Short versus Long-run Demand Elasticities in Asset Pricing*

Philippe van der Beck†

February 1, 2025

Abstract

The recent literature on demand-system asset pricing estimates the slope of investors' demand curves from static logit regressions, implying counterfactual experiments have constant impacts that do not revert over time. Using investors' trades at different horizons, I provide reduced-form evidence that elasticities increase significantly in the long run. To capture these dynamics structurally, I propose a partial adjustment model differentiating between short- and long-run elasticities. I find that the price impact from counterfactuals is four times larger at quarterly horizons than in the long-run equilibrium. The model accounts for investor inertia, captures long-run reversal, and provides insights into return predictability.

*I am grateful to Pierre Collin-Dufresne, Lorenzo Bretscher, Sam Hanson, Adi Sunderam, Semyon Malamud, Jean-Philippe Bouchaud, Paul Huebner, as well as the participants at the AFA 2023, EFA 2024, CFM Workshop on Market Impact, HEC Lausanne PhD Workshop, Toulouse School of Economics, IESE Business School, Universitat Pompeu Fabra, Vrije Universiteit Amsterdam, Warwick Business School, University of Oxford Saïd Business School, Imperial College Business School, HEC Paris, Frankfurt School of Finance, Duke University, Carnegie Mellon University, Ohio State University, Boston University, Harvard Business School, University of Michigan, University of Toronto, Tuck School of Business, NYU Stern, University of Cologne, University of Hamburg, World Finance Conference Turin, and the SFI Research Days for helpful comments and valuable discussions. This paper supersedes a previous version titled "On the Estimation of Demand-based Asset Pricing Models". The research was conducted under financial support from the Swiss National Science Foundation (SNSF).

†pvanderbeck@hbs.edu, Harvard Business School

1 Introduction

Detailed data on investor portfolio holdings has renewed interest in how investor heterogeneity, especially individual demand shocks, affects asset prices. If the aggregate demand curve for a financial security is downward sloping, investors' demand shocks can have a meaningful impact on equilibrium prices. The Demand System Approach to Asset Pricing developed by Koijen and Yogo (henceforth KY, 2019) allows incorporating holdings data into equilibrium asset pricing by estimating investor-specific demand curves. Investors' elasticity of demand with respect to the price (henceforth elasticity) lies at the heart of the demand system approach. It determines the equilibrium price impact of counterfactual experiments within the demand system. If the elasticity of demand is high, then equilibrium prices do not have to move a lot in order to induce investors to rebalance their portfolios, and the price impact of demand shocks is low.

While the idea that elasticities are horizon-dependent is pervasive throughout economics, the literature on asset demand systems has not yet confronted this issue. The standard approach in the demand-system asset pricing literature is to infer static elasticities from logit demand, thereby abstracting from potential horizon effects in the counterfactuals. A natural next step in this literature is to account for horizon-effects and differentiate between short- and long-run equilibrium counterfactuals. This paper addresses this gap, showing — both structurally and through a difference-in-differences estimator – that investors become significantly more elastic over longer horizons. Consistent across both estimations, I find that on average, investors' long-run elasticities are four times larger than their quarterly short-run elasticities. This implies that the impact of counterfactual experiments is four times larger at a quarterly horizon than in the long-term equilibrium. The patterns are in line with the slow-moving capital hypothesis (see Duffie (2010)) and imply that the price impact of counterfactual experiments slowly dissipates over time.

I set the stage by providing reduced-form evidence on increasing elasticities at longer horizons. I compute investors' trades, i.e. changes in their portfolio holdings, and estimate their elasticities over increasing trading horizons in a simple difference-in-difference setting. I find elasticities that increase monotonically at longer horizons. Investors trade 3 to 4 times more elastically over the long run than in the short run.

The reduced-form estimates are closely linked to the standard logit regressions in the literature. The status quo is to estimate elasticities via cross-sectional regressions of portfolio holdings onto instrumented prices. The underlying intuition is that an investor

is more price-elastic in equilibrium if they overweight assets with lower prices. The cross-sectional identification from holdings does not take into account how investors *trade*, i.e., how they adjust their portfolios over time. The static model therefore does not speak to the horizon of investors' elasticities. The reduced-form estimates in this paper are first-difference versions of the static logit specification, with differences computed over varying trading horizons ranging from 1 to 16 quarters. The static logit elasticity identified from portfolio holdings in levels falls between the short- and long-run elasticities identified from investors' trades. This suggests that static logit demand overstates how aggressively investors rebalance their portfolio in the short-run, and understates their rebalancing activity in the long-run.

Motivated by the reduced-form evidence, I propose a simple dynamic extension to KY (2019) that allows differentiating between short- and long-run effects in counterfactual experiments. I micro-found a partial adjustment model following Gârleanu and Pedersen (2013) in which investors gradually shift towards a target portfolio due to trading costs. The optimal portfolio is a weighted average between the passive buy-and-hold portfolio and the target portfolio. The weight on the passive portfolio captures inertia, i.e., persistence in investor demand. I provide an instrumental variables estimator to address the endogeneity of past portfolios and current latent demand arising from persistent demand shocks.

In line with the reduced-form evidence on lower elasticities at short horizons, I find high inertia in institutional portfolios. The average investor's weight on the quarterly buy-and-hold portfolio is 0.84. This implies that investors partially trade towards their long-term target with a quarterly speed of 0.16. Furthermore, consistent with higher trading costs for bigger positions, large investors are more inert than small investors. The short- and long-run elasticities implied by the partial adjustment model are highly consistent with the reduced-form estimates. Slow portfolio adjustments lead to implied long-run elasticities that are around four times higher than in the short-run (quarterly). Over long horizons, large value investors are the most important liquidity providers with a long-run elasticity of 5.2. All results are robust to different instruments for lagged portfolios.

Elasticities that *increase* at longer horizons imply counterfactual price impacts that *decrease* over time. I illustrates this duality in a generalized setting that nests many workhorse asset pricing frameworks. In particular, I show that the equilibrium price effects of counterfactual experiments within a demand system can be approximated as

the product of a demand shock and a multiplier matrix that is the inverse of the aggregate price elasticity of demand. If elasticities increase at longer horizons, the price multiplier decreases and the impact of counterfactuals dissipates over time. Correctly estimating elasticities and their dynamics therefore determines whether the demand-based approach produces realistic counterfactual price paths.

I test the short-term impact and long-run reversal of counterfactual experiments implied by the partial adjustment model. To ensure a true out-of-sample test, I use demand shocks that have not been used to identify the model parameters. Out-of-sample, the dynamic model accurately predicts both the immediate effects of demand shocks and their long-term reversal, both qualitatively and quantitatively. In contrast, the static framework predicts constant price impacts, underestimating the short-run (quarterly) impact while overstating long-run price effects.

Finally, in the partial adjustment model, long-run elasticities – and consequently, long-run price impacts – are not directly estimated but inferred structurally from short-run elasticities and portfolio inertia. This addresses a key challenge in estimating long-horizon elasticities: the absence of plausibly exogenous price instruments that remain relevant over long horizons. The model's long-run price multipliers provide a novel stock-level measure of long-term reversal. Stocks in the habitat of long-term investors with high inertia exhibit strong reversal as these investors gradually rebalance their portfolios and become more elastic over time. This underscores the distinction between short- and long-run dynamics driven by the interaction of inertia and elasticity. Empirically, I find that for stocks classified as reversal stocks, momentum returns disappear entirely, underscoring the incremental information embedded in investor portfolios and their rebalancing behavior.

Related Literature. The paper is related to three strands of literature. Most importantly, it directly relates to the recent and fast-growing literature on demand-based asset pricing initiated by KY (2019). This literature uses data on investors' portfolio holdings for different asset classes to estimate demand curves and, via the price elasticities of demand, link them to equilibrium prices (see Benetton and Compiani (2020), Koijen and Yogo (2020), Haddad et al. (2021), Han et al. (2021), Koijen et al. (2021), Fang et al. (2022), Coqueret (2022), Bretscher et al. (2024), Jiang et al. (2022), Darmouni et al. (2022), Tamoni et al. (2024), Jiang et al. (2024), Koijen et al. (2024), Jansen (2024), Jansen et al. (2024)). All of the papers mentioned above use the demand system approach with horizon-independent elasticities to conduct counterfactual experiments.

Accounting for long versus short-run elasticities can provide valuable additional insights on how the price impact of counterfactual experiments is expected to unfold over time and to what extent it reverts. The partial adjustment model proposed in this paper can be easily integrated into existing frameworks by controlling for lagged portfolios in the demand estimation. By effectively capturing static mandates that may be correlated with prices, this may furthermore sharpen the identification of existing studies. Huebner (2023) estimates a dynamic demand system and shows that if investors trade on momentum, short-term mispricing may worsen over time. The framework proposed in this paper allows separating long-run elasticities from momentum preferences and focuses on the former.

Second, the difference-in-difference estimator used to motivate the partial adjustment model is closely linked to the event-study evidence on price pressure from index reconstitutions (see e.g. Shleifer (1986), Wurgler and Zhuravskaya (2002), Chang et al. (2015), Pavlova and Sikorskaya (2022), Greenwood (2005)), mutual fund flows (see e.g. Edmans et al. (2012), Coval and Stafford (2007), Lou (2012), Wardlaw (2020), Schmickler (2020)) or dividend reinvestments (see e.g. Hartzmark and Solomon (2021) and Schmickler (2020)). All of the measures introduced in these papers can essentially be used as instruments to estimate reduced-form elasticities via the difference-in-difference estimator. This lends additional flexibility to the demand-based approach as the elasticity estimates can be ex-ante verified in a host of different settings.

Finally, this paper contributes to the literature on slow-moving capital (e.g., Mitchell et al. (2007); Duffie (2010); Bogousslavsky et al. (2021)) and gradual portfolio adjustment (e.g., Froot et al. (2001); Bohn and Tesar (1996); Bacchetta and Van Wincoop (2021); Bacchetta et al. (2022)). Following Bacchetta et al. (2023), I propose a partial adjustment model that can be estimated at the investor level, enabling the identification of both long-run and short-run elasticities from portfolio holdings data. Consistent with the theory on dynamic portfolio choice under transaction costs (e.g. Gârleanu and Pedersen (2013), Gârleanu and Pedersen (2016), Collin-Dufresne et al. (2020)), I find that larger investors with bigger positions trade more slowly than smaller investors.

The remainder of this paper is structured as follows. Section 2 provides some background on the role of dynamic demand elasticities in asset pricing. Section 3 describes the data. Section 4 estimates price elasticities in reduced form at different horizons. Section 5 proposes the partial adjustment model. Section 6 provides applications and tests the model out-of-sample. Section 7 concludes.

2 The Role of Demand Elasticities in Asset Pricing

Any asset pricing model is *demand-based* and implies an elasticity of demand with respect to the price that can in general be estimated from observable data on investors' portfolios. The increased availability of investor-level holdings data has revived the interest in estimating the slope of investors' demand curves and linking holdings to equilibrium prices. Neoclassical models generally imply high price elasticities of demand such that demand shocks have small equilibrium price effects.[1] Regardless of the precise magnitude of the elasticity, asset pricing models typically differ in their predictions about how this elasticity varies across horizons. Static models imply elasticities that are independent of the horizon.[2] Instead, dynamic models with slow-moving capital imply elasticities that increase over longer horizons.[3] These models offer tractable price dynamics following supply shocks, implying price impacts that revert over time as inert traders "wake up" or arbitrage/inventory risks decline. They typically solve for equilibrium prices directly and do not produce investor-specific demand curves in closed form. However, equilibrium price multipliers that *decline* over time are just the flip-side of demand elasticities that *increase* over time. Below I provide a generalized framework that highlights this duality.

A Portfolio Choice and Price Elasticities

Consider the following single-period model of portfolio choice. There are N stocks indexed by $n = 1, ..., N$. Shares outstanding of each stock are given by Q^*, normalized to 1 so that the endogenous total market equity of each stock is given by its price $P(n)$. There are I investors indexed by $i = 1, ..., I$ who choose an optimal number of shares $Q_i(n)$ in each stock. For now, I am assuming that the optimal portfolio $Q_i \in \mathbb{R}^N$ is some unknown fund-specific function $f_i(\cdot)$ of the vector of stock prices $P \in \mathbb{R}^N$ and a collection of other exogenous observable and unobservable demand shifters X (such as e.g. expected cash

[1] See the calibrations in Petajisto (2009), Gabaix and Koijen (2023), and Gabaix et al. (2024b). The negligible role of demand shocks in standard models is at odds with extensive empirical evidence from index inclusions mutual fund flows, index reconstitutions, benchmarking (Pavlova and Sikorskaya (2022)), and dividend reinvestments, suggesting that investors have downward-sloping (or inelastic) demand curves. See Wurgler and Zhuravskaya (2002), Lou (2012), Chang et al. (2015), and Hartzmark and Solomon (2021) for respective examples.

[2] See e.g. Wurgler and Zhuravskaya (2002), Petajisto (2009), Pavlova and Sikorskaya (2022), Kashyap et al. (2023), Koijen et al. (2024).

[3] See e.g. Grossman and Miller (1988), Greenwood (2005), Duffie (2010), Bogousslavsky (2016)).

flows, the assets under management, interest rate, risk or investment constraints):

$$Q_i = f_i\big(P, X\big) \tag{1}$$

This general specification allows plugging in a demand curve that matches observable data on portfolio choice. For example KY (2019) use an exponential function for $f_i(\cdot)$ whereas Balasubramaniam et al. (2021) and Betermier et al. (2022) opt for a linear specification. Micro-founding a demand curve that matches observable data is beyond the scope of this paper and is left for future research.[4] Importantly, $f_i(\cdot)$ pins down an investor's $N \times N$ elasticity matrix

$$\zeta_i = -\text{diag}(Q_i)^{-1}\frac{\partial Q_i}{\partial P'}\text{diag}(P) \tag{2}$$

which has own-price elasticities on the diagonal and cross-price elasticities $\zeta_{i,nm} = -\frac{\partial Q_i(n)}{\partial P(m)}\frac{P(m)}{Q_i(n)}$ on the off-diagonal. Appendix A.1 shows that for a CARA utility investor with risk aversion γ_i and a covariance matrix of future cashflows Σ the elasticity matrix is proportional to $(\gamma_i\Sigma)^{-1}$. Risk-averse arbitrageurs require larger price concessions to move away from stocks which contribute more to the risk of the arbitrage portfolio. Note, that in a static model the elasticity does not have a time horizon. It measures the equilibrium adjustment in quantities due to an exogenous price shock, ignoring the speed at which quantities adjust. Appendix A.2 provides a simple dynamic extension to this general framework that differentiates between long and short-run elasticities. Investor inertia can be modeled via a lagged holdings term in the optimal portfolio $Q_{i,t} = f_i\big(P_t, X_t, Q_{i,t-1}\big)$ which implies a long-run elasticity

$$\zeta_i^{LR} = \zeta_i\left(1 + \sum_{s=1}^{\infty}\left(\frac{\partial f_i}{\partial Q_{i,t-1}}\right)^s\right) \tag{3}$$

that exceeds the short-run elasticity ζ_i. In the special case where $f_i(\cdot)$ satisfies a partial adjustment model $Q_{i,t} = \rho_i Q_{i,t-1} + v_i(P_t, X_t)$, as in Gârleanu and Pedersen (2013), the long-horizon elasticity simplifies to $\zeta_i^{LR} = \frac{\zeta_i}{1-\rho_i}$.

