equally in dollar terms, it weights them by their respective total market capitalization. This creates a bias because the long and short legs of most long-short portfolios are not designed to have equal market capitalization. The bias is particularly pronounced for the Fama-French small-minus-big (SMB) portfolio, which explicitly sorts stocks based on market capitalization.

Second, examining only the numerators, the literature's formula implicitly assumes that all assets in the long leg have a beta of 1 to the portfolio, all assets in the short leg have a beta of -1, and all assets outside the long and short legs have a beta of 0. This assumption introduces a bias when assets within each leg exhibit significant variation in their betas. Again, this bias is particularly pronounced for the SMB portfolio, as firm size follows a power-law distribution with extreme variation (Gabaix, 2009).

Having clarified the difference, a natural question for the theory-founded formula (9) is

[4] In practice, the literature sometimes conducts a double sort on two characteristics, as in the Fama-French portfolios.

whether portfolio weights can be used to establish a direct connection between portfolio-level and asset-level ΔQ, similar to how the literature does. The answer is yes, but not in the same way. Specifically, when $K = N$—meaning the number of portfolios equals the number of assets (i.e., no dimension reduction)—equation (9) implies

$$\Delta Q_n = w_{n,1}\Delta Q_1^{\text{port}} + w_{n,2}\Delta Q_2^{\text{port}} + \cdots + w_{n,N}\Delta Q_N^{\text{port}}. \tag{12}$$

This result follows from the fact that the portfolio-weight matrix and the beta matrix are inverses of each other, as first noted by Fama and MacBeth (1973) in their construction of mimicking portfolios for non-tradeable factors.[5]

Equation (12) shows that portfolio weights can indeed be used to relate asset-level and portfolio-level quantities, but in the *opposite* direction of the literature's formula (3). Instead of using portfolio weights to aggregate asset-level quantities into portfolio-level quantities, as in (3), equation (12) does the reverse—it starts from the portfolio level and distributes the change across individual assets. The interpretation is intuitive: the portfolio-level quantity change ΔQ_1^{port} is allocated to asset n according to $w_{n,1}$. Summing across all portfolios from 1 to N then recovers the total quantity change for asset n, as shown in equation (12).

4 Why the Percentage-Demand Measure Does Not Apply to Long-Short Portfolios

This section applies the demand aggregation formula to show that long-short portfolios have net-zero supply ($Q^{\text{port}} = 0$). Consequently, the percentage-demand measure $\Delta Q^{\text{port}}/Q^{\text{port}}$ in the standard elasticity formula is undefined due to division by zero.

[5]To see this, define the $K \times 1$ vector $\Delta\mathbf{Q}^{\text{port}} = (\Delta Q_1^{\text{port}}, \ldots, \Delta Q_K^{\text{port}})^\top$, the $N \times 1$ vectors $\Delta\mathbf{Q} = (\Delta Q_1, \ldots, \Delta Q_N)^\top$ and $\mathbf{r} = (r_1, \ldots, r_N)^\top$, and the $N \times K$ matrix $\mathbf{w} = \{w_{n,k}\}$. Then, equation (9) in vector form reads $\Delta\mathbf{Q}^{\text{port}} = \text{var}(\mathbf{w}^\top\mathbf{r})^{-1}\text{cov}(\mathbf{w}^\top\mathbf{r}, \Delta\mathbf{Q}^\top\mathbf{r})$. In the special case $K = N$, this simplifies to $\Delta\mathbf{Q} = \mathbf{w}\Delta\mathbf{Q}^{\text{port}}$, which, elementwise, corresponds to (12).

Applying (9), the total quantity of any portfolio is given by

$$Q^{\text{port}} = Q_1\beta_{1,\text{port}} + Q_2\beta_{2,\text{port}} + \cdots + Q_N\beta_{N,\text{port}}. \tag{13}$$

Now consider the market portfolio, which invests in each asset in proportion to its total quantity Q_n. The weight of asset n in this portfolio is given by $w_{n,\text{mkt}} = Q_n/(Q_1 + Q_2 + \cdots + Q_N)$. Substituting this into (13) yields

$$\begin{aligned} Q^{\text{port}} &= (Q_1 + Q_2 + \cdots + Q_N)\left(w_{1,\text{mkt}}\beta_{1,\text{port}} + w_{2,\text{mkt}}\beta_{2,\text{port}} + \cdots + w_{N,\text{mkt}}\beta_{N,\text{port}}\right) \\ &= (Q_1 + Q_2 + \cdots + Q_N)\beta_{\text{mkt,port}}. \end{aligned} \tag{14}$$

Thus, Q^{port} equals the total market size times the beta of the market portfolio relative to the portfolio ($\beta_{\text{mkt,port}}$). For the market portfolio itself, this beta is 1 by definition, so $Q^{\text{mkt}} = Q_1 + Q_2 + \cdots + Q_N$, consistent with standard intuition.

