NBER WORKING PAPER SERIES

HOW (NOT) TO IDENTIFY DEMAND ELASTICITIES IN DYNAMIC ASSET MARKETS

Jules H. van Binsbergen
Benjamin David
Christian C. Opp

Working Paper 34528
http://www.nber.org/papers/w34528

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025, Revised June 2026

The authors thank Yu An (discussant), Aditya Chaudhry (discussant), Carter Davis (discussant), Julie Fu (discussant), Willie Fuchs (discussant), Xavier Gabaix, Zhiguo He (discussant), Federico Mainardi (discussant), Aaron Mora (discussant), and seminar participants at Bocconi University, Boston University, Dartmouth College, HEC Montréal, the Wharton School, the University of North Carolina, and the University of Rochester, as well as participants at the ASU Sonoran Winter Finance Conference, the Duke-UNC Asset Pricing Conference, the Financial Research Association Conference, the Finance Theory Group Meeting at UT Austin, the FIRS Conference, the GEA Annual Conference, the HEC-McGill Winter Finance Workshop, the MFA Annual Meeting, the Macro-Finance Tepper-LAEF Conference, the SFS Cavalcade North America, and the SWFA Conference for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Jules H. van Binsbergen, Benjamin David, and Christian C. Opp. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How (Not) to Identify Demand Elasticities in Dynamic Asset Markets
Jules H. van Binsbergen, Benjamin David, and Christian C. Opp
NBER Working Paper No. 34528
December 2025, Revised June 2026
JEL No. E10, G10

ABSTRACT

We evaluate approaches to estimating demand elasticities in dynamic asset markets, both theoretically and empirically. We establish strict necessary conditions that instrumented asset-price variation must satisfy for valid identification. Commonly used instruments violate these conditions, yielding estimates that are largely unrelated to true price elasticities and may even have the wrong sign. We derive transparent formulas for the biases that arise when instrumented price shifts persist or build up and propose an approximate bias-correction factor. Our dynamic general-equilibrium analysis characterizes the mapping from supply shock processes to price-shift processes, revealing severe shortcomings of standard multiplier calculations in multi-period settings.

Jules H. van Binsbergen
University of Pennsylvania
and NBER
julesv@wharton.upenn.edu

Benjamin David
University of Pennsylvania
bdavid97@wharton.upenn.edu

Christian C. Opp
University of Rochester
Simon Business School
and NBER
opp@rochester.edu

1 Introduction

A longstanding literature in finance, beginning with Shleifer (1986) and Harris and Gurel (1986), studies whether demand curves for financial assets are downward sloping and whether the time horizon of measurement matters. More recent work applies tools from the structural industrial organization (IO) literature to estimate institutional demand for long-lived securities. This line of research reports estimated demand elasticities close to zero and contrasts them with theoretical benchmarks that are several orders of magnitude larger. On this basis, it concludes that asset markets are surprisingly inelastic and subject to substantial institutional or behavioral frictions (see, e.g., Koijen and Yogo, 2019, 2021; Gabaix and Koijen, 2023). Quantifying the relative importance of such frictions and other forces shaping the relation between quantities and prices is of first-order importance for understanding allocative efficiency in financial markets and for evaluating the impact of policies such as central bank interventions.

In this paper, we assess the identification strategies underlying existing demand-elasticity estimates and the conclusions drawn from them in settings where agents trade assets dynamically. We show that the instruments used in the existing literature generate large and systematic biases. In fact, existing estimates are largely unrelated to the true price elasticities and the structural parameters they are meant to capture, such as the severity of frictions. A central theme of our analysis is the disconnect between return opportunities and price levels: even highly elastic traders who face no frictions and aggressively trade on return opportunities may trade little, if at all, against persistent price-level shifts. Beyond our analysis of demand-elasticity identification, we examine the relation between quantities and equilibrium prices in a general-equilibrium setting with dynamic trade. We employ this framework to study dynamic asset-market interventions, such as those implemented through central-bank asset-purchase programs, and characterize the intervention paths required to achieve intended price support. This analysis reveals significant shortcomings in the standard multiplier calculations used in the literature to address such questions.

The biases in price elasticity estimates in the existing literature arise from the following conceptual tension. In standard IO settings, a price elasticity measures how demand responds to a marginal price change while holding the product's intrinsic characteristics fixed. In financial markets, however, investors can generally resell securities in the future subject to relatively low trading frictions, so the distribution of future resale prices, $\{P_\tau\}_{\tau>t}$, is itself an important asset characteristic. Whereas exogenously raising today's price increases the *cost* of the asset, raising tomorrow's price increases the *benefit* of owning it, because the investor can sell at that higher price.

Correspondingly, the exclusion restriction does not merely require that an instrument shifting the contemporaneous price at time t leave future dividends unchanged; it also requires that the instrument leave unchanged the distribution of future resale prices. In particular, when price

shifts exhibit momentum or persistence, the resulting estimates can no longer be interpreted as price elasticities. Rather than reflecting structural frictions or preferences, these estimates capture simultaneous shifts in both the costs and benefits of owning the asset. These concerns are not merely theoretical: because financial institutions actively engage in dynamic trading strategies, future resale values are an essential component of the payoff from current asset purchases.

The resulting bias reflects the conflation of two forms of dependence that demand-function estimation must identify separately: the response of demand to the current price, holding fixed the state of the economy, and the dependence of the demand function on those state variables, which summarize the relevant properties of the current environment. Both are economically relevant, but they play different roles. The current-price slope is the price elasticity used for market clearing: it is the slope of the demand schedule with respect to the current price, conditional on the state. By contrast, state dependence captures how the demand schedule itself varies with the economic environment, in particular through conditioning information about the distribution of future dividends and resale values. Thus, demand systems can in principle be used to study the effects of persistent price movements, but doing so requires distinguishing the role of state variables in the demand function from the current-price slope. This distinction is exactly what the exclusion restriction is meant to ensure.

Our contribution is both theoretical and empirical. On the theory side, we develop a model of dynamic trade to identify necessary conditions for valid price elasticity identification using either time-series or cross-sectional holdings data. Specifically, we show that price variation induced by valid instruments must generally (i) affect prices for at most one trading period, and (ii) be unanticipated. While the absence of anticipation is particularly critical for time-series identification (e.g., using index reconstitution event studies),[1] full resolution of price shifts within a single trading period is necessary for both time-series and cross-sectional approaches.

Using our model, we derive formulas that quantify the biases in elasticity estimates when these conditions are violated, covering three forms of price-shift dynamics: (1) resolution of price shifts with a constant per-period probability, (2) no resolution in the first trading period, and (3) predictable price build-up, as in the case of momentum (see Binsbergen et al., 2023). Moreover, we propose an approximate bias-correction factor that maps the estimated elasticity to the true price elasticity, thereby facilitating the recovery of structural parameters with external validity, and establish conditions under which such correction is feasible.