B Market Clearing and the Impact of Demand Shocks

Market clearing implies that all funds must jointly hold the market portfolio $\sum_{i=1}^{I} f_i(\cdot) = Q^*$. Equilibrium prices are only available in closed form for very specific demand curves. For example, linear CARA-normal demand implies the CAPM (see Appendix A.1). For

[4] KY (2019) provide a suggestive micro-foundation of log-linear portfolio weights by using the fact that $w(n) = 1 + X(n) + \frac{1}{2!}X(n)^2 ... \approx e^{X(n)}$.

general demand curves - like the exponential demand specification in KY (2019) - the price impact of counterfactual experiments can be first-order approximated as the product of a demand shock and the market's aggregate price elasticity of demand. To this end, assume that there is an exogenous shock ΔX that leads to a change in aggregate investor demand of Δd (expressed in percent of shares outstanding).

PROPOSITION 1 *A first-order approximation of the equilibrium return Δp due to the demand shock Δd is given by*

$$\Delta p = \mathcal{M}\Delta d \tag{4}$$

where $\mathcal{M} \in \mathbb{R}^{N\times N}$ the inverse of the ownership-weighted average elasticity

$$\mathcal{M} = \left(\sum_{i=1}^{I} diag(Q_i)\zeta_i\right)^{-1}. \tag{5}$$

See Appendix A.3 for a proof. Proposition 1 states that the impact of demand shocks on the cross-section of realized returns is given by the market's aggregate elasticity of demand. $\mathcal{M}$ is the cross-sectional pendant to the scalar multiplier in Gabaix and Koijen (2023) and will henceforth be referred to as the multiplier matrix.[5] Ultimately, elasticities ζ_i and therefore the multiplier matrix $\mathcal{M}$ are pinned down by the specific choice of $f_i(\cdot)$. The CARA-normal setting implies that the multiplier matrix is proportional to $\gamma\Sigma$.

COROLLARY 1 *Let $p = g(A, X, \beta, \epsilon)$ denote the equilibrium vector of log prices as in the demand system by KY(2019, equation 22). Any counterfactual experiment involves a change in assets under management ΔA, characteristics ΔX, demand coefficients $\Delta\beta$ or latent demand $\Delta\epsilon$, and can be represented as an aggregate demand shift $\Delta d \in \mathbb{R}^N$ expressed in percent of shares outstanding. A first-order approximation of the equilibrium price impact of counterfactual experiments is given by equation (4).*

See Appendix A.4 for a proof. The demand-based approach by KY (2019) has been used in various settings to assess the impact of counterfactual experiments on equilibrium prices. Corollary 1 states, that counterfactual experiments can in general be restated in the form of an aggregate demand shock Δd.[6] The quantitative implications of counterfactuals

[5]The more price-elastic investors are (i.e., the larger the diagonal elements in ζ_i), the less prices have to move to accommodate the exogenous demand shock. Cross-price elasticities drive the off-diagonal elements in $\mathcal{M}$ and are responsible for spill-over effects to other stocks.

[6]For example, Koijen et al. (2024) assess the impact of a shift towards passive management by redistributing assets from active to passive funds. Han et al. (2021) evaluate the price impact of mutual

therefore depend on how elastically investors respond to exogenous price changes. Note, that the static framework compares the two equilibria before and after the shock and does not speak to the dynamics of price path following counterfactual experiments. Depending on the way the elasticities and therefore $\mathcal{M}$ are estimated, the price impact captures long versus short-run dynamics.

3 Data and Variable Construction

I obtain institution-level share holdings $S_{i,t}(n)$ from 1990 to 2024 from Thomson's Institutional Holdings Database (s34 file). The data on quarterly stock holdings are subsequently merged with price and fundamentals data from CRSP and Compustat.[7] Institution-level and mutual fund portfolio weights $w_{i,t}(n)$ are constructed as the dollar holdings in each stock (price $P_t(n)$ times shares held $Q_{i,t}(n)$) divided by their assets under management $A_{i,t}$. An institution's assets under management are given by the sum of its dollar holdings. The subscripts t indicate the report date of the 13F filing. Monthly data on mutual funds' net returns and total net assets, as well as other fund-specific characteristics, are obtained from the CRSP survivorship-bias-free mutual fund database. Fund flows relative to lagged assets are denoted as $f_{i,t} = \frac{F_{i,t}}{A_{i,t-1}}$. In line with the theory, I normalize by shares outstanding such that a stock's price and its market capitalization coincide. Thus the number of shares held is equivalent to the fraction of ownership $Q_{i,t}(n) = \frac{S_{i,t}(n)}{\text{Shrout}_t(n)}$, and stock splits do not contaminate changes in shares held from one quarter to the next. I define $\delta_{i,t}(n) \equiv \log \frac{w_{i,t}(n)}{w_{i,t}(0)}$ where $w_{i,t}(0)$ is the weight in the outside asset as in KY (2019). I denote log quantities and market equities using lowercase, i.e. $p_t(n) = \log P_t(n)$ and $q_{i,t}(n) = \log Q_{i,t}(n)$. Δ denotes quarterly changes. Changes beyond single quarters are denoted by Δ^h, where h is the horizon (in quarterly units) over which the change is computed. Percentage changes in prices and quantities are hence given by $\Delta^h p_t(n) = \log P_t(n) - \log P_{t-h}(n)$, $\Delta^h q_{i,t}(n) = \log Q_{i,t}(n) - \log Q_{i,t-h}(n)$ and $\Delta^h \delta_{i,t}(n) = \delta_{i,t}(n) - \delta_{i,t-h}(n)$. If the dollar holdings in the outside asset are constant over time, then $\Delta^h q_{i,t}(n) = \Delta^h \delta_{i,t}(n) - \Delta^h p_t(n)$.

fund risk shifting by redistributing underperforming funds' assets. Bretscher et al. (2024) conduct a battery of counterfactuals for the corporate bond market by changing the wealth distribution across investors A_i and bond characteristics X. Jiang et al. (2024) decompose the US net foreign asset position by iteratively changing its underlying determinants and computing the counterfactual equilibrium. The counterfactual price change is - at least to a first order - driven by $\mathcal{M}$.

[7] See KY (2019) for details on the construction of the database.

4 Reduced-Form Elasticities at Different Horizons

Before diving into the elasticity estimation it is worth noting that virtually all of the reduced-form evidence for inelastic demand is cast in the form of regressions of price changes onto some exogenous demand shift by a specific investor group, such as index demand by passive funds or flow-induced demand by mutual funds. By changing the horizon of the price change, the resulting coefficient captures short- versus long-run price multipliers. In this spirit, I propose a reduced-form estimator for investor-specific elasticities (the inverse of multipliers) using investors' trades at different horizons.

A A simple difference-in-difference estimator

The most general reduced-form estimator, that directly emerges from the definition of the elasticity matrix in equation (2), is evaluating the percentage trade $\Delta^h q_i(n)$ as a response to a percentage change in the price $\Delta^h p(n)$ of all other stocks over different horizons h.

$$\Delta^h q_i(n) = -\sum_{m=1}^{N} \zeta^h_{i,nm} \Delta^h p(m) + d_i(n) \tag{6}$$

I am omitting time t subscripts to avoid notational clutter. $d_i(n)$ captures demand shifters due to e.g. fundamental news, risk, flows, or regulatory constraints. A *causal* identification of elasticities requires *exogenous* variation in prices orthogonal to $d_i(n)$. For now, I will ignore the endogeneity problem (which is addressed in the next section). For trades obtained from quarterly 13F filings, the number of observations for each investor ($N \times T$) is too small to estimate each element of the $N \times N$ elasticities matrix individually. However, the issue can be addressed by parameterizing ζ^i_{nm} as a function of observable stock-specific characteristics. The problem then reduces to estimating interaction effects as in Kelly et al. (2019). For simplicity, I opt for the substitution patterns implied by logit demand (see KY, 2019), which implies homogenous substitution across stocks. There are good reasons as to why cross- and own-price elasticities may differ across stocks. For example, the simple mean-variance portfolio (Appendix A.1) implies heterogeneous cross-elasticities driven by covariances across stocks. Note, that the focus of this paper does not lie on capturing the specific substitution patterns across different stocks, but on assessing how investors' average elasticity changes at longer time horizons. Chaudhary et al. (2023) provide a detailed analysis of heterogenous substitution patterns in corporate bonds. I leave the analysis of horizon effects in cross-elasticities up to future research.

Formally, logit substitution implies

$$\zeta_{i,nm}^h = \begin{cases} \zeta_i^h(1 - w_i(n)), & m = n \\ -\zeta_i^h w_i(m), & m \neq n. \end{cases} \tag{7}$$

Cross-price elasticities are then subsumed by an investor-specific constant and (6) becomes a simple cross-sectional regression for each investor i.

$$\Delta^h q_i(n) = -\zeta_i^h \Delta^h p(n) + c_i^h + d_i(n) \tag{8}$$

where $c_i^h = -\zeta_i^h \sum_{n=1}^N w_i(n)\Delta^h p(n)$ is the scaled portfolio return over h periods, captured by a constant in cross-sectional regressions and by time-fixed effects in panel regressions. This is a simple difference-in-difference estimator of an investor's elasticity, that computes the change in demand for n (in excess of total portfolio demand) relative to the change in price of n (in excess of the total portfolio return). Technically, ζ_i^h measures the difference between own- and cross-price elasticity. However, for diversified portfolios, this difference is negligible as individual portfolio weights are small.

B Identification

Identifying how investors change their holdings, i.e., trade, as a response to changes in the prices, requires instruments for price changes, i.e., returns. To construct exogenous return variation at different frequencies, I follow Lou (2012) and use flow-induced trades by mutual funds.[8] To the extent that fund flows are correlated with the funds' respective portfolio tilts, flow-induced trades may contain fundamental information and violate the exclusion restriction. I, therefore, construct flow shocks $f_{i,t}^{\perp}$ by orthogonalizing flows with respect to fund-level characteristics. See Appendix B for details. Formally, flow-induced demand pressure for stock n is defined as $\Delta z_t(n) = \sum_{i=1}^I f_{i,t}^{\perp} Q_{i,t-1}(n)$.[9] For longer horizons $h > 1$ the flow-based instrument $\Delta^h z_t(n)$ is constructed by summing

[8]The reduced-form literature on price multipliers has proposed several alternative exogenous demand shifts that could be used to instrument for price changes, such as index reconstitutions (see e.g. Chang et al. (2015), Pavlova and Sikorskaya (2022)) and, more recently, signed order flow aggregated to lower frequencies (Li and Lin (2022)).

[9]Because the demand estimation is at the individual institution-level, the correct instrument excludes all flow-induced trades from mutual funds that belong institution i. Formally, $\Delta z_{i,t}(n) = \sum_{j\neq i}^I f_{j,t}^{\perp} Q_{j,t-1}(n)$. Empirically, however, this does not materially affect any of the estimates because total flow-driven trading is not strongly affected by individual funds (even aggregated at the manager level).

$\Delta z_t(n)$ over time. Investor-specific elasticities can then be obtained in a simple two-stage least squares procedure,

$$\begin{aligned} \text{1st Stage:} \quad \Delta^h p_t(n) &= \mathcal{M}^h \Delta^h z_t(n) + u_t(n) \\ \text{2nd Stage:} \quad \Delta^h q_{i,t}(n) &= -\zeta_i^h \Delta^h \hat{p}_t(n) + d_{i,t}(n) \end{aligned} \tag{9}$$

where $u_{i,t}(n)$ and $d_{i,t}(n)$ include the stock-specific control variables, log book equity, dividends to book equity, profitability, investment, and market beta as well as time-fixed effects capturing cross-elasticities. Because flow-induced trading is a weak instrument over the small subsets of stocks held by most investors, I estimate the first stage jointly across all stocks in a two-sample instrumental variables estimation (see Angrist and Krueger (1992), Huebner (2023), and Koijen et al. (2024)). Table 1 reports the first-stage coefficient on the flow-based instrument ($\Delta^h z_t$) across different horizons. Consistent with downward-sloping demand, flow-driven buying pressure by mutual funds is strongly correlated with contemporaneous returns but weakens at longer horizons, limiting its suitability for identifying elasticities beyond four years. Generally, finding strong instruments to estimate *long-run* elasticities is challenging: If a demand shock significantly affects returns over very long horizons, it probably contains fundamental information and therefore violates the exclusion restriction. However, the primary goal of this section is not to pin down investors' exact elasticities, but to provide simple reduced-form evidence of increasing elasticities at larger horizons. The structural model introduced in Section 5 addresses the challenge of weak long-run instruments by identifying horizon-dependent elasticities from portfolio inertia.

C Estimation Results

I estimate investor-specific elasticities over different horizons using the flow-based instrument 2000 to 2024 via two-stage least squares. I remove extensive trades (i.e. inclusions/exclusions from portfolios) and consider horizons h from 1 quarter up to up to 4 years. Due to limited observations from quarterly 13F filings and narrow holdings, elasticities are estimated over the panel of stock-level trades for investors with over 1,000 trades and over 3 years of data. Others, or those with weak first stages, are pooled into equal-sized bins by size. Table 1 reports the estimated coefficients across all horizons.

[Table 1 about here.]

The first column reports the AUM-weighted average elasticity across all investors. At

a quarterly horizon, the average elasticity is 0.58 and increases monotonically at longer horizons. At an annual frequency, the average elasticity reaches 0.72 and increases to 1.68 at the 4-year horizon. Similarly, the median elasticity across investors increases from 0.47 at a quarterly horizon to 2.5 at long horizons. At the 90th percentile among investors, long-run elasticities are over 6 times larger than their short-term counterparts. Unsurprisingly, the AUM-weighted average elasticity approximately equals the inverse of the first-stage coefficient. The first-stage coefficient $\mathcal{M}^h$ represents the average stock-specific multiplier from flow-induced trades. Consistent with the notion of increasing elasticities, $\mathcal{M}^h$ decreases monotonically over longer horizons.[10] Overall, the reduced-form results paint a consistent picture of elasticities increasing substantially over longer horizons, with investors exhibiting roughly 3 to 4 times greater responsiveness in the long run compared to the quarterly frequency.