By construction, long-short portfolios are typically designed to have zero correlation with the market, implying $\beta_{\text{mkt,LS}} = 0$. Substituting this into (14) gives $Q^{\text{LS}} = 0$, confirming that these portfolios have *net-zero supply*. Mathematically, division by zero is undefined in $\Delta Q^{\text{LS}}/Q^{\text{LS}}$; economically, the concept of "percentage change in demand for a long-short portfolio" does not apply.

One might argue that due to sampling noise, $\beta_{\text{mkt,LS}}$ is usually not exactly zero in empirical estimation, so it could still be used in the denominator. However, even if $\beta_{\text{mkt,LS}}$ is only approximately zero, the goal of most long-short portfolio designs is to keep it as close to zero as possible. Dividing by a term that is both noisy and near zero can introduce significant estimation biases.

One might argue that while the percentage change in demand is not well-defined at the aggregate level, individual investors can still have well-defined percentage demand. Suppose there are I investors, indexed by $i = 1, 2, \ldots, I$, each holding Q_i^{LS} dollars of the long-short

portfolio. Then,

$$Q_1^{\mathrm{LS}} + Q_2^{\mathrm{LS}} + \cdots + Q_I^{\mathrm{LS}} = Q^{\mathrm{LS}} = 0. \tag{15}$$

Although Q^{LS} is zero in aggregate, each individual typically holds a nonzero amount, so one can still define $\Delta Q_i^{\mathrm{LS}}/Q_i^{\mathrm{LS}}$ at the investor level.

However, this approach can be misleading. Since long-short portfolios have net-zero supply, if some investors go long, others must go short. Consequently, a 1% change in $\Delta Q_i^{\mathrm{LS}}/Q_i^{\mathrm{LS}}$ has different implications depending on whether Q_i^{LS} is positive or negative. For example, if $Q_i^{\mathrm{LS}} > 0$ and $\Delta Q_i^{\mathrm{LS}}/Q_i^{\mathrm{LS}} = 1\%$, the investor's long position increases by 1%. However, if $Q_i^{\mathrm{LS}} < 0$, the same 1% change actually reflects a decrease in demand. To see why, consider moving from -100 to -101:

$$\frac{-101 - (-100)}{-100} = \frac{-1}{-100} = 1\%. \tag{16}$$

Despite the fraction being 1%, moving from -100 to -101 clearly represents a larger short position, reducing overall demand. Thus, a "1% increase" in a short position actually corresponds to a decrease in the investor's demand. In general, the sign of Q_i^{LS} (long versus short) reverses the meaning of a percentage change in demand. Simply put, percentage changes in individual holdings are not directly comparable across investors because the long versus short sign fundamentally changes their economic meaning.

5 Alternative Measures of Elasticity for Long-Short Portfolios

Having shown that $(\Delta Q/Q)/(\Delta P/P)$ is ill-defined for long-short portfolios, this section explores alternative elasticity measures. One option is the semi-elasticity $\Delta Q/(\Delta P/P)$, often used in industrial organization to measure the absolute dollar change in a portfolio's demand when its price rises by 1%. Although intuitive, this measure is not scale-invariant and thus not directly comparable across different portfolios. For example, if a portfolio is levered by a factor of 2 (i.e., each asset's weight is doubled), the percentage price change $(\Delta P/P)$

doubles, but the quantity change ΔQ is halved (because each asset's beta to the levered portfolio is halved due to the portfolio's increased volatility). Consequently, $\Delta Q/(\Delta P/P)$ drops to one-fourth of its original value. Because the semi-elasticity can be arbitrarily altered by scaling, it does not permit an apples-to-apples comparison across portfolios.