We show that the relative bias in the estimated elasticity approximately equals the expected

[1] When price changes induced by index inclusion effects are either anticipated or resolve over longer horizons, the resulting elasticity estimates are generally biased. Early studies, such as Harris and Gurel (1986), documented mean reversion within two weeks, but more recent work suggests that the index inclusion effect has either disappeared for the S&P 500 index (Greenwood and Sammon, 2025) or exhibits longer-lasting effects for other indices. See Chang et al. (2015); Beneish and Whaley (1996); Lynch and Mendenhall (1997); Wurgler and Zhuravskaya (2002).

fraction of the price shift that resolves by an investor's first feasible trading date after portfolio data are observed. If a price shift does not resolve by an investor's next feasible trading date, the estimated elasticity will be approximately zero even if the true elasticity is infinite. That is, the estimate is entirely uninformative: infinite frictions and zero frictions yield identical estimates, implying that a bias correction is infeasible. Naturally, an estimate that cannot distinguish between these extremes does not have external validity and cannot be employed in counterfactual analyses. Yet this issue is pertinent to the instrument proposed by Koijen and Yogo (2019) (KY19 hereafter), which, as we show, induces price shifts that do not resolve by investors' next feasible trading date.

How misleading elasticity estimates can be when price-shift dynamics violate the exclusion restriction is also illustrated by the momentum case: the estimated elasticity takes the opposite sign of the true elasticity. When the variation inducing a positive price shift also predicts higher future returns, agents rationally tilt *toward* the security, creating the appearance of a negative price elasticity.

Crucially, we show that the horizon over which price-shift persistence or build-up must be measured to evaluate instrument validity is a structural characteristic of the investor, determined by trading technology, regulations, or behavioral frictions, not by the econometrician's choice of data frequency. For instance, a price shift that persists for two weeks before fully resolving might be viewed as highly transitory in many econometric contexts. Yet for an institutional investor that can trade intraday, such a shift is effectively fully persistent at the relevant horizon.

The following example illustrates the external validity problems of estimates that violate the conditions for valid price elasticity identification we establish. Suppose the Federal Reserve unexpectedly and credibly announces that it will do "whatever it takes" to raise the price of an asset by 1% for one year, relative to the counterfactual price path absent intervention. After one year, the program is gradually unwound over the course of the following week. To streamline the argument, assume the asset pays no dividends over the entire intervention horizon and that investors have log utility.

How does this exogenous increase in prices affect investors' optimal portfolio weights on impact? Not at all. Once the surprise announcement is made and the price has shifted upward, the *return* distribution going forward is unaffected by the intervention until the program is unwound one year later. As a result, both static cross-sectional and event-study approaches using data on portfolio weights around the time of the event yield elasticity estimates of zero.

But does this imply that asset markets are inelastic? Not at all. The crux is that the announcement alters the asset's future resale value without changing the distribution of returns going forward — it does not speak to how elastic investors are. To see why, note that investors will in fact elastically tilt their portfolios away from the asset once the program is unwound, responding to the abnormal price decline of 1% over the ensuing week. For the program not to unravel by backward

induction, the Federal Reserve must therefore purchase large quantities of the asset at that time, reflecting log investors' true, highly elastic behavior.[2]

This example illustrates that in dynamic trading environments, variation in holdings observed around the time of an exogenous price shift may provide no meaningful information about either the true demand elasticity or the quantity the Federal Reserve must purchase. Both static cross-sectional identification and event-study approaches fail. The example also highlights that the identification conditions we characterize cannot be relaxed simply by the econometrician choosing a lower data sampling frequency for asset holdings, as this frequency is distinct from the one at which agents can adjust their portfolios. Investors' portfolio shares are constant throughout the whole year and only briefly shift for one week before reverting to their original level.[3]

Our empirical results are closely connected to these issues. In particular, we find that cross-sectional variation in KY19's instrumented price variable, $\widehat{me}$ (instrumented market equity), does not satisfy the identification conditions we derive. Price shifts induced by the instrument at time t do not resolve even after a month. In fact, point estimates suggest that they might exhibit momentum. Using such variation as an instrument mechanically yields elasticity estimates that are close to zero (entirely unrelated to the severity of frictions and true elasticities), or even negative, the opposite sign of the true price elasticity. These estimates therefore do not recover the current-price slope of the demand schedule, which is required for market clearing and hence for endogenizing equilibrium prices in counterfactual analyses based on demand systems. As discussed above, this conclusion is not in tension with using demand systems to study persistent price shocks; rather, in a correctly specified demand function, such shocks must be captured through state variables, not through the variation used to identify the current-price slope.

Consistent with our model's prediction that estimated elasticities are negative when price shifts build up, we find that a large fraction of the price-elasticity estimates in the KY19 framework are negative when the estimation is not constrained to produce positive values.[4] The estimates are orders of magnitude lower than the theoretical benchmarks to which they have been compared

[2] A closely related policy intervention is a forward commitment by the Federal Reserve to undertake any purchases necessary, one year from now, to lift the asset's price by 1% for a single trading period, *relative to the case of no intervention*. By backward induction, prices will already rise at the time of the announcement, as agents anticipate being able to resell the asset at a 1% higher price at the end of the year. By construction, this intervention involves no trade between the Federal Reserve and other market participants at the time of the announcement. As a result, the incorrectly inferred unit price elasticity of demand is zero.

[3] Similarly, our points about elasticity identification are not resolved by appealing to the notion of long-run elasticities. In standard product markets, a long-run elasticity considers a longer time horizon over which consumers are able to adjust, for example, by switching to alternative products. This typically leads to *more* elastic responses, not lower elasticities. The key difference is that whereas a regular good can plausibly maintain its intrinsic characteristics in the face of persistent price changes, this is infeasible for financial assets, since their resale value is itself a central demand-relevant characteristic.

[4] Specifically, this occurs when we relax KY19's assumption that the estimated beta on instrumented market equity in the demand equation is bounded above by one.

in the literature simply because the instrument violates the exclusion restriction.[5] Indeed, as our theoretical analysis shows, if the price shift fully persists in the first trading period of institutional investors, the estimates are in fact entirely uninformative. In general, these estimates measure how demand responds to a stochastic price-shift process induced by the instrument. That is, they identify what we term a shifter-process elasticity, which is specific to the underlying shock process and does not represent a structural object.

Our results also emphasize that the concept of exogeneity typically invoked in economic models is insufficient for the proper identification of demand elasticities in financial markets. Even if a shock that shifts prices would be regarded as truly exogenous within the context of a model (e.g., if induced by a supply shock randomly chosen by nature), it still cannot be used to identify demand elasticities if it alters the intrinsic characteristics of the asset. However, an exogenous shock that affects prices across multiple trading dates induces exactly such a change. As a result, it introduces or alters dynamic state variables in traders' optimization problems, causing elasticity estimates to reflect changes in these state variables rather than the contemporaneous price.