D Relation to Structural Estimates from Logit Demand

In a conceptual deviation from the event-study evidence for inelastic demand, KY (2019) and all papers using their approach identify elasticities structurally using cross-sectional regression in *levels* as opposed to *changes*. Motivated by the fact that investors' portfolio weights are log-normally distributed in the cross-section, portfolio weights are modeled as $w_i(n) = \frac{\exp(\delta_i(n))}{1+\sum_n \exp(\delta_i(n))}$ where

$$\delta_i(n) = \log \frac{w_i(n)}{w_i(0)} = \beta_i p(n) + u_i(n) \tag{10}$$

is estimated using quarterly cross-sectional regressions for each investor. Subtracting $p(n)$ from both sides and noting that the weight in the outside asset is subsumed by a constant c_i in a cross-sectional regression, this can be rewritten as

$$q_i(n) = -\zeta_i p(n) + c_i + d_i(n) \tag{11}$$

[10] Hence, in a pure reduced-form setting, a decreasing first-stage price impact is sufficient to document higher average elasticities at longer horizons. Estimating investor-specific elasticities in a second stage nevertheless has two distinct benefits. First, it gives insights into the underlying investor-specific determinants of the multiplier. Second, one can obtain stock-specific elasticities and time-varying effects even by estimating a scalar elasticity for every investor. Because ownership varies across stocks and time, ownership-weighted sums across elasticities vary across stocks and time.

where $\zeta_i = 1 - \beta_i$.[11] The reduced-form estimator (8) for horizon h is hence a straightforward first-difference version of the level specification in KY (2019), with differences computed over h periods. The logit demand curve in levels implies high price elasticities for investors tilting towards cheaper (controlling for fundamental value) stocks in the cross-section. It estimates constant equilibrium elasticities that are independent of the horizon. When conducting counterfactual experiments, the price change implied by the estimated elasticities is an *equilibrium* response that does not have a time-series dimension. The bottom row of Table 1 reports the elasticities obtained from logit demand.[12] Flow-induced trades impact returns but are not significantly related to market equities in levels, even over long horizons. I, therefore, instrument endogenous market equity $p(n)$ with KY's mandate-based instrument.[13] The AUM-weighted average elasticity is 0.73. Despite using different estimators and instruments, the logit and reduced-form elasticities are both low and comparable in magnitude. This is another piece of evidence that institutional demand is indeed far more inelastic than implied by neoclassical frameworks (see Gabaix et al. (2024b)). Nevertheless, logit elasticities are static and capture an average horizon-independent effect that lies between the short- and long-run elasticity from the reduced-form estimation. Note, however, that the estimates are not directly comparable as they rely on different instruments. The partial adjustment model proposed below estimates both horizon-specific and static elasticities with the same instrument and confirms that the static logit estimate indeed lies between the short- and long-run elasticity.

What drives the difference between reduced-form elasticities, which use trades at varying horizons, and logit elasticities, which use log portfolio holdings in levels? The elasticities may diverge for two concrete reasons. First, trades introduce a specific time-series dimension that is not present in the static logit estimation. Second, trades effectively control for lagged portfolios and hence remove hidden investment mandates that may lead to biased elasticities (if correlated with market equity). Appendix Section C illustrates the two effects in a simulated laboratory environment. The structural model presented in the next section allows disentangling the two effects empirically.

[11] Note, that the budget constraint, $\sum_n w_i(n) = 1 - w_i(0)$, implies own-price elasticities are given by $-\frac{\partial q_i(n)}{\partial p(n)} = \zeta_i(1 - w_i(n))$ as opposed to ζ_i. However, since $w_i(n)$ is small empirically, this adjustment has a negligible impact on the implied elasticities.

[12] I use the same empirical setup as for the reduced-form specification, i.e., the same binning across investors, panel estimation, and pooled first stage.

[13] Formally, the mandate-based instrument is given by $z_{i,t}(n) = \sum_{j \neq i} A_{j,t} \frac{\mathbb{1}_{j,t}(n)}{1+\mathcal{N}_{j,t}}$ where $\mathcal{N}_{j,t}$ denotes the set of stocks in investor j's investment universe and $\mathbb{1}_{j,t}(n)$ is an indicator equal to 1 if $n \in \mathcal{N}_{j,t}$.

5 Dynamic Elasticities: A Structural Approach

Motivated by the horizon-dependent elasticities from the reduced-form estimator, I propose a simple dynamic extension to the logit specification in KY (2019). This allows incorporating horizon-effects in the counterfactual experiments of existing and future applications of the demand-system approach. Beyond improved counterfactuals, the structural approach has three other key benefits. First, long-run elasticities are not directly estimated but inferred from short-run elasticities combined with portfolio inertia. This mitigates a key challenge in estimating long-horizon elasticities: The lack of plausibly exogenous instruments that are relevant at long horizons. Second, the first-difference estimator can obscure valuable insights, particularly for slow-moving holdings. Intuitively, many long-run investors respond slowly to new investment opportunities. For example, a value investor may rebalance annually or gradually trade towards their aim portfolio. Over-differencing their slow-moving portfolio holdings may remove valuable long-run information contained in the data causing the reduced-form estimator to fit noise from quarters in which the institution only partially rebalances or does not trade at all. Last, the reduced-form estimator can only be applied to investors with a large time series of available trades.

A Partial Adjustment Model

Portfolio inertia can originate from both behavioral factors, such as cognitive constraints (Enke and Graeber (2023)), or structural factors, such as adjustment costs or regulatory constraints. In asset pricing, a natural micro-foundation for slow portfolio adjustments is transaction costs (Gârleanu and Pedersen (2013)). Let $\delta^*_{i,t}(n)$ denote the desired target portfolio the investor would trade towards instantaneously in the absence of adjustment costs. Following Koijen and Yogo (2019), I assume that the target portfolio is linear in log market equity $p_t(n)$ and a demand shifter $d_{i,t}(n)$. Formally,

$$\delta^*_{i,t}(n) = \beta_{i,t} p_t(n) + d_{i,t}(n) \tag{12}$$

As before, $d_{i,t}(n) = \sum_{k=1}^{K} X^k_t(n)\beta^k_{i,t} + \varepsilon_{i,t}(n)$ captures demand for log book equity, dividends to book equity, profitability, investment, and market beta as well as a constant that accounts for investor-time fixed effects. The desired portfolio contains the long-run elasticity $\zeta^{LR}_{i,t} = 1 - \beta_{i,t}$, which captures the extent to which the investor bets against deviations from fundamental value in the long run. Next, let $\delta^p_{i,t}(n) = \delta_{i,t-1}(n) + \Delta p_t(n)$

denote the passive buy-and-hold portfolio the investor would hold in the absence of active portfolio rebalancing between $t-1$ and t. For simplicity of notation, I assume no valuation gains in the outside asset.[14] Appendix A.5 shows that under quadratic trading costs that are proportional to risk, constrained mean-variance optimization yields a partial adjustment model: The optimal portfolio $\delta_{i,t} \in \mathbb{R}^N$ is a weighted mix of the desired long-run optimal portfolio $\delta^*_{i,t}$ and the passive buy-and-hold portfolio $\delta^p_{i,t}$,

$$\delta_{i,t} = \rho_{i,t}\delta^p_{i,t} + (1-\rho_{i,t})\delta^*_{i,t} \tag{13}$$

where the weight on the passive buy and hold portfolio $\rho_{i,t} \in (0,1]$ captures an investor's inertia, which is driven by portfolio adjustment costs and risk aversion.[15] Note, that one could conceive alternative microfoundations for (13) based on, for example, inattention.

B Estimation

The passive portfolio can be decomposed into lagged ownership relative to shares outstanding, $q_{i,t-1}(n)$, and current market equity. Formally, $\delta^p_{i,t}(n) = q_{i,t-1}(n) + p_t(n) - h_{i,t}$ where $h_{i,t} = \log(A_{i,t}w_{i,t}(0))$ are the log dollar holdings in the outside asset. Consequently, equation (13) can be rewritten as a simple dynamic panel model for every investor i,

$$\delta_{i,t}(n) = \underbrace{\rho_{i,t}q_{i,t-1}(n)}_{\text{Partial Adjust.}} + \alpha_{i,t}p_t(n) + u_{i,t}(n) \tag{14}$$

where $\alpha_{i,t} = (1-\rho_{i,t})\beta_{i,t} + \rho_{i,t}$ and $u_{i,t}(n) = (1-\rho_{i,t})d_{i,t}(n) - \rho_{i,t}h_{i,t}$. Note, that this is the same cross-sectional specification as in KY (2019) except for the lagged ownership term $q_{i,t-1}(n)$. The reduced-form coefficients $\alpha_{i,t}$ and $\rho_{i,t}$ allow backing out the investor's long and short-term elasticity as

$$\zeta^{SR}_{i,t} \approx 1 - \alpha_{i,t} \tag{15}$$

$$\zeta^{LR}_{i,t} \approx (1-\rho_{i,t})^{-1}(1-\alpha_{i,t}) \tag{16}$$

[14] Any portfolio-level growth is subsumed by time-fixed effects in the empirical specification.

[15] Calvet et al. (2009) propose an analogous specification for the cross-section of household rebalancing behavior. In an earlier version of their paper, Gabaix et al. (2024a) introduce a similar model for capturing inertia in the time series of household demand. Similar partial adjustment models (PAMs) have been employed in capital structure regressions by Fama and French (2002) and Kayhan and Titman (2007).

The long-run elasticity is larger than the short-run elasticity by the partial adjustment factor $(1-\rho_{i,t})^{-1}$. Between the instantaneous and long-term impact are h-quarter elasticities given by $\zeta_{i,t}^{h} \approx \sum_{s=0}^{h} \rho_{i,t}^{s}(1-\alpha_{i,t})$. This approximation holds for diversified portfolios with small $w_{i,t}(n)$. Appendix A.6 provides exact expressions including the full elasticity matrix including the cross-elasticities implied by the budget constraint. To rule out upward-sloping demand curves, I follow KY (2019) and impose the coefficient constraint $\alpha_{i,t} < 1$ for all investor-quarter pairs. To maximize comparability to the static estimation, I use the mandate-based demand in levels to instrument for endogenous price levels in (14). I estimate both the partial adjustment model and the static logit specification (i.e. excluding $\rho_{i,t}q_{i,t-1}(n)$), quarterly from 1990 to 2024 for every investor using the same stock universes and pooling of narrow holdings across the estimations. Under valid instruments for market equity, we can consistently estimate the short-run elasticity $1-\alpha_{i,t}$. However, extrapolating from the short to the long run by dividing by $(1-\rho_{i,t})$ is only valid if $\rho_{i,t}$ is well identified.

C Identifying Inertia

If unobservable demand shocks are persistent, then $\rho_{i,t}$ is not consistently estimated. Intuitively, a long-run value investor may hold the same portfolio in the current quarter as last quarter not because she is inert, but because her value signals have not changed. However, note, that $q_{i,t-1}(n)$ are not past portfolio weights but past ownership *relative* to shares outstanding, i.e., total supply. The endogeneity concern is therefore less severe than it may initially appear. Persistent demand shocks affect portfolio weights but not necessarily relative ownership. I nevertheless provide two possible instruments for $q_{i,t-1}$ that are plausibly exogenous to current latent demand. First, I instrument $q_{i,t-1}$ by its long-run (3-year) lag. Since trading signals are unlikely to persist over several years, this helps mitigate the problem of omitted persistent factors that jointly affect past portfolios and current latent demand. Second, I use a model-based instrument by estimating the counterfactual lagged portfolio held in the absence of latent demand shocks, i.e., the portfolios implied by lagged observables. Because the results are largely unaffected by the choice of the instrument, the main text reports the results using raw lagged ownership $q_{i,t-1}(n)$. Appendix B provides further details and shows that all results are qualitatively and quantitatively robust to the instrument choice. Last, in all specifications I impose the constraint that $\rho_{i,t} \in [0,1)$. While for the majority of investors, the unconstrained $\hat{\rho}_{i,t}$ already lies between 0 and 1, this ensures finite long-run elasticities and unique equilibria.

D Estimation Results

Table 2 reports the estimated coefficients from the partial adjustment model across investor groups relative to the benchmark logit estimates without the partial adjustment term.

Portfolio Inertia. The first column reports inertia $\rho_{i,t}$ across investor groups. The average investor is relatively inert and partially trades towards their long-term desired portfolio with a speed of $1 - \rho_{i,t} = 0.16$. Bacchetta et al. (2023) find similarly low adjustment speeds for international mutual funds using an alternative econometric specification with different instruments. The overall high inertia in institutional portfolios is consistent with the reduced form evidence on increasing elasticities at longer horizons. Gradual portfolio adjustments lead to a large impact of demand shocks and excess volatility at short horizons that only slowly dissipate over time. Intuitively, changes in large positions should take longer to execute as trading greater fractions of current volume implies higher price impacts. In line with higher trading costs for bigger positions, I find that large investors are on average 23% more inert than small investors. The bottom rows of Table 2 confirm that results are robust to using different instruments for lagged portfolios, with adjustment speeds of 0.23 and 0.15 using long-run lags and the model-based instrument, respectively.

Short-run Elasticities. The second column reports the short-run elasticity $\zeta_{i,t}^{SR}$ from the partial adjustment model. The AUM-weighted average across all investors is 0.32. Among investors, elasticities exhibit an intuitive pattern: small investors are more elastic than large investors, value investors are more elastic than growth investors, and investors with high active shares (Cremers and Petajisto (2009)) are more elastic than those with low active shares. Large investors with a low active shares are the least elastic investor group with an average short-run elasticity of 0.12. Across all groups, short-run elasticities are lower than the static logit estimates.[16] Note that both estimators use the same instrument for endogenous market equities. This suggests that - conditional on the instrument - logit elasticities overstate how aggressively investors trade *quarterly* and instead capture what investors trade towards over longer horizons. Intuitively, an equal-weighted investor that rebalances their portfolio every year has a logit elasticity of around 1, but

[16]Note, the slight difference to the average coefficient reported in Table 1, which use panel regressions and a pooled first stage to facilitate comparison with the reduced-form estimates.

a quarterly short-run elasticity of 0 in most quarters.

Long-run Elasticities. Slow portfolio adjustments lead to implied long-run elasticities $\zeta_{i,t}^{LR} = \frac{\zeta_{i,t}^{SR}}{1-\rho_{i,t}}$ that are around 4 times higher than their short-term counterpart. The AUM-weighted average long-run elasticity implied by the partial adjustment model is 1.2. While many large investors are not very elastic at short horizons, they provide considerable liquidity at long horizons. Over the long run, large value investors are the most important liquidity providers with a long-run elasticity of 5.2. This emphasizes the importance of value investors in stabilizing equity markets over the long term. The static logit elasticity lies between the short- and the long-run elasticity from the partial adjustment model. This implies that the impact of counterfactuals based on logit demand systems may *underestimate* the true price response in the short run and *overestimate* the long-run effect. Note, that the average long-term elasticity implied by the partial adjustment model is slightly lower than the reduced-form long-run elasticity of 1.68 in Table 1. Naturally, the overall magnitudes of elasticities depend on the price instrument used. Less persistent price instruments such as flow-driven trades correspond to higher elasticities. However, the *ratios* between long and short-term elasticities are quantitatively consistent. Both, reduced-form and structural estimates suggest that long-term elasticities are 3 to 4 times larger than quarterly elasticities. This has important implications for the price impact of counterfactual experiments, which I explore in Section 6.