The portfolio theory for aggregating demand, developed earlier, is based on general first principles and does not rely on specific equilibrium assumptions. To derive an elasticity measure that reflects deeper economic forces, this section introduces additional equilibrium assumptions standard in the literature. Specifically, assume that marginal investors are competitive with mass μ, have CARA utility with risk-aversion parameter γ, and that asset returns $r_1, \ldots, r_n$ are jointly normally distributed.

When noise traders demand ΔQ_n dollars of asset $n = 1, \ldots, N$, prices adjust from P_n to $P_n + \Delta P_n$ to induce marginal investors to sell $\Delta Q_n/\mu$ dollars of asset n, thereby clearing the market. Thus, the marginal investor's optimality condition satisfies

$$\begin{aligned} &\{-\Delta Q_1/\mu, \ldots, -\Delta Q_N/\mu\} \\ &= \arg \max_{\{y_1,\ldots,y_N\}} \mathbb{E}\left[-\exp\left(-\gamma \sum_{n=1}^{N}\left(y_n(r_n - \frac{\Delta P_n}{P_n}R_F) + S_n(r_n + R_F)\right)\right)\right]. \end{aligned} \tag{17}$$

Each marginal investor freely chooses y_n (the amount to buy) to maximize utility. The term $y_n r_n$ represents the additional payoff risk taken, where r_n is the excess return from Section 3. When the asset's time-0 price changes from P_n to $P_n + \Delta P_n$, the total purchase cost is no longer just y_n but instead $y_n + y_n \Delta P_n/P_n$. The term $R_F y_n \Delta P_n/P_n$ accounts for this adjustment and compounds it to time 1. Finally, S_n denotes the amount of asset n that each marginal investor held before the shock, which earns $r_n + R_F$ per dollar at time 1.

Applying the first-order condition of (17), one derives the price impact equation:[6]

$$\frac{\Delta P_n}{P_n} = \lambda \left[\text{cov}(r_n, r_1)\Delta Q_1 + \text{cov}(r_n, r_2)\Delta Q_2 + \cdots + \text{cov}(r_n, r_N)\Delta Q_N \right], \tag{18}$$

where $\lambda := \gamma/(\mu R_F)$ approximates the per-capita risk aversion of marginal investors, a fundamental economic object.[7] Equation (18) represents the price impact of a demand shock: the percentage price change of asset n is a linear combination of demand shocks to all assets, weighted by their return covariances with asset n. This is a standard result in the theoretical literature (Rostek and Yoon, 2023).

Now consider N portfolios whose returns are uncorrelated, i.e.,

$$\text{cov}\left(\frac{\Delta P_k^{\text{port}}}{P_k^{\text{port}}}, \frac{\Delta P_j^{\text{port}}}{P_j^{\text{port}}} \right) = 0 \text{ for all } k \neq j. \tag{19}$$

Focusing on uncorrelated portfolios eliminates cross-factor interactions. In turn, each univariate beta with respect to a given portfolio remains the same even when controlling for other portfolios.

For any portfolio k, let $w_{n,k}$ denote the weight of asset n in that portfolio. Then,

$$\begin{aligned}
\frac{\Delta P_k^{\text{port}}}{P_k^{\text{port}}} &= w_{1,k}\frac{\Delta P_1}{P_1} + \cdots + w_{N,k}\frac{\Delta P_N}{P_N} && (20)\\
&= \lambda \left[\text{cov}\left(\sum_{n=1}^{N} w_{n,k} r_n, r_1 \right) \Delta Q_1 + \cdots + \text{cov}\left(\sum_{n=1}^{N} w_{n,k} r_n, r_N \right) \Delta Q_N \right] && (21)\\
&= \lambda \left[\text{cov}\left(r_k^{\text{port}}, r_1 \right) \Delta Q_1 + \cdots + \text{cov}\left(r_k^{\text{port}}, r_N \right) \Delta Q_N \right] && (22)\\
&= \lambda \text{var}\left(r_k^{\text{port}} \right) \left[\beta_{1,k}\Delta Q_1 + \cdots + \beta_{N,k}\Delta Q_N \right] && (23)\\
&= \lambda \text{var}\left(r_k^{\text{port}} \right) \Delta Q_k^{\text{port}}. && (24)
\end{aligned}$$