Our analysis of asset demand in a dynamic general equilibrium setting yields sharp predictions for the price effects of dynamic asset-purchase programs and other shocks that affect asset supply over time, such as demand from irrational retail traders or liquidity traders. The results have important implications for how easily such shocks can generate substantial *price-level* shifts even when the rest of the market is populated by fully rational traders who can trade frictionlessly. The reason is that, if such shocks are expected to shift prices persistently, they have little effect on the return distribution, which is the object that governs the portfolio choices of rational investors.

It would therefore be a major misconception to infer that, because elastic traders largely eliminate abnormal return opportunities, they must also eliminate mispricing in the sense of price-level dislocations (see also Binsbergen et al., 2023). This distinction is economically important because price levels, rather than abnormal returns, are the object that matters for real allocative efficiency (see Binsbergen and Opp, 2019). Thus, even fully rational and frictionless financial market participants may have little incentive to correct the price-level dislocations that matter for real allocations.[6]

These principles also have direct implications for the design of dynamic policy interventions, such as those implemented by central banks: the required intervention quantities depend critically on the persistence of the targeted price support. For instance, under a constant per-period reso-

[5] Koijen and Yogo (2021) and Gabaix and Koijen (2023), for example, compare their estimates to the theoretical price elasticity implied by Petajisto's (2009) calibration of the capital asset pricing model (CAPM), which yields price elasticities for individual stocks in excess of 5,000.

[6] This finding contrasts with Gabaix and Koijen (2023), who argue that market participants' limited responsiveness to persistent price shifts must be due to substantial frictions in financial markets, such as fixed-share portfolio mandates and inertia.

lution probability, increasing the expected duration of a 1% price support from one week to one year reduces the required residual supply intervention at inception by 98%. This large difference arises because rational, optimizing traders sell the asset much more aggressively when the induced price shift is expected to reverse quickly. Required intervention quantities also differ substantially depending on whether policy commitments have a fixed expiration date or resolve randomly.

More generally, we characterize the entire path of dynamic asset-market interventions required to sustain a given Markovian price-support process. This analysis reveals significant limitations of conventional multiplier calculations. Even in an economy with only a single risky asset and a safe asset, there is no single multiplier that characterizes the relationship between purchase quantities and the resulting price-path response, including when the targeted price effect is permanent. Moreover, purchase programs that solely promise future interventions naturally generate infinite announcement multipliers, the mirror image of an erroneously estimated zero price elasticity. Our findings therefore provide constructive guidance for the design of central-bank asset-purchase programs, such as quantitative easing, by clarifying how the timing and stochastic properties of interventions shape their dynamic price effects.

Our contribution is related to several strands of literature, including the heterogeneous-agent general-equilibrium literature, recent work on demand-system asset pricing, and the dynamic-demand literature in industrial organization.

The issues we examine naturally connect to the heterogeneous-agent asset-pricing literature, in which the cross-sectional allocation of assets, consumption, or institutional capital is a state variable, and equilibrium prices are functions of that state (see, e.g., Dumas, 1989; Basak and Cuoco, 1998; Heaton and Lucas, 1996; He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014; Bhamra and Uppal, 2014; Gârleanu and Panageas, 2015; Borovička, 2020). The KY19 framework is a multi-period general-equilibrium model with heterogeneous investors. However, rather than solving for the equilibrium state variables that govern price and return dynamics, potentially including the distribution of assets under management (AUM), KY19 specify agents' demands and beliefs about returns as functions of exogenously stipulated asset characteristics. Our results reveal the econometric consequence of this shortcut: when the instrumented price variation is persistent, the price-elasticity coefficient identified by the KY19 demand system reflects demand variation induced by movements in state variables, rather than the current-price slope needed for market clearing.

This issue also arises in Davis et al. (2025), who argue that, relative to models such as Petajisto (2009), incomplete price pass-through in frictionless portfolio-choice models reduces price elasticities from around 7,000 to 800, a factor of 9, and that unspanned returns reduce them further to approximately 5. The object they characterize is again a derivative of demand with respect to state variables or deep parameters that shift the distribution of future prices. Such derivatives must

not be used in place of the price elasticity entering the demand functions of the demand-systems literature following KY19, since market clearing occurs in the current price, conditional on the state, not in state variables.

He et al. (2025) also examine asset-demand estimation in a dynamic general equilibrium setting, but their analysis pertains to a different object: agents' optimal quantity responses to changes in expected returns. They emphasize the important fact that observed demand slopes with respect to expected returns can differ from true slopes because exogenous supply shifts plausibly affect return volatility and comovement with future investment opportunities. Our object of interest, the price elasticity of demand, differs from theirs in important ways, as a shock's impact on the price level is distinct from its effect on returns over the next period. We show that this disconnect can even cause price-elasticity estimates to take the opposite sign of the true elasticity, particularly in cases of price-shift build-up. Moreover, the key identification condition we derive for price elasticities—full resolution of price shifts within one trading period—is generally not required when estimating demand slopes with respect to expected returns. In sum, our analyses pertain to distinct objects and provide complementary insights for the literature on demand estimation.

Gabaix and Koijen (2023) examine multipliers both empirically and theoretically. They explain the large magnitudes of empirically estimated multipliers, particularly for the stock market, within a dynamic general equilibrium framework by introducing two frictions in capital allocation through institutions: fixed-share mandates and inertia. Moreover, Gabaix and Koijen (2023) show that persistence in flows leads to higher multiplier estimates and define a "market multiplier" as the multiplier associated with a fully persistent inflow. The authors argue that rational models without frictions predict that this market multiplier is around $1/20$, whereas empirical estimates in the literature, as well as those obtained using their approach, yield values between 1.5 and 5.

Our paper contributes to this literature on multipliers by establishing several key results. First, we show that multipliers used to describe price impact at the inception of multi-period purchases are insufficient and misleading in characterizing the relationship between quantities and prices. Even in the case of the most stylized dynamic intervention, one that induces a fully permanent price shift by a fixed percentage, there is no single multiplier that would describe the purchases needed for such a shift: A one-time quantity purchase cannot shift the price permanently by a fixed percentage, nor can a permanently higher price level be sustained by any single purchase. Going beyond the ill-defined notion of multipliers, we provide formulas that characterize the actual path of incremental holdings required to sustain a given path of price shifts.

Second, we show quantitatively that permanent interventions can plausibly exhibit pseudo multipliers of around 2 or higher without requiring additional frictions. Moreover, estimates should be expected to take high values, potentially unbounded, in practically relevant settings involving position build-up, such as central bank asset purchase programs. Importantly, such evidence on

multipliers is fully consistent with market participants being highly elastic, with no meaningful frictions limiting their quantity responses.