[Table 2 about here.]

E Decomposing the Difference in Elasticities

What drives the difference between long-run elasticities from the partial adjustment model $\zeta_{i,t}^{LR}$ and the elasticities from static logit demand $\zeta_{i,t}^{KY}$? On the one hand, controlling for lagged holdings mitigates the problem of omitted persistent demand shocks that are correlated with long-run price levels. On the other hand, the partial adjustment model introduces heterogeneous trading horizons which may partly drive the divergence of $\zeta_{i,t}^{LR}$ and $\zeta_{i,t}^{KY}$.[17] Figure 1 shows that the difference in coefficients $\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}$ is strongly related to investor size. Large investors are considerably more elastic in the long run than what their static logit estimate implies. Note that because $\zeta_{i,t}^{LR} = \frac{\zeta_{i,t}^{SR}}{1-\rho_{i,t}}$ any difference between the two estimates can either be attributed to trading horizons ($\rho_{i,t} \neq 0$) or omitted

[17] See simulation in Appendix Section C.

variables ($\zeta_{i,t}^{SR} \neq \zeta_{i,t}^{KY}$). I therefore decompose the difference in coefficients into two components

$$\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY} = TH_{i,t} + OV_{i,t} \tag{17}$$

where $TH_{i,t} = \zeta_{i,t}^{LR} - \zeta_{i,t}^{SR}$ captures variation due to trading horizons and $OV_{i,t} = \zeta_{i,t}^{SR} - \zeta_{i,t}^{KY}$ captures variation from omitted variables. In order to see why $OV_{i,t}$ captures omitted variables, note that the partial adjustment model (14) can be rewritten as

$$\tilde{\delta}_{i,t}(n) = \alpha_{i,t} p_t(n) + u_{i,t}(n).$$

where $\tilde{\delta}_{i,t} \equiv \delta_{i,t} - \rho_{i,t} q_{i,t-1}$ are quasi-holdings corrected for persistent demand shifts. Therefore any difference between $\zeta_{i,t}^{SR}$ and $\zeta_{i,t}^{KY}$ is driven by correcting the left-hand side for persistent demand shifts that are correlated with prices. The grey and red bars in Figure 1 plot the decomposition of $\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}$ into $TH_{i,t}$ and $OV_{i,t}$. The difference in elasticities is primarily driven by trading horizons, i.e. by inertia $\rho_{i,t}$. A simple variance decomposition reveals that 84% (std. error 1.4%) of the variation in $\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}$ is driven by $TH_{i,t}$.[18] While the omitted variable effect is 0 for large investors, it is negative for small investors. The portfolios of smaller investors are tilted towards small-cap stocks, which gives the impression of elasticity when failing to control for past holdings. Therefore static logit elasticities overstate short-run trading activity for smaller institutions.

[Figure 1 about here.]

6 The Impact and Reversal of Counterfactual Experiments

A Short-run versus Long-run Multipliers

As implied by (5), the price multiplier matrices based on short and long-run elasticities are respectively given by the ownership-weighted sum of the investor-specific elasticities

$$\mathcal{M}_t^{SR} \equiv \left(\sum_{i=1}^{I} \text{diag}(Q_{i,t}) \zeta_{i,t}^{SR} \right)^{-1} \tag{18}$$

$$\mathcal{M}_t^{LR} \equiv \left(\sum_{i=1}^{I} \text{diag}(Q_{i,t}) \zeta_{i,t}^{LR} \right)^{-1}. \tag{19}$$

[18] Formally, the variance decompositioin is given by $1 = \frac{\text{cov}\left(\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}, TH_{i,t}\right)}{\text{var}\left(\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}\right)} + \frac{cov(\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}, OVB_{i,t})}{var(\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY})}$.

Similarly, medium-term multipliers over h quarters are given by $\mathcal{M}_t^h \equiv \left(\sum_{i=1}^{I} \text{diag}(Q_{i,t})\zeta_{i,t}^h\right)^{-1}$. Panel a) of Figure 2 plots the stock-specific multiplier implied by the dynamic demand system. The average short-run impact across stocks is 4.5 and strongly declines over time to around 1. This suggests that the impact of counterfactual experiments is over 4 times larger in the short run than over long horizons. Again, while the overall magnitude of counterfactual price changes likely depends on the instrument with which elasticities are identified, Figure 2 suggests that there are important differences in the *relative* impact over time. Note, that while investor-specific elasticities do not vary across stocks (assuming $w_{i,t}(n) \approx 0$), ownership shares $Q_{i,t}(n)$ vary considerably. This implies cross-sectional variation in multipliers across stocks. Panel b) plots the distribution of the dynamic multiplier. There is considerable heterogeneity in price impact and reversal in the cross-section of stocks. Some stocks are more inert than others which results in heterogeneous reversal speeds. Section C explores the return predictability from this novel measure of stock-specific long-term reversal.

[Figure 2 about here.]

B Testing the Demand System Out-of-Sample

The counterfactual prices obtained from the demand system depend on the multiplier matrix implied by the elasticities (see Proposition 1). A simple out-of-sample test of the demand system is whether it predicts the price path following *exogenous* demand shocks that have not been used for identifying the system. Note, that the elasticities from the partial adjustment model have been identified from the mandate-based instrument. Therefore the impact of mutual funds' flow-induced trades provides a true out-of-sample test of the estimated elasticities and the long-term reversal implied by the demand system. Under the assumption, that flow-induced trades are *exogenous* demand shifts, one can test the demand system by comparing the simulated price impact to the observed impact. The black bars in Figure 3 plot the cumulative impact of flow-induced trades over time.[19] I then simulate the price impact of the same flow-induced trades using the horizon-specific multiplier matrices from above. The counterfactual impact implied by the partial adjustment model impressively matches the observed impact and its reversal both qualitatively and quantitatively. Instead, static logit demand implies a constant price impact that underestimates the short-run and overestimates the long-run effect.

[19] I estimate the cumulative future price impact following a one time demand shock at $t-h$ using a distributed lag model given by $\Delta p_t(n) = \sum_{l=1}^{h} \beta_l \Delta z_{i,t-l}(n) + u_{i,t}(n)$ including time-fixed effects.

[Figure 3 about here.]

C *Inertia and Return Predictability*

Because investors are inert, elasticities are larger in the long run and the price impact from demand shocks reverts over time. I define long-term reversal as the difference in long- and short-run multipliers

$$\mathcal{R}_t(n) \equiv \mathcal{M}_t^{SR}(n) - \mathcal{M}_t^{LR}(n). \tag{20}$$

which is a novel stock-level measure of reversal. Because $\rho_{i,t} \in [0, 1)$ and $\alpha_{i,t} < 1$, short-run price impact (ceteris paribus) always exceeds long-run price impact and $\mathcal{R}_t(n) \geq 0$. Two factors drive reversal. First, stocks in the habitat of short-term investors (with low $\rho_{i,t}$) have similar long and short-term impacts and therefore exhibit less reversal. Second, stocks in the habitat of long-term investors (with high $\rho_{i,t}$) are particularly exposed to reversal if these investors have high elasticities. This encapsulates the difference in short- and long-run dynamics which is critically driven by the interaction of inertia and elasticity. An intuitive example is stocks in the habitat of long-run value investors. The initial price impact may considerably exceed the long-run impact if the primary investors in the stock become more elastic over time.

I test whether $\mathcal{R}_t(n)$, a new measure of long-term reversal, contains additional information beyond standard return predictors. Reversal should only affect returns after a demand shock. If momentum is partly driven by non-fundamental shocks, then stocks with higher $\mathcal{R}_t(n)$ should exhibit weaker momentum: As investors become more elastic over time, they trade more aggressively against the initial demand shocks over time, thereby acting contrarian against momentum. Figure 4 shows that among reversal stocks – those held by increasingly elastic investors, such as long-run value investors – momentum disappears entirely. In contrast, non-reversal stocks exhibit over 3% higher annualized momentum returns, with the Sharpe ratio of the unconditional strategy rising from 0.37 to 0.54.

[Figure 4 about here.]

Table 3 confirms this using monthly Fama-MacBeth regressions. Conditioning on reversal the coefficient on momentum doubles from 0.26% to 0.47%. This underscores the incremental information embedded in the stock-specific reversal measure and the broader role of portfolio rebalancing dynamics in shaping return predictability.

[Table 3 about here.]

7 Conclusion

A growing literature uses portfolio holdings to quantify the impact of demand shocks on equilibrium prices via investors' price elasticity of demand. Existing studies estimate demand slopes using static logit regressions, implying constant price impacts in counterfactual experiments. In this paper, I provide both reduced-form and structural evidence that investors' price elasticity of demand increases substantially at longer horizons, reaching up to four times its short-term value. This suggests a key distinction between what investors hold in the cross-section and how they trade in the time series. The price impact of counterfactual demand shocks is therefore larger at high frequencies but dissipates over time as investors adjust their portfolios. To capture these dynamics, I develop a simple yet flexible extension of the logit framework – a partial adjustment model – that links holdings and trades within a unified demand system. I show that quadratic transaction costs are a possible micro-foundation for this model, rationalizing investors' gradual portfolio rebalancing. The dynamic model accurately predicts both the short-run impact of demand shocks and their long-run reversal out-of-sample. These findings emphasize the importance of accounting for trading horizons and portfolio adjustment speeds into demand-based asset pricing models, ensuring more accurate counterfactual predictions.

References

Angrist, J. D. and A. B. Krueger (1992). The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. *Journal of the American statistical Association 87*(418), 328–336.

Bacchetta, P., M. Davenport, and E. Van Wincoop (2022). Can sticky portfolios explain international capital flows and asset prices? *Journal of International Economics 136*, 103583.

Bacchetta, P., S. Tieche, and E. Van Wincoop (2023). International portfolio choice with frictions: Evidence from mutual funds. *The Review of Financial Studies 36*(10), 4233–4270.

Bacchetta, P. and E. Van Wincoop (2021). Puzzling exchange rate dynamics and delayed portfolio adjustment. *Journal of International Economics 131*, 103460.

Balasubramaniam, V., J. Y. Campbell, T. Ramadorai, and B. Ranish (2021). Who owns what? a factor model for direct stock holding. Technical report, National Bureau of Economic Research.

Benetton, M. and G. Compiani (2020). Investors' beliefs and asset prices: A structural model of cryptocurrency demand. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-107).

Betermier, S., L. E. Calvet, S. Knüpfer, and J. Kvaerner (2022). What do the portfolios of individual investors reveal about the cross-section of equity returns? *Available at SSRN 3795690*.

Bogousslavsky, V. (2016). Infrequent rebalancing, return autocorrelation, and seasonality. *The Journal of Finance 71*(6), 2967–3006.

Bogousslavsky, V., P. Collin-Dufresne, and M. Sağlam (2021). Slow-moving capital and execution costs: Evidence from a major trading glitch. *Journal of Financial Economics 139*(3), 922–949.

Bohn, H. and L. L. Tesar (1996). Us equity investment in foreign markets: portfolio rebalancing or return chasing? *The American Economic Review 86*(2), 77–81.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma (2024). Institutional corporate bond pricing. *Swiss Finance Institute Research Paper* (21-07).

Calvet, L. E., J. Y. Campbell, and P. Sodini (2009). Fight or flight? portfolio rebalancing by individual investors. *The Quarterly journal of economics 124*(1), 301–348.

Chang, Y.-C., H. Hong, and I. Liskovich (2015). Regression discontinuity and the price effects of stock market indexing. *The Review of Financial Studies 28*(1), 212–246.

Chaudhary, M., Z. Fu, and J. Li (2023). Corporate bond multipliers: Substitutes matter. *Available at SSRN*.

Collin-Dufresne, P., K. Daniel, and M. Sağlam (2020). Liquidity regimes and optimal dynamic asset allocation. *Journal of Financial Economics 136*(2), 379–406.

Coqueret, G. (2022). Characteristics-driven returns in equilibrium. *arXiv preprint arXiv:2203.07865*.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

Cremers, K. M. and A. Petajisto (2009). How active is your fund manager? a new measure that predicts performance. *The review of financial studies 22*(9), 3329–3365.

Darmouni, O., K. Siani, and K. Xiao (2022). Nonbank fragility in credit markets: Evidence from a two-layer asset demand system. *Available at SSRN 4288695*.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of finance 65*(4), 1237–1267.

Edmans, A., I. Goldstein, and W. Jiang (2012). The real effects of financial markets: The impact of prices on takeovers. *The Journal of Finance 67*(3), 933–971.

Enke, B. and T. Graeber (2023). Cognitive uncertainty. *The Quarterly Journal of Economics 138*(4), 2021–2067.

Fama, E. F. and K. R. French (2002). Testing trade-off and pecking order predictions about dividends and debt. *Review of financial studies*, 1–33.

Fang, X., B. Hardy, and K. K. Lewis (2022). Who holds sovereign debt and why it matters. Technical report, National Bureau of Economic Research.

Froot, K. A., P. G. O'connell, and M. S. Seasholes (2001). The portfolio flows of international investors. *Journal of financial Economics 59*(2), 151–193.

Gabaix, X. and R. S. Koijen (2023). In search of the origins of financial fluctuations: The inelastic markets hypothesis. *Available at SSRN 3686935*.

Gabaix, X., R. S. Koijen, F. Mainardi, S. Oh, and M. Yogo (2024a). Asset demand of us households. Technical report, National Bureau of Economic Research.

Gabaix, X., R. S. Koijen, F. Mainardi, S. Oh, and M. Yogo (2024b). Limited risk transfer between investors: A new benchmark for macro-finance models. *Available at SSRN*.

Gârleanu, N. and L. H. Pedersen (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance 68*(6), 2309–2340.

Gârleanu, N. and L. H. Pedersen (2016). Dynamic portfolio choice with frictions. *Journal of Economic Theory 165*, 487–516.

Greenwood, R. (2005). Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics 75*(3), 607–649.

Grossman, S. J. and M. H. Miller (1988). Liquidity and market structure. *the Journal of Finance 43*(3), 617–633.

Haddad, V., P. Huebner, and E. Loualiche (2021). How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing. *Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing (April 7, 2021)*.

Han, X., N. L. Roussanov, and H. Ruan (2021). Mutual fund risk shifting and risk anomalies. *Available at SSRN 3931449*.

Hartzmark, S. M. and D. H. Solomon (2021). Predictable price pressure. *Available at SSRN 3853096*.

Hou, K., C. Xue, and L. Zhang (2015). Digesting anomalies: An investment approach. *The Review of Financial Studies 28*(3), 650–705.

Huebner, P. (2023). The making of momentum: A demand-system perspective. In *Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting*.

Jansen, K. A. (2024). Long-term investors, demand shifts, and yields. *The Review of Financial Studies*, hhae071.