[6] To see this, define the $N \times 1$ vectors $\Delta\mathbf{Q} = (\Delta Q_1, \ldots, \Delta Q_N)^\top$, $\mathbf{r} = (r_1, \ldots, r_N)^\top$, $\mathbf{S} = (S_1, \ldots, S_N)^\top$, and $\Delta\mathbf{p} = (\Delta P_1/P_1, \ldots, \Delta P_N/P_N)^\top$. The first-order condition of (17) then takes the form $R_F\Delta\mathbf{p} - \mathbb{E}[\mathbf{r}] = \gamma \text{var}(\mathbf{r})(\Delta\mathbf{Q}/\mu - \mathbf{S})$. Setting $\Delta\mathbf{Q} = \mathbf{0}$ yields $\Delta\mathbf{p} = \mathbf{0}$. Taking the difference then gives $\Delta\mathbf{p} = \text{var}(\mathbf{r})\Delta\mathbf{Q}\gamma/(R_F\mu)$, which elementwise corresponds to (18).

[7] The division by the gross risk-free rate R_F appears because the model is discrete. In a continuous-time setting, this division disappears.

Step (20) expresses the portfolio's price change in terms of individual asset changes (by the LOOP). Step (21) substitutes the asset-level price impact equation (18) from equilibrium conditions. Step (22) again applies the LOOP to relate the portfolio's excess return to individual asset excess returns. Step (23) uses the uncorrelated-factors assumption to ensure that univariate betas match multivariate betas. Finally, step (24) applies the aggregation formula (9) for portfolio quantity.

Equation (24) implies that

$$\frac{\Delta Q_k^{\text{port}} \text{var}(r_k^{\text{port}})}{\Delta P_k^{\text{port}} / P_k^{\text{port}}} = \frac{1}{\lambda}. \tag{25}$$

The left-hand side of (25) provides a risk-based elasticity for long-short portfolios. The term $\text{var}(r_k^{\text{port}})$ represents the portfolio's risk. In effect, this elasticity quantifies the change in the portfolio's *demanded risk*—rather than just its dollar quantity—when its price increases by 1%. As equation (25) shows, this elasticity equals $1/\lambda$, the reciprocal of the marginal investor's per-capita risk-aversion parameter. As a result, it is scale-invariant: if a portfolio is levered by a factor of 2, then $\text{var}(r_k^{\text{port}})$ quadruples, leaving the elasticity unchanged.

This risk-based elasticity is also related to the canonical price of risk in asset pricing. Specifically,

$$\frac{\Delta P_k^{\text{port}} / P_k^{\text{port}}}{\Delta Q_k^{\text{port}} \text{var}(r_k^{\text{port}})} = \frac{\text{change in return due to quantity shift}}{\text{change in risk due to quantity shift}}. \tag{26}$$

The canonical price of risk captures the unconditional risk–return tradeoff, whereas the risk-based elasticity quantifies the risk–return tradeoff *conditional on a change in quantity.*

Two remarks are in order. First, since portfolios k can be freely chosen from 1 to N, equation (25) implies that all portfolios must share the same elasticity, $1/\lambda$. This means that in the base model, one cannot distinguish, say, a "value factor elasticity" from a "size factor elasticity"—they are identical. This is a limitation of the model, not of the elasticity concept. This limitation arises from the standard theoretical assumption that marginal investors have CARA utility across all assets. More complex utility functions (e.g., CRRA or recursive utility) also imply this restriction: while they allow risk aversion to vary over time via state

variables, they remain mean–variance in the local cross-section.

However, modern asset pricing posits that different portfolios may have distinct risk–return tradeoffs, thus making the empirical search for an optimal portfolio meaningful (Fama and MacBeth, 1973). If one already accepts that different portfolios can have different unconditional risk–return tradeoffs, it is natural to allow these portfolios to also have different conditional risk–return tradeoffs when quantities change. Under CARA utility, $\lambda = \gamma/(\mu R_F)$, so distinct elasticities can arise if pricing kernels vary by risk dimension (implying different γ) or if marginal investors are segmented along different risk dimensions (implying different μ). Hence, comparing the estimated λ across different portfolios provides a consistent, apples-to-apples view of marginal risk-bearing capacity. An and Huber (2025) theoretically derive and empirically estimate such a flexible elasticity structure. Estimating these portfolio-specific elasticities requires time-series demand instruments for each portfolio, as also noted by Haddad, He, Huebner, Kondor, and Loualiche (2025).