It is further important to distinguish the price elasticities estimated in the demand-systems literature following KY19 from price-impact measures in the market microstructure literature. In canonical microstructure models, order flow moves prices because it is potentially informative about future payoffs (e.g., Kyle, 1985; Glosten and Milgrom, 1985).[7] Empirical measures of price impact similarly capture the extent to which trades forecast future price revisions (e.g., Hasbrouck, 1991). This is a different object from the current-price slope required in demand-system estimation. The exclusion restriction in the KY19 framework requires identifying variation in the current price that does not also shift demand through payoff-relevant channels. By contrast, price impact in asymmetric-information models is exactly the equilibrium response of prices to order flow that conveys information about those future payoffs.

Our paper also relates to the dynamic-demand literature on durable and storable goods, which has long recognized that static elasticity estimates can be biased when consumers make choices in multi-period settings (see, e.g., Hendel and Nevo, 2006; Chen et al., 2008; Schiraldi, 2011; Gowrisankaran and Rysman, 2012; Fukasawa, 2024). Several papers quantify these biases by estimating or simulating a dynamic model and comparing it with a misspecified static counterpart. However, this literature generally does not provide a closed-form or otherwise transparent analytical expression for the relative bias induced by instrumented price variation with persistent dynamics, including price momentum, nor does it derive the corresponding debiasing factor that we obtain in our setting. These formulas emerge in the low-trading-friction case of our continuous-time setting, which is a natural benchmark for many financial markets because investors can often rebalance much more quickly than consumers can adjust durable-goods holdings. More generally, we show that increased trading frictions attenuate the bias: the longer investors must wait before they can trade again, the longer an instrument-induced price shift can persist without violating the identification condition. To our knowledge, the durable-goods literature has not derived this property, the closed-form bias formula, or the corresponding debiasing factor for static elasticity estimates.

Finally, our key results pertaining to dynamics are distinct from the channels emphasized in the literature on the implications and validity of different substitution structures assumed in static models (Fuchs et al., 2025a,b; Kocherlakota, 2025; Haddad et al., 2025; Koijen and Yogo, 2025; An and Huber, 2025). Whereas these papers examine limitations and biases that may emerge from assumptions regarding cross-sectional substitution, our contribution is to isolate dynamic violations of identification assumptions and to provide direct empirical tests to verify them. Going

[7] In the limiting case in which trading motives are driven only by payoff-relevant information, the no-trade theorem of Milgrom and Stokey (1982) applies under its standard common-knowledge and rational-expectations assumptions.

beyond elasticity estimation, our analysis further characterizes the properties of dynamic purchase programs needed to implement dynamic price-support policies, such as quantitative easing. While Fuchs et al. (2025a) focus on cross-sectional substitution and spillovers, the authors also consider a two-period extension where agents' tastes might drive the future price. In that context, they conclude that observed demand elasticities do not reveal whether an investor is buying based on her own tastes or her beliefs about market-wide tastes.

2 Model

We consider a continuous-time environment with two groups of agents: an investor group that explicitly maximizes CRRA utility and an "outside group" that introduces exogenous residual supply shocks. The latter group can be interpreted, for instance, as noise traders or central banks. We begin by specifying the CRRA investors' trading and consumption opportunities conditional on postulated equilibrium asset-price dynamics. We then characterize the outside group's residual supply dynamics. We intentionally reverse-engineer supply shocks to obtain a minimally complex state space and a clean characterization of the shock dynamics required either to identify demand elasticities or to induce targeted equilibrium price effects.

Preferences. There is a measure one of CRRA investors whose utility from consumption is given by

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}, \tag{1}$$

where $\gamma > 0$ is the coefficient of relative risk aversion. The subjective discount rate applied to future utility is denoted by ρ.

Assets. Agents can invest in multiple risky assets $j = 1, ..., J$ and in a risk-free asset with return r_f. We postulate that absent price-shifter shocks detailed below, prices of the risky assets each follow:

$$\frac{d\tilde{P}_{j,t}}{\tilde{P}_{j,t}} = \mu_j dt + \sigma_{j,A} dB_{A,t} + \sigma_{j,I} dB_{j,t}, \tag{2}$$

where $dB_{A,t}$ represents a common "aggregate" shock and $dB_{j,t}$ an asset-specific shock. The Brownian motions $dB_{A,t}$ and $\{dB_{j,t}\}_{j=1}^{J}$ are mutually independent. The actual trading prices can be further affected by price-shifter shocks. Specifically, the equilibrium trading price $P_{j,t}$ of asset j

satisfies

$$P_{j,t} = e^{\beta_j s_t} \tilde{P}_{j,t}, \text{ with } s_t \in \Omega \tag{3}$$

where $\beta_j s_t$ represents a log price shift applying to asset j, as similarly defined in Binsbergen and Opp (2019) and Binsbergen et al. (2023), Ω denotes the discrete set of possible shifter states s_t, and β_j governs the exposure of asset j to the shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$. In Section 3.3, we formally derive the outside group's residual supply dynamics that yield the postulated prices $P_{j,t}$ in equilibrium.

Event dates. To allow for well-defined demand elasticity estimation using discrete price shifts in a continuous-time setting, we introduce discrete event dates. At these dates, price shifts may build up or change, and agents may consume (disperse capital). While we initially consider the case in which agents can trade continuously, we also analyze a setting in which trading is restricted to these discrete event dates, reflecting frictions in the trading process. Event dates occur at the jump times τ_n of a Poisson process N_t with intensity $\lambda > 0$. The parameter λ allows us to flexibly adjust the average frequency at which these events occur.

At each event date, the following ordered events occur:

1. **Shifter-state transitions.** The shifter state s moves to s' with the Markov transition probabilities $q(s'|s)$.
2. **Consumption.** With probability π_C, existing CRRA investors exit and are replaced by a new cohort with identical post-jump wealth W'. Exiting investors liquidate their assets and consume their wealth W'. At the same time, newly entering investors purchase assets with their wealth W'.
3. **Portfolio choice.** Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_J)^\top$ be the vector of portfolio shares invested in the J risky assets; the fraction in the risk-free asset is $\theta_0 = 1 - \mathbf{1}^\top \boldsymbol{\theta}$. CRRA investors choose their risky shares $\boldsymbol{\theta}_{\tau_n}$.