Jansen, K. A., W. Li, and L. Schmid (2024). Granular treasury demand with arbitrageurs. Technical report, National Bureau of Economic Research.

Jiang, Z., R. J. Richmond, and T. Zhang (2022). Understanding the strength of the dollar. Technical report, National Bureau of Economic Research.

Jiang, Z., R. J. Richmond, and T. Zhang (2024). A portfolio approach to global imbalances. *The Journal of Finance 79*(3), 2025–2076.

Kashyap, A. K., N. Kovrijnykh, J. Li, and A. Pavlova (2023). Is there too much benchmarking in asset management? *American Economic Review 113*(4), 1112–1141.

Kayhan, A. and S. Titman (2007). Firms' histories and their capital structures. *Journal of financial Economics 83*(1), 1–32.

Kelly, B. T., S. Pruitt, and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics 134*(3), 501–524.

Koijen, R. S., F. Koulischer, B. Nguyen, and M. Yogo (2021). Inspecting the mechanism of quantitative easing in the euro area. *Journal of Financial Economics 140*(1), 1–20.

Koijen, R. S., R. J. Richmond, and M. Yogo (2024). Which investors matter for equity valuations and expected returns? *Review of Economic Studies 91*(4), 2387–2424.

Koijen, R. S. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy 127*(4), 1475–1515.

Koijen, R. S. and M. Yogo (2020). Exchange rates and asset prices in a global demand system. Technical report, National Bureau of Economic Research.

Lettau, M., S. C. Ludvigson, and P. Manoel (2018). Characteristics of mutual fund portfolios: Where are the value funds? Technical report, National Bureau of Economic Research.

Li, J. and Z. Lin (2022). Price multipliers are larger at more aggregate levels. *Available at SSRN 4038664*.

Lou, D. (2012). A flow-based explanation for return predictability. *The Review of Financial Studies 25*(12), 3457–3489.

Mitchell, M., L. H. Pedersen, and T. Pulvino (2007). Slow moving capital. *American Economic Review 97*(2), 215–220.

Pavlova, A. and T. Sikorskaya (2022). Benchmarking intensity.

Petajisto, A. (2009). Why do demand curves for stocks slope down? *Journal of Financial and Quantitative Analysis 44*(5), 1013–1044.

Schmickler, S. (2020). Identifying the price impact of fire sales using high-frequency surprise mutual fund flows. *Available at SSRN 3488791*.

Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance 41*(3), 579–590.

Tamoni, A., S. Sokolinski, and Y. Li (2024). Which investors drive anomaly returns and how? *Available at SSRN 4242745*.

Wardlaw, M. (2020). Measuring mutual fund flow pressure as shock to stock returns. *The Journal of Finance 75*(6), 3221–3243.

Wurgler, J. and E. Zhuravskaya (2002). Does arbitrage flatten demand curves for stocks? *The Journal of Business 75*(4), 583–608.

Figure 1: **Long-run Elasticities versus Logit Elasticities**.
The figure decomposes the difference between long-run elasticities from the partial adjustment model $\zeta_{i,t}^{LR}$ and the static elasticities from logit demand $\zeta_{i,t}^{KY}$. The black bars plot $\zeta_{i,t}^{LR} - \zeta_{i,t}^{KY}$, averaged across AUM deciles. The red and grey bars decompose the difference into heterogeneity in trading horizons $TH_{i,t} = \zeta_{i,t}^{LR} - \zeta_{i,t}^{SR}$ and omitted variables $OV_{i,t} = \zeta_{i,t}^{SR} - \zeta_{i,t}^{KY}$.

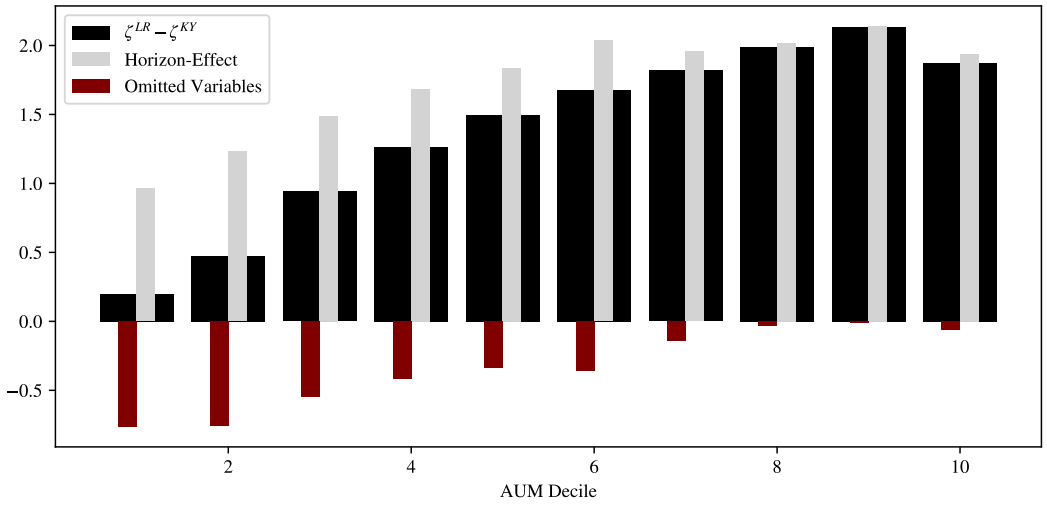
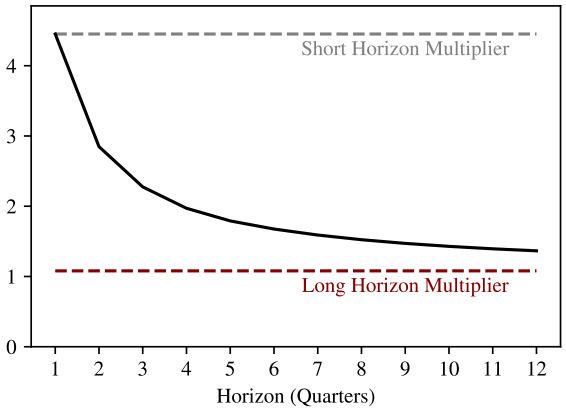


Figure 2: **Stock-specific Dynamic Multipliers**.
Panel a) of the figure plots the multiplier (i.e., price impact of counterfactual experiments) over different horizons $\mathcal{M}_t^h \equiv \left(\sum_{i=1}^{I} \text{diag}(Q_{i,t})\zeta_{i,t}^h\right)^{-1}$. I report the median $\mathcal{M}_t^h$ across stocks and quarters for horizons ranging from 1 to 12 quarters. Panel b) plots the distribution of the multiplier across stocks.

(a) Short versus Long-run Impact



(b) Distribution of Multiplier

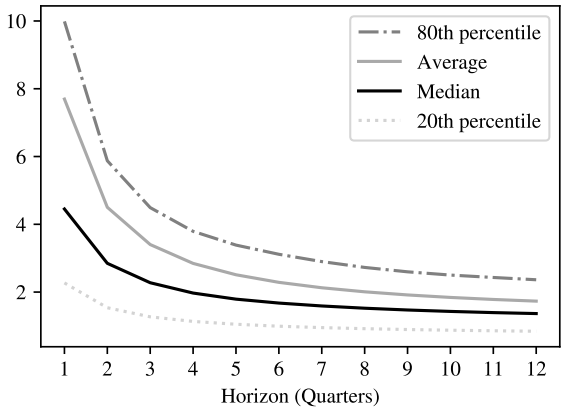


Figure 3: **Predicting the Impact of Flow-Induced Trades**.
The figure plots the price impact of flow-induced trades against the simulated price impact from the partial adjustment model. I estimate the long-run impact using a distributed lag model $\Delta p_t(n) = \sum_{l=1}^{h} \beta_l \Delta z_{i,t-l}(n) + u_{i,t}(n)$ including time-fixed effects and stock-specific control variables given by log book equity, market beta, profitability, investment and dividends-to-book equity. I plot the cumulative coefficients and 95% confidence intervals from a Wald test with standard errors clustered by time. The red and blue lines report the structurally implied impact, $\frac{1}{T}\sum_{t=1}^{T} \text{cov}_t\left(\mathcal{M}^h(n), \Delta z(n)\right) / \text{var}_t\left(\Delta z(n)\right)$, which is a weighted average of the stock-specific multiplier, where stock-time observations are weighted proportional to their importance in the OLS regressions.

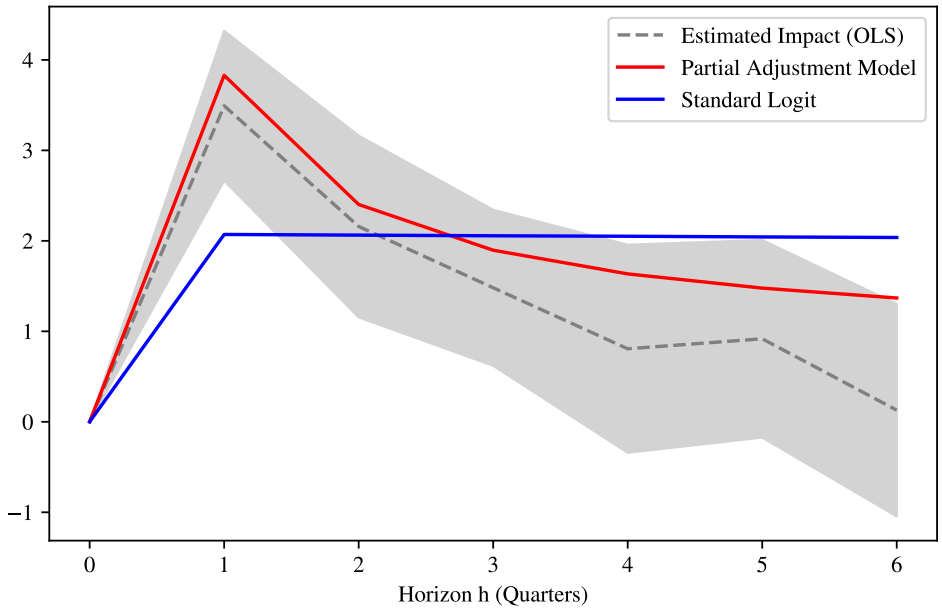


Figure 4: **Inertia and Momentum Returns**.
The figure plots the cumulative returns of monthly rebalanced momentum portfolios. Momentum is defined as the past cumulative returns from 12 months to 1 month prior to the portfolio formation date $[t-12, t-1]$. The black line is the return on a portfolio that goes long (short) the top (bottom) tercile of momentum stocks. The blue and red lines report conditional momentum returns, by splitting the sample in two halves based on the median stock-specific reversal $\mathcal{R}_t(n)$ in the formation month and computing conditional momentum terciles.

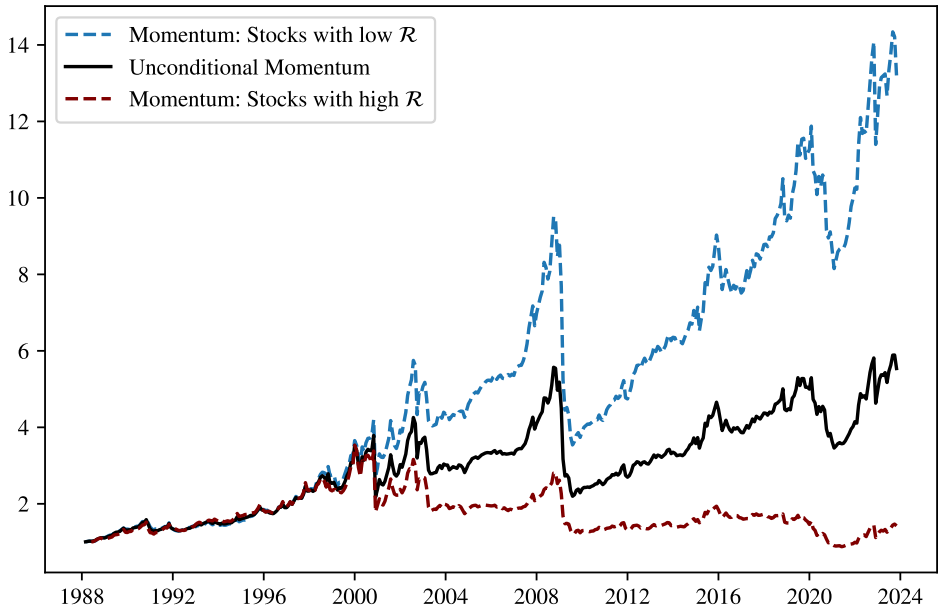


Table 1: **Reduced-Form Elasticities at Different Horizons.**
The table reports the estimated elasticities from the simple difference-in-difference estimator

$$\text{1st Stage:} \quad \Delta^h p_t(n) = \mathcal{M}^h \Delta^h z_t(n) + u_t(n)$$
$$\text{2nd Stage:} \quad \Delta^h q_{i,t}(n) = -\zeta_i^h \Delta^h \hat{p}_t(n) + d_{i,t}(n)$$

for horizons ranging from $h = 1$ quarter up to 4 years. $d_{i,t}(n)$ and $u_t(n)$ include time-fixed effects and stock-specific control variables, log book equity, dividends to book equity, profitability, investment, and market beta. I follow KY (2019) and impose $\zeta_i^h > 0$ for all investors at all horizons. Column 1 reports the AUM-weighted average elasticity across all investors where AUM-weights are averaged over the sample period. The standard error of the AUM-weighted average elasticity is computed by bootstrapping 100 draws from the standard errors on the investor-specific elasticities and constructing the AUM-weighted average for each sample. Columns 2-4 report the distribution of elasticities across all investors, excluding investors with binding constraints ($\zeta_i^h = 0$). The rightmost column reports the first-stage coefficient, $\mathcal{M}^h$, estimated from a pooled regression across all stocks and time periods. Standard errors are clustered by time. The bottom row reports the static logit elasticity identified from KY's mandate-based instrument.