Second, merely requiring uncorrelated returns does not uniquely determine a set of portfolios.[8] Under the CARA assumption (where all portfolios share the same elasticity), selecting any particular set of uncorrelated portfolios is inconsequential, as they all yield the same underlying elasticity structure.

However, if different portfolios are allowed to have different elasticities, the choice of uncorrelated portfolios affects the elasticity structure of the underlying assets. An illustrative analogy is that, for a perfect circle, choosing any two axes identifies the same shape. In contrast, for an ellipse, selecting different axes yields different ellipses. An (2024) shows that there is a unique set (up to scaling) of portfolios satisfying both uncorrelated returns (19) and uncorrelated demand:

$$\operatorname{cov}\left(\Delta Q_k^{\text{port}}, \Delta Q_j^{\text{port}}\right) = 0 \text{ for all } k \neq j, \tag{27}$$

[8]For any set of N portfolios with uncorrelated returns, there are $N(N-1)/2$ degrees of freedom to rotate to an alternative set of uncorrelated portfolios, even after factoring in scaling.

then examines the implications for asset-level elasticity and demonstrates that this choice of portfolios is uniquely optimal under suitable conditions.

6 Conclusion

This paper examines whether the standard demand elasticity measure applies to long-short portfolios and shows that it is not well-defined in this context. Prior studies aggregate portfolio demand by taking a weighted average of each asset's percentage demand change (with weights equal to the asset's portfolio weight). This heuristic approach lacks a formal theoretical foundation.

To address this issue, this paper develops a theoretical principle for demand aggregation and derives an alternative method consistent with the "portfolio betas" practice in risk management. A key implication of this approach is that the conventional notion of "percentage change in demand for a long-short portfolio" is not economically meaningful, as such portfolios have net-zero supply. Consequently, the standard elasticity measure, which relies on dividing demand changes by total portfolio quantity, leads to a division-by-zero problem.

To advance empirical work on demand-based asset pricing, this paper explores alternative elasticity measures that remain valid for long-short portfolios. It particularly emphasizes the importance of incorporating risk exposure into elasticity measurement. Since long-short portfolios are constructed to capture risk factors rather than absolute demand levels, a proper elasticity measure should reflect the sensitivity of demanded risk—not just quantity—to price changes. In this vein, a measure that accounts for risk—$\Delta Q \times \text{var}(r)/(\Delta P/P)$—remains both well-defined and meaningful for long-short portfolios. This insight offers a more robust approach to measuring demand elasticity in asset pricing and has potential implications for how researchers empirically assess the impact of shocks on priced risk factors.

References

An, Yu, 2024, An arbitrage foundation for demand effects in asset pricing, Working paper.

An, Yu, and Amy W. Huber, 2025, Demand propagation through traded risk factors, Working paper.

Ben-David, Itzhak, Jiacui Li, Andrea Rossi, and Yang Song, 2022, Ratings-driven demand and systematic price fluctuations, *Review of Financial Studies* 35, 2790–2838.

Fama, Eugene F, and Kenneth R French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F, and James D MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Fuchs, William, Satoshi Fukuda, and Daniel Neuhann, 2024, Demand-system asset pricing: Theoretical foundations, Working paper.

Gabaix, Xavier, 2009, Power laws in economics and finance, *Annual Review of Economics* 1, 255–293.

Gabaix, Xavier, and Ralph SJ Koijen, 2022, In search of the origins of financial fluctuations: The inelastic markets hypothesis, Working paper.

Haddad, Valentin, Zhiguo He, Paul Huebner, Pèter Kondor, and Erik Loualiche, 2025, Causal inference for asset pricing, Working Paper.

Koijen, Ralph SJ, and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Li, Jiacui, 2022, What drives the size and value factors? *Review of Asset Pricing Studies* 12, 845–885.

Li, Jiacui, and Zihan Lin, 2025, Price multipliers are larger at more aggregate levels, Working paper.

Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance* 7, 77–91.

Peng, Cameron, and Chen Wang, 2024, Factor rebalancing, Working paper.

Ross, Stephen A, 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–60.

Rostek, Marzena J, and Ji Hee Yoon, 2023, Imperfect competition in financial markets: Recent developments, Working paper.