A portfolio policy $\boldsymbol{\theta}_t$ is $\{\mathscr{F}_t\}$-adapted and square-integrable. To guarantee positive wealth in the presence of price jumps, the admissible set in the case of continuous trade is defined as

$$\Theta(s) \equiv \{\boldsymbol{\theta} : \ 1 + \boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'} > 0 \ \forall s' \text{ with } q(s'|s) > 0\}, \tag{4}$$

where:

$$\kappa_{s,s'}^{(j)} \equiv e^{\beta_j (s'-s)} - 1, \qquad \boldsymbol{\kappa}_{s,s'} \equiv \left(\kappa_{s,s'}^{(1)}, \ldots, \kappa_{s,s'}^{(J)}\right)^\top. \tag{5}$$

3 Analysis

In this section, we first characterize CRRA investors' optimal portfolio policies under the postulated price dynamics. Second, we examine the difference between true demand elasticities and estimators derived from holdings data. Third, we close the model by characterizing the residual supply shock dynamics induced by the second group of agents that yield the postulated price dynamics in equilibrium.

3.1 Value Function and Optimal Trading

In this subsection, we characterize CRRA investors' value function and optimal trading behavior when they can trade continuously between Poisson event dates. In later subsections, we also consider the discrete-trading case, where trading is restricted to event dates only. The detailed analyses for that case are presented in the Appendix sections referenced in the propositions.

Over a small horizon dt, absent an event date, only diffusive risk is borne by the CRRA investors and wealth dynamics are given by:

$$\frac{dW_t}{W_t} = r_f dt + \sum_{j=1}^{J} \theta_{j,t}\Big[(\mu_j - r_f)dt + \sigma_{j,A}\, dB_{A,t} + \sigma_{j,I}\, dB_{j,t}\Big]. \tag{6}$$

Define the covariance matrix

$$\Sigma = \boldsymbol{\sigma}_A \boldsymbol{\sigma}_A^\top + \operatorname{diag}\big(\sigma_{1,I}^2, \ldots, \sigma_{J,I}^2\big), \quad \boldsymbol{\sigma}_A \equiv (\sigma_{1,A}, \ldots, \sigma_{J,A})^\top.$$

Then $Var\big[dW_t/W_t\big] = \boldsymbol{\theta}^\top \Sigma\, \boldsymbol{\theta} dt$. Upon the arrival of an event date, a jump in the shifter state s may occur. Finally, with probability π_C, a CRRA agent liquidates its assets and consumes post-jump wealth W'. Otherwise, with complementary probability $(1-\pi_C)$, the agent continues and a new portfolio is chosen.

Let $V(W,s)$ be the continuation value immediately after an event sequence. We obtain the Hamilton–Jacobi–Bellman (HJB) equation

$$\begin{aligned} 0 = \max_{\boldsymbol{\theta}} \Big\{ & -\rho V + V_W W\Big[r_f + (\boldsymbol{\mu} - r_f \mathbf{1})^\top \boldsymbol{\theta}\Big] + \tfrac{1}{2} V_{WW} W^2 \boldsymbol{\theta}^\top \Sigma\, \boldsymbol{\theta} \\ & + \lambda\, \mathbb{E}\big[\pi_C\big(u(W') - V(W,s)\big) + (1-\pi_C)\big(V(W',s') - V(W,s)\big)\big]\Big\}, \end{aligned} \tag{7}$$

where W' denotes post-jump wealth:

$$W' = W\Big[1 + \boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}\Big], \qquad s' \sim q(\cdot\,|s). \tag{8}$$

We conjecture and verify that the value function takes the form

$$V(W,s) = u(W)\cdot A(s). \tag{9}$$

Substituting $V_W = W^{-\gamma}A(s)$ and $V_{WW} = -\gamma W^{-\gamma-1}A(s)$ into (7) and dividing by $u(W)$ yields the following system of equations solved by the function $A(s)$ for $s \in \Omega$:

$$\begin{aligned} 0 = &-\rho A(s) + (1-\gamma)A(s)\left[r_f + (\boldsymbol{\mu} - r_f\mathbf{1})^\top\boldsymbol{\theta} - \tfrac{1}{2}\gamma\boldsymbol{\theta}^\top\boldsymbol{\Sigma}\boldsymbol{\theta}\right] \\ &+ \lambda\left[\sum_{s'} q(s'|s)\left[\pi_C + (1-\pi_C)A(s')\right]\cdot\left(1+\boldsymbol{\theta}^\top\boldsymbol{\kappa}_{s,s'}\right)^{1-\gamma} - A(s)\right]. \end{aligned} \tag{10}$$

Differentiating (10) with respect to each θ_j and dividing by $(1-\gamma)$ yields the first-order conditions for all $j = 1,\ldots,J$:

$$\begin{aligned} 0 = &A(s)\left[(\mu_j - r_f) - \gamma(\boldsymbol{\Sigma}\boldsymbol{\theta})_j\right] \\ &+ \lambda\sum_{s'} q(s'|s)\left[\pi_C + (1-\pi_C)A(s')\right]\cdot\kappa^{(j)}_{s,s'}\cdot\left(1+\boldsymbol{\theta}^\top\boldsymbol{\kappa}_{s,s'}\right)^{-\gamma}. \end{aligned} \tag{11}$$

When none of the assets are exposed to the shifter state ($\beta_j = 0$ for all j), equation (11) collapses to the standard Merton (1969) condition for optimal portfolio shares. For instance, in the case of a single risky asset, the optimal portfolio share then takes the familiar form $\theta_1^* = \frac{\mu_1 - r_f}{\sigma_1}$. More generally, however, additional terms related to shifter-state transitions arise. Importantly, analogous terms also appear under log utility (see Appendix A.1), so they should not be interpreted as merely reflecting intertemporal hedging motives. Rather, they capture changes in the shifter state s at event times that affect the return distribution over the next instant. Together, equation (10) and the J first-order conditions in (11) determine the function $A(s)$ (for $s \in \Omega$) and the optimal risky-share vector $\boldsymbol{\theta}^*(s)$.

Note that agents can perfectly maintain such optimal portfolio shares only in the continuous-trading case. In contrast, in the discrete-trading case analyzed below, portfolio shares evolve passively between event dates as prices change and can be rebalanced only on event dates.

3.2 Demand Elasticity Estimates and Biases

In this section, we examine under what conditions true demand elasticities can be accurately estimated based on holdings data and what biases arise if such conditions are violated.

True price elasticities and trading frictions. As a first step, we show that the frequency with which agents can trade has relevant implications for the magnitudes of true demand elasticities. In

our baseline setup, agents can trade continuously, which is a useful benchmark that represents a fully frictionless trading environment. Yet, given various practical limitations to truly continuous trade, we are also interested in examining elasticities when institutions can trade only at discrete dates, such as at a daily frequency. Correspondingly, we start by characterizing the true own-price elasticity in two cases: (i) continuous trade and (ii) discrete trade occurring only upon the event dates τ_n.

Consistent with the literature, we define the demand elasticity as the partial derivative of the log demand function at time t with respect to the $\mathscr{F}_t$-measurable log price, $\ln P_{j,t}$ (see, e.g., Koijen and Yogo, 2019). We generally focus on the portfolio share $\theta_{j,t}$ as the demand variable, but we will also analyze the quantity demanded, denoted by $n_{j,t}$.