Horizon h (Quarters)	Avg. Elasticity ζ_i^h		Elasticity Distribution			First Stage	
			10th pctl.	Median	90th pctl.		
1	0.58	(0.2)	0.14	0.47	0.8	2.57	(0.38)
2	0.68	(0.21)	0.35	0.61	1.01	2.2	(0.3)
3	0.85	(0.22)	0.24	0.73	1.25	1.94	(0.25)
4	0.72	(0.16)	0.36	0.8	1.39	1.71	(0.23)
5	0.81	(0.17)	0.25	0.92	1.55	1.53	(0.19)
6	0.87	(0.18)	0.3	1.07	1.88	1.37	(0.15)
7	0.99	(0.15)	0.24	1.07	2.21	1.23	(0.12)
8	1.0	(0.11)	0.42	1.2	2.41	1.15	(0.11)
9	1.02	(0.12)	0.38	1.21	2.6	1.06	(0.11)
10	1.14	(0.13)	0.42	1.22	2.52	0.97	(0.13)
11	1.35	(0.14)	0.47	1.53	3.17	0.89	(0.14)
12	1.41	(0.13)	0.5	1.78	3.86	0.82	(0.15)
13	1.43	(0.13)	0.51	1.91	4.34	0.77	(0.16)
14	1.55	(0.15)	0.3	1.92	3.94	0.72	(0.16)
15	1.59	(0.16)	0.67	2.19	4.67	0.66	(0.17)
16	1.68	(0.17)	0.87	2.5	5.02	0.61	(0.18)
Static Logit (Levels)	0.73	(0.01)	1.52	2.18	2.78	0.58	(0.02)

Table 2: **Structural Elasticities from Partial Adjustment Model.**
The table reports the results of estimating the constrained dynamic demand curve $\delta_{i,t}(n) = \rho_{i,t} q_{i,t-1}(n) + \alpha_{i,t} p_t(n) + X_t \beta_{i,t} + \epsilon_{i,t}(n)$ from 1990 to 2024 quarterly for every investor. The moment condition is given by

$$\mathbb{E}_t[\epsilon_{i,t}(n) | q_{i,t-1}(n), z^M_{i,t}(n), X_t] \quad \text{s.t.} \quad \rho_{i,t} \in [0,1), \alpha_{i,t} < 1 \quad \forall i,t.$$

where X_t is a vector of controls including a constant, log book equity, market beta, dividends to book equity, investment, and profitability, and $z^M_{i,t}(n)$ is KY's mandate-based instrument. Short- and long-run elasticities are given by $\zeta^{SR}_{i,t} = 1 - \alpha_{i,t}$ and $\zeta^{LR}_{i,t} = (1 - \rho_{i,t})^{-1}(1 - \alpha_{i,t})$ respectively. I group investors into deciles by size, active share, and value/growth tilts and report average coefficients for each group. Active shares are computed as $\frac{1}{2}\sum_n |w_{i,t}(n) - w^M_{i,t}(n)|$ where $w^M_{i,t}(n)$ are value-weights within the investor's subset of stocks held. Value/growth tilts are computed as portfolio-weighted averages of market-to-book ratios following Lettau et al. (2018). The bottom rows report the AUM-weighted average elasticity across investors when instrumenting $q_{i,t-1}(n)$ with its 3-year lag and the model-based instrument respectively. The rightmost column reports the elasticities obtained from a static quarterly logit estimation following KY (2019).

	Dynamic Estimates from Partial Adjustment Model			
	Inertia $\rho_{i,t}$	Short-Run Elasticity $\zeta^{SR}_{i,t}$	Long-Run Elasticity $\zeta^{LR}_{i,t} = \zeta^{SR}_{i,t}/(1-\rho_{i,t})$	Logit Elasticity $\zeta^{KY}_{i,t}$
All Investors (AUM-weighted)	0.84	0.31	1.19	0.5
Large Investors	0.8	0.73	2.66	0.79
Active	0.8	1.05	3.72	1.23
Passive	0.87	0.12	0.7	0.11
Growth	0.76	1.09	3.59	1.08
Value	0.75	1.58	5.18	1.69
Small & Med. Investors	0.65	0.91	2.61	1.28
Active	0.64	1.02	2.83	1.33
Passive	0.67	0.76	2.26	1.09
Growth	0.63	0.98	2.68	1.43
Value	0.61	1.14	2.97	1.5
Robustness to Instruments for Past Portfolios				
Long-run Lags	0.77	0.32	1.35	
Model-based Instrument	0.85	0.43	1.21	

Table 3: **Inertia and Momentum Returns: Fama-McBeth Regressions**
The table plots the coefficients to Fama-McBeth (FMB) regressions of monthly stock returns (in excess of the 1-month Treasury-bill rate) onto lagged stock-specific characteristics. The left column reports the baseline FMB regression. The right column interacts momentum with a dummy $1_{\mathcal{R}}$ equal to 1 if a stock has above median reversal $\mathcal{R}_{t-1}(n)$ in the prior month $t-1$.

Characteristics	Baseline		Conditional	
Log market equity	-0.27	(0.08)	-0.28	(0.08)
Book-to-market equity	0.08	(0.05)	0.08	(0.05)
Profitability	0.19	(0.05)	0.2	(0.05)
Investment	-0.36	(0.03)	-0.36	(0.03)
Market beta	0.08	(0.09)	0.08	(0.09)
Momentum	0.26	(0.09)	0.47	(0.1)
Momentum$\times 1_{\mathcal{R}}$			-0.52	(0.08)

Appendix A Proofs and Derivations

A.1 Intuition from a CARA Utility Investor

Assume that the stocks $n = 1, ..., N$ pay risky future cashflows $D(n)$ that are jointly normally distributed with a covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$. The investors' initial endowments $A_{i,0}$ are exogenously determined e.g. by the money that households allocate to them. The funds may invest in the N stock, as well as a riskless asset whose exogenous interest rate is normalized to 0. Under CARA preferences with risk aversion equal to γ_i, the institution maximizes $\mathbb{E}[e^{-\gamma_i A_i}]$, subject to the budget constraint on future wealth $A_i = A_{i,0} + Q_i'(\mathbb{E}[D] - P)$, by choosing the optimal portfolio

$$Q_i = (\gamma_i \Sigma)^{-1}(\mathbb{E}[D] - P). \tag{21}$$

Formally, the elasticity is given by $\text{diag}(Q_{i,t})^{-1}(\gamma_i \Sigma)^{-1}\text{diag}(P_t)$. For CARA preferences it is more convenient to define elasticities in absolute terms as opposed to percentages, in which case the elasticity matrix is given by $\tilde{\zeta}_i = -\frac{\partial Q_i}{\partial P'} = (\gamma_i \Sigma)^{-1}$. Market clearing implies that $\sum_i Q_i = Q^*$. The linear demand function implies that equilibrium market clearing prices are available in closed form and given by the CAPM

$$P = \mathbb{E}[D] - \gamma \Sigma Q^* \tag{22}$$

where $\gamma \equiv 1/\big(\sum_{i=1}^{I} 1/\gamma_i\big)$ is the market's effective risk aversion. In the CARA case, the most intuitive way to model exogenous demand ΔD is via a change in supply Q^*, such as the inclusion of some stocks in an index. The resulting price change is given by $\Delta P = \gamma \Sigma \Delta Q^*$. Thus the CARA setting implies that the multiplier matrix is proportional to $\gamma\Sigma$. It is driven by the covariance of cashflows and the market's effective risk aversion. Riskier stocks have a larger multiplier because they contribute more to the arbitrage portfolio. Note, that the channel through which demand shocks affect equilibrium prices in the CAPM is entirely risk-driven: If investors are risk neutral, or future cashflows are risk-free then $P = \mathbb{E}[D]$ and exogenous demand shocks do not affect prices. Greenwood (2005) shows that in a multi-period CARA setting the price impact of demand shocks decreases over time as uncertainty is resolved and the risk of future cashflows converges to 0. Formally, $\Delta P = \gamma \Sigma \Delta Q^*(T - t)$ where $T - t$ is the time left until liquidation time T. In other words, investors become more elastic over time as future payouts become more predictable. In the limit, two assets with equal riskless cashflows trade at identical

prices.

A.2 Dynamic Extension to the Generalized Framework

To illustrate the difference between short- and long-run dynamics, I propose a simple dynamic extension to the general framework in Section 2. Time is indexed by t. I assume that an investor's optimal dynamic portfolio is given by

$$Q_{i,t} = f_i\big(P_t, X_t, Q_{i,t-1}\big) \tag{23}$$

Via the lagged holdings term $Q_{i,t-1}$, the optimal portfolio depends on infinitely many lagged characteristics and prices. This introduces path dependence and can be motivated by e.g. portfolio adjustment costs. Section A.5 provides a possible microfoundation. The short-run elasticity captures the immediate response to price changes given by $\zeta_{i,t}^{SR} = -\text{diag}(Q_{i,t})^{-1}\frac{\partial Q_{i,t}}{\partial P_t'}\text{diag}(P_{i,t})$ as in the static model. The long-horizon elasticity accounts for the cumulative impact of a change in prices at time t, including both the direct impact at t and all indirect impacts via $Q_{i,t+1}, Q_{i,t+2}, \ldots$. The chain rule implies $\frac{\partial Q_{i,t+s}}{\partial P_t} = \frac{\partial f_i(P_{t+s}, X_{t+s}, Q_{i,t+s-1})}{\partial Q_{i,t+s-1}'}\frac{\partial Q_{i,t+s-1}}{\partial P_t'}$. Solving recursively yields

$$\frac{\partial Q_{i,t+s}}{\partial P_t'} = \prod_{k=1}^{s}\frac{\partial f_i}{\partial Q_{i,t+k-1}'}\frac{\partial Q_{i,t}}{\partial P_t'} \tag{24}$$

While this gives the instantaneous response to a lagged price change, the long-horizon elasticity is defined as the total cumulative response as s goes to infinity

$$\zeta_{i,t}^{LR} = -\sum_{s=0}^{\infty}\text{diag}(Q_{i,t})^{-1}\frac{\partial Q_{i,t+s}}{\partial P_t'}\text{diag}(P_{i,t}) \tag{25}$$

For simplicity, I assume that the inertia matrix $\frac{\partial f_i}{\partial Q_{i,t+k-1}'}$ is diagonal.[20] Plugging (24) into (25) yields

$$\zeta_{i,t}^{LR} = \zeta_{i,t}^{SR}\left(1 + \sum_{s=1}^{\infty}\prod_{k=1}^{s}\frac{\partial f_i}{\partial Q_{i,t+k-1}}\right) \tag{26}$$

If inertia is time-independent, i.e. $\frac{\partial f_i}{\partial Q_{i,t+k-1}'} = \frac{\partial f_i}{\partial Q_{i,t-1}'}$ one obtains (3).

[20]This implies that yesterday's holding in stock n affects both today's holding in n and m, but the spillover effect only comes through the short-term cross-elasticities, i.e. off-diagonal elements in $\zeta_{i,t}^{SR}$.

A.3 Proof of Equation (4)

The market clearing condition implies that $\sum_{i=1}^{I} Q_i = Q^*$ where $Q_i = f_i(P, X)$. Recall that the elasticity matrix is given by $\zeta_i = -\text{diag}(Q_i)^{-1}\frac{\partial Q_i}{\partial P'}\text{diag}(P)$. For simplicity, I also define the elasticity in absolute terms (instead of percentages) as $\tilde{\zeta}_i = -\frac{\partial Q}{\partial P'}$. We want to approximate the effects of an exogenous shock ΔX on equilibrium prices P. Differentiating both sides of the market clearing condition with respect to X yields

$$\frac{\partial Q^*}{\partial X'} = \sum_{i=1}^{I} \frac{\partial Q_i}{\partial X'} + \frac{\partial Q_i}{\partial P'}\frac{\partial P}{\partial X'} \tag{27}$$

Shares outstanding are normalized to 1. Therefore $\frac{\partial Q^*}{\partial X'} = 0$. Rewriting (27) in terms of elasticities yields and solving for $\frac{\partial P}{\partial X'}$ yields

$$\frac{\partial P}{\partial X'} = \left(\sum_{i=1}^{I} \tilde{\zeta}_i\right)^{-1} \sum_{i=1}^{I} \frac{\partial Q_i}{\partial X'}. \tag{28}$$

Let $\Delta d = \sum_{i=1}^{I} \frac{\partial Q_i}{\partial X'}\Delta X$ be the first order approximation to the aggregate demand shock caused by the exogenous shock ΔX. Because shares outstanding are normalized to 1, Δd is measured in percent of total supply. A first-order approximation to the equilibrium increase in prices is given by

$$\Delta P = \left(\sum_{i=1}^{I} \tilde{\zeta}_i\right)^{-1} \Delta d. \tag{29}$$

Left-multiplying both sides by $\text{diag}(P)^{-1}$ and expressing the resulting return $\Delta p = \text{diag}(P)^{-1}\Delta P$ in terms of elasticities in percentages ζ_i yields

$$\Delta p = \left(\sum_{i=1}^{I} \text{diag}(Q_i)\zeta_i\right)^{-1} \Delta d = \mathcal{M}\Delta d. \tag{30}$$

A.4 Counterfactual Experiments in KY (2019)

Motivated by the fact that portfolio weights are log-normally distributed in the data, KY (2019) use an exponential function for $f_i(P, X)$ and model investor demand as

$$w_i(n) = \frac{\exp(\delta_i(n))}{1 + \sum_n \exp(\delta_i(n))} \tag{31}$$

where $\delta_i(n) = \beta_{i,0}p(n) + u_i(n)$ and $u_i(n) = \sum_{k=1}^{K} \beta_{i,k}X_k(n) + \varepsilon_i(n)$. The sum of all investors must hold the market portfolio, thus

$$\log \sum_{i=1}^{I} A_i w_i(p) = p. \tag{32}$$

where $p \in \mathbb{R}^N$ is the vector of log market equities. Due to the exponential demand specification, equilibrium market equities are not available in closed form. However, a first-order approximation to the equilibrium price *change* following any counterfactual experiment is given by

$$\Delta p \approx \sum_{i=1}^{I} \left(\frac{\partial p}{\partial A_i} \Delta A_i + \frac{\partial p}{\partial u_i} \Delta u_i \right) \tag{33}$$

The demand shocks Δu_i capture demand changes due to changes in characteristics ΔX, preferences $\Delta \beta_i$ and latent demand $\Delta \varepsilon_i$. By differentiating the market clearing condition (32) with respect to the wealth distribution $\{A_i\}_{i=1}^{I}$ and demand shocks $\{u_i\}_{i=1}^{I}$ one can obtain closed-form solutions for the partial derivatives. I start with the assets under management. The derivative of (32) with respect to the assets under management of investor j, A_j, is given by:

$$\frac{\partial p}{\partial A_j} = \mathbf{H}^{-1} \left(\sum_{i=1}^{I} A_i \frac{\partial w_i}{\partial p'} \frac{\partial p}{\partial A_j} + w_j \right), \tag{34}$$

where $\mathbf{H} = \sum_{i=1}^{I} A_i \text{diag}(w_i)$. Solving for $\frac{\partial p}{\partial A_j}$ yields

$$\frac{\partial p}{\partial A_j} = \underbrace{\left(\mathbf{I} - \sum_{i=1}^{I} A_i \mathbf{H}^{-1} \frac{\partial w_i}{\partial p'} \right)}_{\mathcal{M}} \mathbf{H}^{-1} w_j = \mathcal{M} \mathbf{H}^{-1} w_j \tag{35}$$

Similarly, the derivative with respect to the other demand shocks is given by

$$\frac{\partial p}{\partial u_j} = \mathcal{M} \mathbf{H}^{-1} A_j \frac{\partial w_j}{\partial d_j} \tag{36}$$

The partial derivatives of w_j with respect to p and d_j are available in closed form. Let $\mathbf{G}_j = \text{diag}(w_j)(\mathbf{I} - \mathbf{1}w_j')$. The partial derivatives can be expressed as $\frac{\partial w_j}{\partial p'} = \mathbf{G}_j \beta_{j,0}$ and

$\frac{\partial w_j}{\partial d_j'} = \mathbf{G}_j$. Plugging into (33) yields

$$\Delta p \approx \mathcal{M}\left(\sum_{i=1}^{I} \mathbf{H}^{-1} w_i \Delta A_i + \mathbf{H}^{-1} \mathbf{G}_i \Delta u_i \right) \tag{37}$$

Note, that $\mathbf{H}^{-1}(w_i A_i)$ is the vector of ownership shares of investor i in all stocks. Because shares outstanding are normalized to 1, this is equivalent to $\text{diag}(Q_i)$. Thus we can express the equilibrium return as the product of a demand shock Δd and the multiplier matrix

$$\Delta p = \mathcal{M} \Delta d \tag{38}$$

where

$$\Delta d = \sum_{i=1}^{I} \text{diag}(Q_i) \left(\mathbf{1} \frac{\Delta A_i}{A_i} + (\mathbf{I} - \mathbf{1} w_j') \Delta u_i \right) \tag{39}$$

Because shares outstanding are normalized to 1, Δd is the aggregate demand shift in percent relative to total supply due to flows $\frac{\Delta A_i}{A_i}$ and portfolio changes $(\mathbf{I} - \mathbf{1} w_j') \Delta d_i$. Note, lastly, that as in KY (2019) the relationship between an investor's elasticity ζ_i and β_i is given by

$$\zeta_i = \mathbf{I} - \beta_i (\mathbf{I} - \mathbf{1} w_i') \tag{40}$$

See equation (14) in their paper. Plugging into the expression for $\mathcal{M}$ in (35) yields

$$\mathcal{M} = \left(\sum_{i=1}^{I} \text{diag}(Q_i) \zeta_i \right)^{-1} \tag{41}$$

A.5 Microfoundation: Partial Adjustment Model

Consider investor i who maximizes a mean-variance objective by choosing a vector of portfolio weights $\delta_t \in \mathbb{R}^{\mathbb{N}}$. This represents a slight abuse of notation with respect to the empirical specification, which defines δ_t as log portfolio weights relative to an outside asset. Expected returns and covariances are given by μ_t and Σ_t respectively. For simplicity, there are two periods $t-1$ and t. The multiperiod extension can be found in e.g. Gârleanu and Pedersen (2013) or Bacchetta et al. (2023) but does not add incremental insights for the context of this paper. The investor faces portfolio adjustment costs (AC_t) when trading $\Delta\delta_t = \delta_t - \delta_t^p$ dollars, where δ_t^p are the passive buy-and-hold dollar holdings at time t in the absence of trading between $t-1$ and t. Following Gârleanu and

Pedersen (2013), I assume quadratic portfolio adjustment costs that are proportional to the covariance matrix of returns

$$AC_t(\Delta\delta_t) = \frac{a}{2}\Delta\delta_t'\Sigma_t\Delta\delta_t. \tag{42}$$

Quadratic adjustment costs are a direct consequence of transacting with linear market impact.[21] The investor maximizes the standard mean-variance objective

$$\max_{\delta_t}\left\{\delta_t'\mu_t - \frac{\gamma}{2}\delta_t'\Sigma_t\delta_t - AC_t(\Delta\delta_t)\right\} \tag{43}$$

The optimal portfolio of shares at time t is a weighted mix of the past portfolio δ_t^p and the mean-variance optimal portfolio in the absence of trading costs $\delta_t^* = \frac{1}{\gamma}\Sigma_t^{-1}\mu_t$,

$$\delta_t = \rho\delta_t^p + (1-\rho)\delta_t^*, \tag{44}$$

where $\rho = \frac{a}{\gamma+a}$ captures the inertia in investors' portfolios. Note, that this microfounds a partial adjustment model in terms of portfolio weights. Instead, the empirical specification uses log portfolio weights, which tend to be normally distributed in the cross-section. I acknowledge this inconsistency and present this simplified microfoundation to enhance clarity in exposition and to offer some intuition regarding the underlying sources of portfolio inertia. It is possible to develop a micro-foundation for portfolio holdings that are log-linear in characteristics, as demonstrated by Koijen and Yogo (2019).

A.6 Long versus Short-Run Elasticities

Plugging the long-run desired portfolio $\delta_{i,t}^*(n) = \beta_{i,t}p_{i,t}(n) + d_{i,t}(n)$ in (13) yields

$$\delta_{i,t}(n) = \rho_{i,t}\delta_{i,t}^p(n) + (1-\rho_{i,t})\beta_{i,t}p_{i,t}(n) + (1-\rho_{i,t})d_{i,t}(n) \tag{45}$$

Let $h_{i,t} = \log(A_{i,t-1}w_{i,t-1}(0))$ denote the log holdings in the outside asset. Note, that $\delta_{i,t}^p(n) = \log(Q_{i,t-1}(n)P_t(n)) - h_{i,t-1}$. For simplicity of notation, I assume no valuation gains in the outside asset. Any portfolio-level growth is subsumed by time-fixed effects in the empirical specification. Subtracting log prices and adding lagged holdings in the

[21] Buying δ_t dollars worth of shares moves prices by $\delta_t'\Lambda_t$ where Λ_t is a multidimensional version of Kyle's lambda. Trading costs are given by δ_t times the price impact resulting in the quadratic term $\Delta\delta_t'\Lambda_t\Delta\delta_t$. If price impact is proportional to risk as in Greenwood (2005), this gives $\Delta\delta_t'\Sigma_t\Delta\delta_t$.

outside assets on both sides, we can rewrite the optimal portfolio as

$$q_{i,t}(n) = (1-\rho_{i,t})h_{i,t-1} + \rho_{i,t}q_{i,t-1}(n) + (1-\rho_{i,t})(\beta_{i,t}-1)p_{i,t}(n) + (1-\rho_{i,t})d_{i,t}(n) \quad (46)$$

Solving forward recursively results in the following long-run demand

$$q_{i,\infty}(n) = \frac{(1-\rho_{i,t})h_{i,t-1}}{(1-\rho_{i,t})} + \frac{(1-\rho_{i,t})(\beta_{i,t}-1)}{1-\rho_{i,t}}p_{i,t}(n) + \frac{(1-\rho_{i,t})d_{i,t}(n)}{1-\rho_{i,t}}. \quad (47)$$

and adding log prices $p_{i,t}$ on both sides and simplifying yields

$$\delta_{i,\infty}(n) = \beta_{i,t}p_{i,t}(n) + d_{i,t}. \quad (48)$$

Thus, over an infinitely long time horizon the investor holds the long-run desired portfolio $\delta_{i,\infty}(n) = \delta^*_{i,t}(n)$ with a long-run price response of $\beta_{i,t}$. Note, that $(1-\beta_{i,t})(1-\rho_{i,t})$ and $1-\beta_{i,t}$ are only approximations for the true short and long-run elasticities. The true elasticities are obtained by imposing the budget constraints $\sum_n w_{i,t}(n) = 1 - w_{i,t}(0)$ and $\sum_n w_{i,\infty}(n) = 1 - w_{i,\infty}(0)$ respectively, and differentiating $\delta_{i,t}(n)$ and $\delta_{i,\infty}(n)$ with respect to current prices. This yields

$$\zeta^{LR}_{i,t} = \mathbf{I} - \beta_{i,t}\big(\mathbf{I} - \mathbf{1}w'_{i,t}\big) \quad (49)$$

$$\zeta^{SR}_{i,t} = \mathbf{I} - (\beta_{i,t}(1-\rho_{i,t}) + \rho_{i,t})\big(\mathbf{I} - \mathbf{1}w'_{i,t}\big) \quad (50)$$

where the long-run elasticity $\zeta^{LR}_{i,t}$ is the same as in KY (equation 14, 2019) and the short-run elasticity $\zeta^{SR}_{i,t}$ is roughly smaller by a factor $(1-\rho_{i,t})$. When investors hold diversified portfolios (i.e. the portfolio weights in individual stocks are small), the elasticity matrices simplify to $\zeta^{LR}_{i,t} \approx (1-\beta_{i,t})\mathbf{I}$ and $\zeta^{SR}_{i,t} \approx (1-\beta_{i,t})(1-\rho_{i,t})\mathbf{I}$ respectively.

Appendix B Estimation Details

B.1 Flow-Induced Trading

To construct exogenous price variation at a quarterly frequency, I use flow-induced trades by domestic equity mutual funds as in Lou (2012).[22]. I set quarterly flows less than -

[22]The reduced-form literature on price multipliers has proposed several alternative exogenous demand shifts that could be used to instrument for price changes, such as index reconstitutions (see e.g. Chang et al. (2015), Pavlova and Sikorskaya (2022)) and, more recently, signed order flow aggregated to lower frequencies (Li and Lin (2022))

100% or greater than 200% to missing and only include funds for which the total assets computed from their portfolio holdings is between 75% and 120% of the total net assets (TNA) reported by CRSP. To the extent that fund flows are correlated with the funds' respective portfolio tilts, flow-induced trades may contain fundamental information and violate the exclusion restriction. For example, deteriorating earnings of technology firms may cause flows out of mutual funds with a technology tilt. I therefore construct orthogonalized flows $f_{i,t}^{\perp}$ using quarterly cross-sectional regressions of flows $f_{i,t}$ onto fund-level characteristics $C_{i,t}$. Formally,

$$\forall t : f_{i,t} = c_t + C_{i,t}\beta_t + f_{i,t}^{\perp} \tag{51}$$

where $C_{i,t}$ is a vector of fund-level characteristics. Fund-level characteristics are constructed using fund-level characteristic scores as in Lettau et al. (2018), which are portfolio-weighted averages of the stock characteristics $C_t(n)$ within a fund's portfolio, $C_{i,t} = \sum_{n=1} w_{i,t}(n)C_t(n)$. For every fund, I compute scores for greenness (based on MSCI ESG scores), value, size, momentum, profitability, investment, and beta.[23] Exogenous flow-induced demand is then simply given by summing over all hypothetical trades:

$$\Delta z_{i,t}(n) = \sum_{j \neq i}^{I} f_{j,t}^{\perp} Q_{j,t-1}(n) \tag{52}$$

Because the demand estimation is at the individual institution level, the correct instrument excludes all flow-induced trades from mutual funds that belong to institution i. Empirically, however, this does not materially affect any of the estimates because total flow-driven trading is not strongly affected by individual funds (even aggregated at the manager level). I therefore opt for the simpler version of the instrument given by $\Delta z_t(n) = \sum_{i=1}^{I} f_{i,t}^{\perp} Q_{i,t-1}(n)$. The long-horizon version of the flow-based instrument, $\Delta^h z_t$, is constructed by summing Δz_t over time.

B.2 Instruments for lagged portfolios

Under persistent demand shocks $q_{i,t-1}(n)$ is potentially endogenous, i.e., $\text{cov}(\varepsilon_{i,t}(n), q_{i,t-1}(n)) \neq 0$. It is possible that the normalization by shares outstanding does not fully remove the effect of persistent demand shocks on both the current portfolio and lagged ownership.

[23] Including dummy variables for all CRSP objective codes at the most granular level and controlling for past flow lags leaves all results in the paper unchanged.

I, therefore, propose two simple instruments denoted by $z^q_{i,t-1}(n)$.

Long-run lagged Ownership.
First, I instrument $q_{i,t-1}$ by its 3-year lag. This instrument satisfies the exclusion restriction if persistent signals only affect demand for a limited amount of time (here <3 years). While this instrument is clearly imperfect, the exogeneity can be gauged by how much the resulting inertia coefficient changes. Moving from a quarterly to a 3-year lag does not materially affect the coefficients suggesting that long-run persistent demand shocks do not pose a strong potential bias.

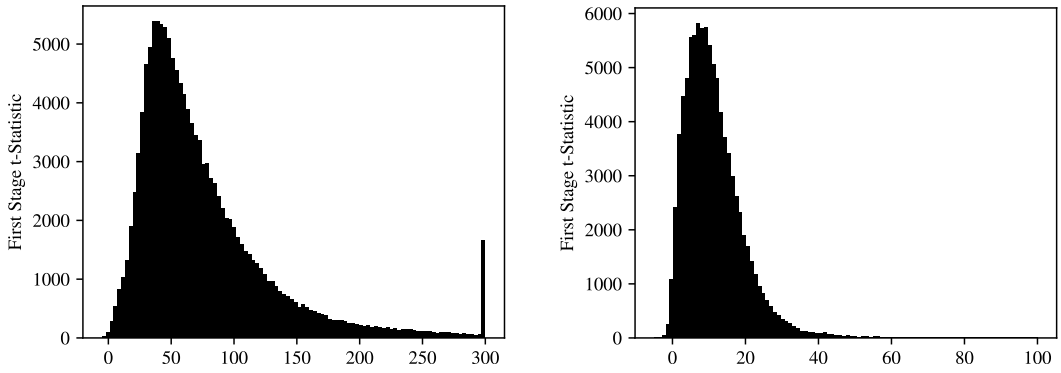
Model-based Instrument.
Second, I propose a simple model-based instrument. Let $\delta^X_{i,t-1}(n)$ denote the fitted value from a cross-sectional regression of portfolios $\delta_{i,t-1}(n)$ onto characteristics *excluding* lagged ownership $q_{i,t-2}(n)$.[24] I then construct the instrument for $q_{i,t-1}(n)$ as the counterfactual ownership (dollar holdings relative to market cap) in the absence of latent demand shifts, formally $z^q_{i,t-1}(n) = \delta^X_{i,t-1}(n) - p_{t-1}(n)$. Note, this is similar to instrumenting $q_{i,t-1}(n)$ directly with all lagged characteristics $X_{t-1}(n)$ and the lagged mandate-based instrument $z^M_{i,t-1}(n)$. The exclusion restriction holds under the assumption that latent demand is not driven by lagged characteristics, i.e. lagged characteristics only affect current demand via the inertia channel. A violation of this assumption could be momentum demand to the extent that momentum is not subsumed by profitability (see Hou et al. (2015)). Controlling for momentum does not materially affect the resulting inertia coefficient.[25]

Figure B.1 plots the first-stage t-statistic for both instruments over all investor-quarter pairs. Both instruments are highly relevant for the vast majority of investor-quarter pairs. Pooling investor-quarter pairs with weak instruments and re-estimating their demand leaves all results unchanged.

[24] Effectively, $\delta^X_{i,t-1}(n)$ is the fitted value from the level specification in KY (2019).
[25] The results are available upon request.

Figure B.1: **Relevance Inertia Instrument**.
The figure plots the first stage t-statistic on the model-based instrument (Panel a) and long-run lagged ownership (Panel b) for all investor-quarter pairs. The first-stage regression is given by $q_{i,t-1}(n) = \beta_{i,t} z^q_{i,t-1}(n) + \mathcal{M}_{i,t} z^q_{i,t}(n) + Controls + \epsilon_{i,t}(n)$ where the controls include log book equity, dividends to book equity, profitability, investment, and market beta as of t, as well as an investor-time-specific constant.