Proposition 1. *Suppose there is one risky asset. In the case of continuous trading, the own-price elasticity is infinite,*

$$\varepsilon_C \;=\; -\frac{\partial \ln \theta_t^*}{\partial \ln P_t} = \infty. \tag{12}$$

In contrast, if agents can trade only upon the discrete event dates τ_n, the elasticity takes finite values

$$\varepsilon_D = -\frac{\partial \ln \theta_t^*}{\partial \ln P_t} < \infty, \tag{13}$$

and satisfies for short expected trading intervals: $\varepsilon_D/\lambda \approx \frac{[\pi_C+(1-\pi_C)\cdot A]}{(\mu-r_f)\cdot A}$ with

$$A \equiv \lambda\,\pi_C \cdot \left(\lambda\,\pi_C + \rho - (1-\gamma)\left[r_f + \frac{(\mu-r_f)^2}{2\gamma\sigma^2}\right]\right)^{-1}. \tag{14}$$

Proof. See Appendix A.2. □

Proposition 1 establishes that in a fully frictionless environment where agents can trade continuously, demand elasticities are infinite. Critically, this result holds even in the case of a "macro" demand elasticity where agents just trade in two asset classes that are not close substitutes, such as bonds and stocks. Conversely, the presence of a finite demand elasticity is direct evidence of trading frictions, such as the trade inertia considered in our model. Correspondingly, a correctly identified demand elasticity is in principle revealing the presence of such frictions.

The fact that demand elasticities are infinite in a frictionless world obtains for the following reason. An own-price elasticity quantifies how a marginal log-change to the contemporaneous price affects demand, keeping unchanged future prices at which an agent can resell the asset. In the case of continuous trade, this definition means that a log price shift $\xi \equiv d\ln P$ at time t needs

to resolve by the next instant $(t+dt)$, implying an incremental instantaneous resolution return $-\xi$ that lowers the expected rate of return relative to the baseline level of μ. This discrete change in the expected return of the risky asset from μ to $(\mu-\xi)$ implies a discrete (non-marginal) tilt in the optimal portfolio share θ^*.[8] Correspondingly, a *marginal* price shift causes *discrete* portfolio share adjustment, leading to the result of an infinite price elasticity under continuous trade.

In contrast, the second part of Proposition 1 shows that elasticities are finite when agents can trade only at discrete dates. The reason is that a marginal price shift $\xi = d\ln P$ at time t then resolves over a discrete period of time, leading to only a marginal adjustment in the trading period's expected return. Correspondingly, now a *marginal* price shift induces only a *marginal* tilt in the portfolio share θ, yielding a finite elasticity.

In sum, substantial differences between discrete trading and continuous trading emerge since the resolution of a given contemporaneous price shift has a more muted impact on returns per unit of time when agents can trade less frequently.

Shifter-process elasticities. We now characterize the elasticity estimates an econometrician obtains when using empirical variation that shifts not only the contemporaneous price, P_t, but also the distribution of future prices, $\{P_\tau\}_{\tau>t}$. In this context, the following lemma clarifies the mathematical nature of such objects.

Lemma 1. *Since investor demand θ_t is a function of $\mathscr{F}_t$-measurable objects, its partial derivatives are well-defined only with respect to such objects. Therefore, any elasticity that measures demand responsiveness to shifts affecting future prices $\{P_\tau\}_{\tau>t}$ represents a derivative of demand with respect to $\mathscr{F}_t$-measurable state variables or deep parameters that summarize the conditional distribution of those future prices.*

Lemma 1 highlights that objects measuring demand responsiveness to changes in the distribution of future prices, for example a price shock lasting for more than one trading period, are sensitivities of demand with respect to state variables or deep parameters. As such, these objects are distinct from a price elasticity, which is a *partial* derivative with respect to the current price, $\frac{\partial \ln \text{demand}_{j,t}}{\partial \ln P_{j,t}}$, conditional on the state and deep parameters. This distinction is essential for the literature estimating demand systems following KY19 and is at the heart of the corresponding exclusion restriction. The reason is that solving for equilibrium prices requires the slope of the demand schedule with respect to the current price, conditional on the state. This is the slope that the Walrasian auctioneer uses to clear the market. After all, the Walrasian auctioneer clears the market in the current price, conditional on the state of the economy. It does not clear the market in state variables.

[8] Note that the tilt is not infinite in magnitude due to the asset's risk exposures.

The need to identify this current-price slope in no way implies that demand systems cannot be used to investigate the effects of shocks that generate persistent price movements. Rather, such shocks are captured through the demand function's dependence on state variables.

A correctly specified demand system must therefore include both forms of dependence: the slope with respect to the current price, holding fixed the state, and the dependence on state variables that shift the economic environment. The exclusion restriction for demand-slope identification is not that these variables must be irrelevant for demand. Rather, to identify the price elasticity, the variation used for identification must isolate movements in the current price and must not simultaneously shift demand through states that affect the asset's future payoffs, namely dividends and resale values.

Going forward, to clarify this distinction, we refer to the class of elasticities that measure demand responses to shifts in state variables or deep parameters that persistently affect price levels as *shifter-process elasticities*. To define a shifter-process elasticity formally in the context of our model, consider the case where the asset price becomes marginally exposed to the shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$. Because $\ln P_{j,t} = \beta_j s_t + \ln \tilde{P}_{j,t}$, a one-unit change in the exposure β_j causes an s_t-unit change in the log price of asset j at time t.

Definition 1. *The shifter-process elasticity is defined as:*

$$\eta_j(s_t) \equiv -\frac{1}{s_t} \frac{d \ln \theta_j^*(s_t; \beta_j)}{d\beta_j}\bigg|_{\beta_i = 0 \ \forall i}. \tag{15}$$

The shifter-process elasticity measures how the date-t optimal portfolio share, $\theta_{j,t}^*$, responds to a marginal exposure to the shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$. Importantly, while the initial price shift is $d\beta_j s_t$, the response captured by this derivative reflects the effect of shifting the distribution of future price paths by $\{d\beta_j s_\tau\}_{\tau=t}^{\infty}$. Whereas the price elasticity reflects structural trading frictions (see Proposition 1 and its discussion), a shifter-process elasticity is specific to the dynamic properties of the shifter process associated with a given instrument. As we show below, a shifter-process elasticity can take any value, regardless of the presence of frictions.

While empirical variation introducing a change in the distribution of future prices can help identify the shifter-process elasticity defined in equation (15), the requirements for own-price elasticity identification are considerably more stringent. The following lemma details three conditions under which econometricians can infer the own-price elasticity from an incremental exposure of an asset to a shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$.

Lemma 2. *The own-price elasticity for asset j can be inferred from holdings data via variation in the asset's exposure to the shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$ if the following properties are satisfied:*

1. *The incremental price shift at time t fully and permanently resolves by the next trading date.*

2. *The price shift is unanticipated.*

3. *Only the price of asset j is shifted, while leaving all other prices unaffected, $\beta_{\neq j} = 0$.*

The lemma lays out three conditions for identification. First, variation used for estimation must leave prices at future trading dates (resale values) unaffected. Otherwise, it alters the fundamental characteristics of the asset by changing the future cash flows it can deliver to the investor. This requirement applies equally to time-series and cross-sectional identification strategies. The issue is that *any* variation inducing changes in resale values *after* the date when portfolio holdings are observed and used for estimation effectively alters an asset's intrinsic characteristics (its exposure to state dynamics), not just the current asset price, and thus directly violates the exclusion restriction.

Second, time-series variation inducing a price change at time t and used for estimation must be unanticipated by investors since otherwise, the associated holdings change does not just reflect a date-t price change but also previous expectations of that price change, resulting in biased elasticity estimates. For instance, if a positive price shift was already expected to occur, then the optimal pre-shock holding is elevated due the prospect of an incremental positive return. As a result, after the positive price shift has realized, the portfolio weight would decline more than it otherwise would, and the elasticity estimate based on time-series changes would be biased upwards (if this were the only deviation from Lemma 2).

Third, to isolate the own-price elasticity, the variation must exclusively affect the price of the asset under consideration. A simultaneous change in other asset prices would generally introduce cross-asset substitution effects that confound elasticity estimation. In Section 3.3, we characterize the residual supply-shock dynamics across all assets $j = 0, \ldots, J$ that would be required to induce any given price-shift process, including a short-lived shift that changes only the price of asset j. As highlighted there, such supply-shock processes must account for both cross-sectional substitution and dynamic wealth effects. Related to this issue, Fuchs et al. (2025b) establish in a static setting that, absent targeted simultaneous shocks to multiple assets, there are theoretical limits to identifying asset demand from purely observational data. Moreover, when Condition 3 of Lemma 2 is not satisfied, significantly more restrictive structural assumptions must be imposed to identify own-price elasticities, as emphasized by the existing literature examining these issues in static settings (see Fuchs et al., 2025a; Haddad et al., 2025; Koijen and Yogo, 2025). Rather than investigating the empirical plausibility of such structural assumptions, our paper focuses on identification issues arising in environments in which assets can be resold at the end of a period. Given this focus, our analysis is particularly concerned with the theoretical and empirical implications of violations of Condition 1 of Lemma 2.

First, we theoretically investigate how violations of Condition 1 of Lemma 2 affect elasticity estimates. The following proposition compares the true price elasticity under discrete trading,

ε_D, with the elasticity estimates obtained from a shock that introduces a price shift persisting for multiple trading periods, thereby violating Condition 1 of Lemma 2. The price-shift dynamics are illustrated in Figure 1.

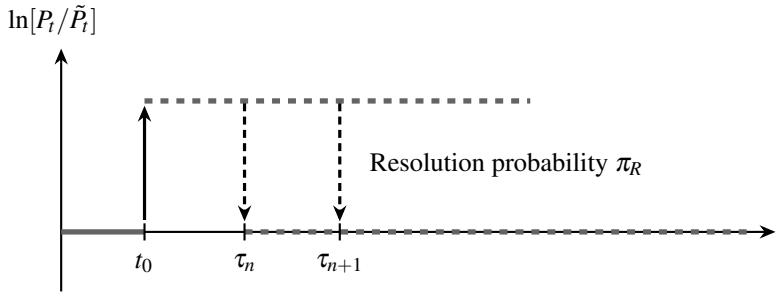


Figure 1: Stochastic price-shift resolution. The figure illustrates the price-shift dynamics considered in Proposition 2. At time t_0, the price level P_t is shifted up relative to the baseline level $\tilde{P}_t$, and at each Poisson arrival date $\tau_n > t_0$, this price shift is permanently resolved with probability π_R.

Proposition 2 (Constant per-period price shift resolution)**.** *Suppose conditions 2 and 3 of Lemma 2 are satisfied and that there is a single risky asset. Further, trade can occur only on event dates τ_n, and a marginal log price shift away from the price level $\tilde{P}$ resolves fully and permanently with per-event probability π_R. The elasticity estimated from such a price shift, $\hat{\varepsilon}$, recovers the shifter-process elasticity η and satisfies the following relations with the true price elasticity ε_D:*

1. ***One-period shifts:*** *For $\pi_R = 1$, condition 1 of Lemma 2 is satisfied and the estimated elasticity coincides with the true elasticity,*

$$\hat{\varepsilon} = \varepsilon_D. \tag{16}$$

2. ***Persistent shifts:*** *For $\pi_R < 1$, the estimated demand elasticity is biased toward zero. For short trading periods (large λ), the following relation obtains:*

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = \pi_R + O(\lambda^{-2}) \;\to\; \pi_R \quad as \;\; \lambda \to \infty. \tag{17}$$

Proof. See Appendix A.2. □

The proposition reveals that the per-period resolution probability π_R is directly linked to the bias in elasticity estimates. As an illustration, suppose an institutional investor can trade, on average, at a daily frequency and that a price shift used for identification persists for two years on average. In this case, the estimated elasticity would be only $1/730 \approx 0.14\%$ of the true elasticity.

One might infer from this that large biases require substantial persistence, but this inference would be incorrect. The $1/730$ figure is specific to the stylized assumption of a constant per-

period resolution probability, which embeds a nonzero probability of resolution in the investor's very first trading period. As the following proposition shows, a price shift that fully persists for only one trading period (say of expected length of 1 day) and resolves thereafter produces in fact a more extreme bias. Consequently, as we discuss below, persistence measured in the usual sense is misleading in this context of biases. The case where there is no price shift resolution in the first trading period is illustrated in Figure 2.

Proposition 3 (No resolution in the first trading period)**.** *Suppose conditions 2 and 3 of Lemma 2 are satisfied and that there is a single risky asset. Further, trade can occur only on event dates τ_n, and a marginal log price shift away from the price level $\tilde{P}$ remains in place for one trading period. Thereafter, it resolves fully and permanently with per-event probability π_R. The elasticity estimated from such a price shift, $\hat{\varepsilon}$, recovers the shifter-process elasticity η. For short trading periods (large λ), we obtain the following relation between the true elasticity ε_D and the estimated elasticity $\hat{\varepsilon}$:*

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = 0 + O(\lambda^{-2}) \;\to\; 0 \quad as \;\; \lambda \to \infty. \tag{18}$$

Under log utility and continuous trade (i.e., when trade is not restricted to event dates), we obtain exactly $\hat{\varepsilon} = \eta = 0$, whereas the true price elasticity is $\varepsilon_C = \infty$.