(a) Model-based Instrument (b) Long-run lagged Ownership



B.3 Constrained Estimation of Partial Adjustment Model

Putting the instruments and coefficient constraints together, the moment condition for identifying the partial adjustment model is given by

$$\begin{aligned} &\mathbb{E}_t\left[\varepsilon_{i,t}(n) \mid z^q_{i,t-1}(n), z^M_{i,t}(n), X_t(n)\right] \\ &\qquad \text{s.t.} \quad \rho_{i,t} \in [0,1), \alpha_{i,t} < 1. \end{aligned} \tag{53}$$

where $\varepsilon_{i,t}(n) = \delta_{i,t}(n) - \rho_{i,t} q_{i,t-1}(n) - \alpha_{i,t} p_t(n) - \sum_{k=1}^{K} X_{k,t}(n)\beta_{i,t,k}$. I estimate (53) via constrained GMM. However, estimation via unconstrained OLS and imposing the coefficient constraints ex-post yields roughly equivalent results.

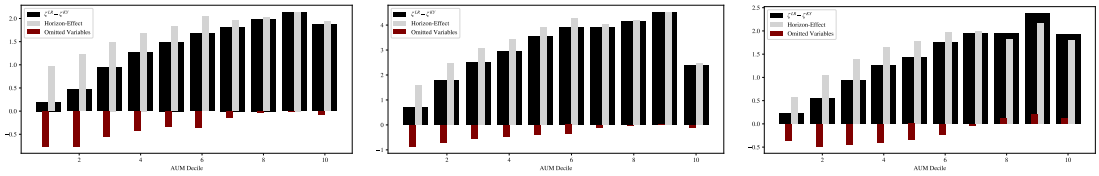
B.4 Results for different instruments

Figure B.2 illustrates the robustness of the results with respect to different instruments. Regardless of the instrument, investors are inert with an average $\rho_{i,t}$ of around 0.65-0.85 (see Table 2). Regardless of the instrument, the coefficient difference between holding and trading $\zeta^{SR}_{i,t} - \zeta^{KY}_{i,t}$ is positively correlated with investor size (Panel b). Short-run multipliers greatly exceed long-run multipliers and are of the same magnitude across instruments (Panel c). The reversal implied by the long-run elasticities is in line with

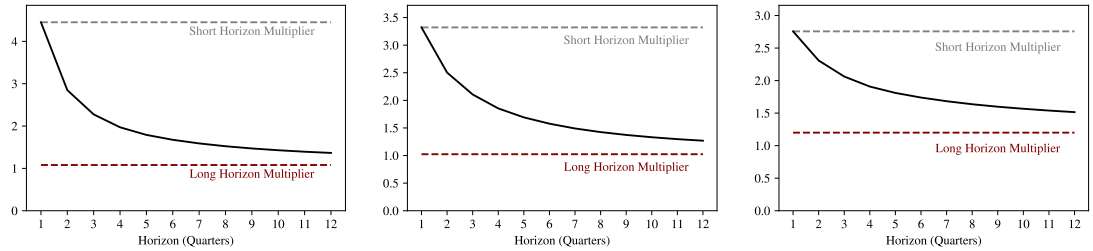
the impact reversal from flow-induced trades for all instruments (Panel d). Last, conditioning on model-based reversal $\mathcal{R}_t(n)$ improves unconditional momentum returns for all instruments.

Figure B.2: **Robustness**.
The figure reports all results for different instruments for $q_{i,t-1}(n)$. The left-most column reports the baseline OLS results without instruments for $q_{i,t-1}(n)$. The middle column reports results for the model-based instrument. The right-most column reports results for using long-run lagged ownership as an instrument. Panel a) plots the distribution of the coefficient difference between holding and trading $\zeta_{i,t}^{SR} - \zeta_{i,t}^{KY}$. The remaining panels plot all results from Section 6 for the different instruments: Panel b) plots the median short- and long-run multipliers, Panel c) reports the out-of-sample test, and Panel d) plots momentum returns conditional on model-based reversal $\mathcal{R}_t(n)$.

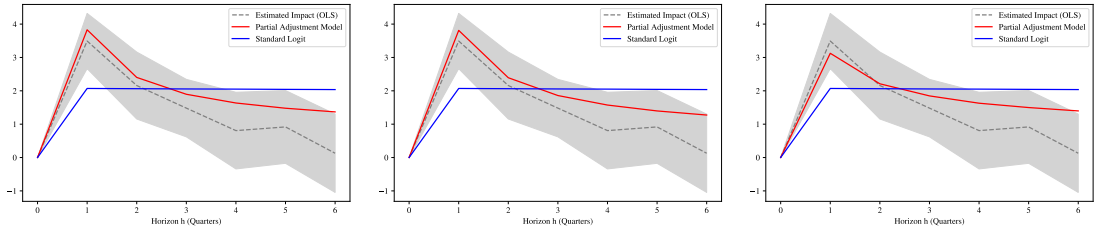
(a) Distribution $\zeta_{i,t}^{SR} - \zeta_{i,t}^{KY}$



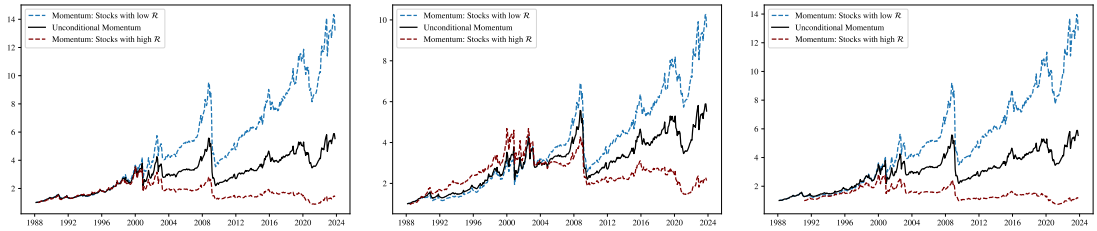
(b) Short- versus Long-run Impact



(c) Predicting the Impact of Flow-Induced Trades



(d) Inertia $\mathcal{R}_t(n)$ and Momentum Returns



Appendix C Simulation

To illustrate the effects of the omitted variable bias (b) and trading horizon (h), I provide simple simulations from two stylized investors $i = \{b, h\}$. Both investors enter the market in $t = 2010$ with \$1 under management. Investor b (for bias) is used to illustrate the potential omitted variable bias from unobservable portfolio tilts in the cross-section. Investor h (for horizon) is used to illustrate the effects of trading horizons on the estimated elasticities. I simulate their portfolios using observable market data from 2010 to 2024 and estimate their elasticities cross-sectionally in reduced form (i.e. from changes) and structurally (i.e. from levels). Formally, the regression specifications are given by $\log w_{i,t}(n) = \beta_{i,t} p_t(n) + c_t + \epsilon_{i,t}(n)$ and $\Delta \log w_{i,t}(n) = \beta_{i,t} \Delta p_t(n) + c_t + \epsilon_{i,t}(n)$. To this end, I construct a balanced panel of stock-quarter observations by taking the subset of all US stocks as of January 2024 that were already publicly traded in 2010. At every quarter t, the cross-section of both investors' (log) portfolio weights can be decomposed linearly

$$\log w_{i,t}(n) = \beta_i p_t(n) + \epsilon_{i,t}(n) \quad \text{for} \quad t \geq 2010. \tag{54}$$

where $\epsilon_{i,t}(n)$ are unobserved demand shifts coming from e.g. differential starting portfolios or dependence on lagged weights (inertia). β_i measures the slope of the demand curve, i.e. the cross-sectional relationship between portfolio weights and prices. The OLS estimate of $1 - \beta_i$ is an unbiased estimator of the elasticity of investor i as long as $\epsilon_{i,t}(n)$ and $p_t(n)$ are cross-sectionally uncorrelated.

C.1 Investor b: Omitted variables Bias

Investor b enters the market with portfolio tilted towards technology stocks. For simplicity, assume that as of $t = 2010$ the technology tilt perfectly coincides with market equity. Thus she enters the market with a value-weighted portfolio. Formally, her starting portfolio is given by $\log w_{b,2010}(n) = p_{2010}(n) - p_{2010}$ where $p_{2010} = \log \sum_{n=1}^{N} P_{2010}(n)$ is the log total market cap as of 2010. She then has no other demand shifters, and henceforth trades *quarterly* against price moves with an elasticity of 1, thereby keeping her portfolio weights constant at their starting level $w_{b,2010}(n)$. In the generic demand curve (54), investor b's starting portfolio implies that $\epsilon_{b,t}(n) = p_t(n) - p_{2010}$ and $\beta_{b,t}(n) = 0$. Estimating elasticities from portfolio holdings in levels leads to a coefficient bias $\zeta_{b,t}^{\text{Levels}} - \zeta_{b,t}$ that is driven by how cross-sectionally correlated current market equities are with the starting

values $p_{2010}(n)$ in 2010. Formally, the coefficient bias is given by

$$\zeta_{b,t}^{\text{Levels}} - \zeta_{b,t} = \frac{\text{cov}_t(p_{2010}(n), p_t(n))}{\text{var}_t(p_t(n))}, \tag{55}$$

which tends to be positive because cross-sections of market equities are quite persistent. When $t = 2010$ the coefficient bias is exactly 1. In first differences however, $\Delta\epsilon_{b,t}(n) = 0$, and $\zeta_{b,t}^{\text{Changes}} - \zeta_{b,t} = 0$. Figure C.3 illustrates the relative advantages of the estimation in levels and changes. Investor b's elasticity in levels is "perfectly" biased. She starts with a value-weighted portfolio that is *cross-sectionally* passive but trades *actively* against price movements. Over time, the initially value-weighted portfolio moves further and further away from true market prices resulting in a smaller coefficient bias. However, cross-sections of market equities are so persistent, that a substantial bias remains even after 12 years. In first differences, the bias from the omitted variable (the value-weighted starting portfolio) is exactly removed.

C.2 Investor h: Trading Horizon

Investor h enters the market with an equal-weighted portfolio. Hence, there is no omitted variable bias. However, instead of assuming that she trades quarterly against price movements, she slowly trades back to her target portfolio at a rate $1 - \rho$. Formally, her starting portfolio is given by $\log w_{h,2010}(n) = \log(1/N)$. Because she slowly trades back to her target portfolio at a rate $1 - \rho$, her portfolio weights are given by $\log w_{h,t}(n) = \rho \log w_{h,t}^p(n) + (1 - \rho)\log(1/N)$, where $w_{h,t}^p(n) = w_{h,t-1}(n)\frac{P_t(n)}{P_{t-1}(n)}$ is the passive buy-and-hold portfolio the investor would hold in the absence of trading. Solving this recursive equation yields $\log w_{h,t}(n) = \sum_{s=1}^{t} \rho^{t-s+1} \log \frac{P_s(n)}{P_{s-1}(n)} - N$. For simplicity, consider the simple case of a single price change directly after $t = 2010$ and no price changes thereafter. The portfolio weight at any time $t \geq 2010$ is given by $\log w_{h,t}(n) = \rho^{(t-2010)}(p_t(n) - p_{2010}(n)) - \log N$ and the coefficient bias from portfolio holdings in levels is

$$\zeta_{h,t}^{\text{Levels}} - \zeta_{h,t} = \rho^{(t-2010)}\left(1 - \frac{\text{cov}_t(p_{2010}(n), p_t(n))}{\text{var}_t(p_t(n))}\right). \tag{56}$$

As long as prices remain close to their starting values, i.e. $\frac{\text{cov}_t(p_{2010}(n), p_t(n))}{\text{var}_t(p_t(n))} \approx 1$, the level specification gives an unbiased estimate of the true long-run elasticity $\zeta_{h,t} = 1$. Similarly, if prices change but $\rho \to 1$, the level specification remains unbiased. Estimating

elasticities from trades, however, greatly understates the true long-run elasticity. To this end, note that changes in holdings are given by

$$\Delta \log w_{h,t}(n) = \rho \Delta p_t(n) + (1-\rho) \sum_{s=1}^{t-1} \rho^{t-s} \Delta p_s(n) \tag{57}$$

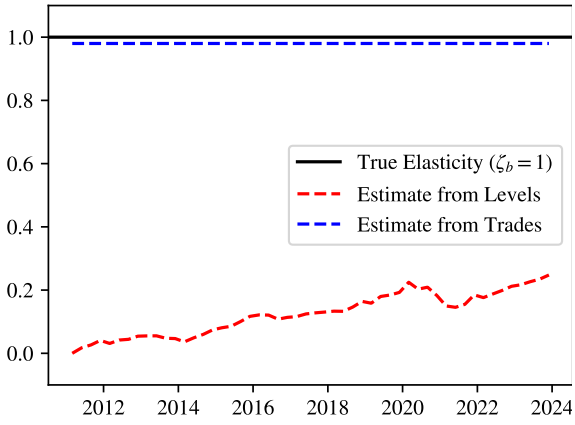
For simplicity, assume that there is no cross-sectional correlation in returns. Cross-sectional regressions of $\Delta \log w_{h,t}(n)$ onto $\Delta p_t(n)$ therefore yield an implied elasticity of $\zeta_{h,t}^{\text{Changes}} = 1 - \rho$. The coefficient bias is given by

$$\zeta_{h,t}^{\text{Changes}} - \zeta_{h,t} = \rho \tag{58}$$

The slower the investor trades, the larger the coefficient bias. In the limit, as $\rho \to 1$, the investor trades infinitely slow and the estimation in changes is "perfectly" biased. The right panel of Figure C.3 shows that the level specification captures the true long-term elasticity of 1. The specification in changes, however, only captures the quarterly short-term responses of $1 - \rho$. Therefore, while investor b's elasticity is better captured using trades, investor h is better captured using levels.

Figure C.3: **Omitted Variable Bias versus Trading Horizon**.
The figure plots the estimated elasticities from simulated portfolio data using true market prices from 2010 to 2024. I simulate the portfolio holdings $q_{i,t}(n)$ for two investors $i = b, h$ with equal (long-run) elasticities $\zeta_b = \zeta_h = 1$. I estimate their elasticities quarterly from cross-sectional regressions in levels, i.e. $q_t^s(n) = -\zeta_t^s p_t(n) + c_t + \epsilon_t(n)$, and quarterly cross-sectional regressions in changes, i.e. $\Delta q_t^s(n) = -\zeta_t^s \Delta p_t(n) + c_t + \epsilon_t(n)$. Panel a) plots the estimated elasticities for investor b, who enters the market in 2010 with a value-weighted portfolio and trades back to her starting portfolio quarterly. Panel b) plots the estimated elasticities for investor h, who enters the market in 2010 with an equal-weighted portfolio and gradually trades back to her starting portfolio with speed $\rho = 0.95$.

(a) Omitted Variable Bias



(b) Trading Horizon