Proof. See Appendix A.3. □

Proposition 3 shows that the estimates are virtually unrelated to the true elasticity when price shifts do not resolve by the investor's next feasible trading date, a plausible scenario for a variety of candidate instruments. In this case, infinite frictions and zero frictions yield identical estimates. Therefore, these estimates have no external validity. This occurs because rational investors set portfolio weights primarily based on the return distribution over their next trading period, and that distribution is unaffected by the initial price shift. Under log utility with continuous trading, this result is exact: the estimated elasticity $\hat{\varepsilon} = \eta$ is zero even though the true price elasticity ε is infinite. While this sharp result holds when the asset does not pay a dividend in the first trading period, we show, based on a calibration presented in Case Study 2 of Section 3.3.2, that the elasticity is still very close to zero when that assumption is relaxed.

Proposition 3 also highlights that empirically examining price-shift persistence at low frequencies, such as at the annual frequency, is not the relevant object for assessing the biases we document. This is because the demand response depends primarily on whether the price shift resolves by the investor's next feasible trading date. A price shift that fully resolves after being in place for one week might be typically viewed as highly transitory. Yet for an institutional investor that can trade intraday, this shift does not resolve by the next feasible trading date at all. Notably, the feasible trading frequency varies across investor types: for a high-frequency trading firm, it might

be as short as a millisecond, whereas for a pension-fund investor subject to regulatory or institutional adjustment frictions the feasible trading frequency might be significantly lower. The relevant benchmark for evaluating instrument validity is thus the investor's trading frequency, a structural feature determined by technology, regulation, or behavioral constraints, not the econometrician's choice of data frequency.

The following corollary to Propositions 2 and 3 characterizes an approximate bias-correction factor that maps the estimated shifter-process elasticity to the true price elasticity, and establishes when such a correction is feasible.

Corollary 1 (Bias correction)**.** *When estimating an elasticity with an instrument that shifts both the contemporaneous price and future prices, the estimate recovers the shifter-process elasticity,* $\hat{\varepsilon} = \eta$*, which differs from the true price elasticity* ε*. For short trading periods (high* λ*), the bias-correction factor* $\mathscr{C}$ *satisfying*

$$\mathscr{C} \times \hat{\varepsilon} = \varepsilon \tag{19}$$

is approximately equal to

$$\mathscr{C} \approx \frac{\textit{Price Shift}}{\textit{Expected Resolution by Next Feasible Trading Date}}. \tag{20}$$

This correction factor requires that the price shift is expected to at least partially resolve by the investor's next feasible trading date, a horizon that is investor-specific; otherwise the denominator is zero and the correction is infeasible (see Proposition 3).

The bias correction factor $\mathscr{C}$ is a model-based approach designed to address a failure of the exclusion restriction: the instrument shifts future prices (which largely determine the benefit the investor receives from the asset) in addition to P_t (which governs the cost of the asset). We note that since both the numerator and the denominator of the correction factor $\mathscr{C}$ are estimated with noise, any empirical implementation must properly account for the joint distribution of estimation error to determine quantities such as standard errors for the debiased estimate. In particular, when price-shift resolution is small, the correction factor $\mathscr{C}$ is large, and the same level of absolute estimation error in the denominator translates into large relative error in $\mathscr{C}$, leading to a noisier debiased estimate.

Compounding the correction factor noise discussed above, low resolution also weakens the demand response signal, a point already established in the limiting case of no resolution, analyzed in Proposition 3: when the price shift does not resolve at all over the first trading period, the estimated elasticity is zero and entirely uninformative about frictions. More generally, low resolution generates a smaller expected return shift over the next trading period, attenuating the demand response.

Low resolution thus causes both problems simultaneously.

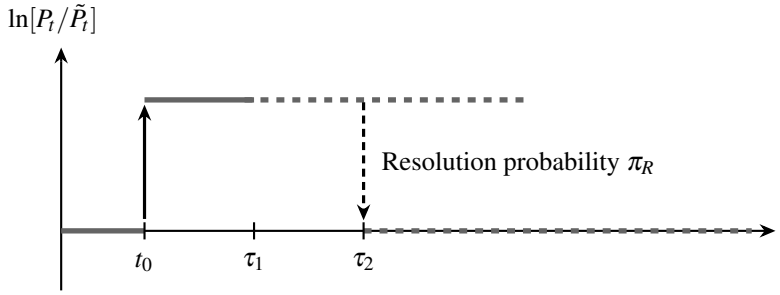


Figure 2: No resolution in the first trading period. The figure illustrates the price-shift dynamics considered in Proposition 3. At time t_0, the price level P_t is shifted upward relative to the baseline level $\tilde{P}_t$. This shift is not resolved in the first trading period, implying that it is still in place at the next Poisson arrival date τ_1. Thereafter, at each Poisson arrival date $\tau_n > \tau_1$, the shift is permanently resolved with probability π_R.

Finally, we turn to another important scenario, momentum-type price-shift dynamics, to highlight their particularly severe implications for elasticity mismeasurement. These dynamics are illustrated in Figure 3. As we will show in our empirical analysis in Section 4, such dynamics are potentially relevant in the context of KY19's instrument.

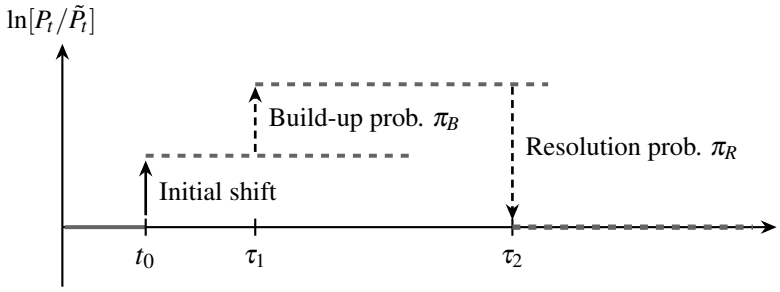


Figure 3: Price-shift build-up (momentum). The figure illustrates the price-shift dynamics considered in Proposition 4. Initially, at time $t = 0$, the price level P_t is shifted up relative to the baseline level $\tilde{P}_t$. Thereafter, at each subsequent Poisson arrival date $\tau_n > 0$, the shift builds up to twice its initial magnitude with probability π_B. Once that level is reached, the price level resolves back to the baseline level $\tilde{P}_t$ with resolution probability π_R at each subsequent Poisson arrival date.

Proposition 4 (Momentum dynamics)**.** *Suppose that conditions 2 and 3 of Lemma 2 are satisfied and that there is one risky asset. Further, trade can occur only on event dates τ. After an unanticipated log price shift $d\beta$ away from $\tilde{P}$, the price shift builds up to $2d\beta$ with probability π_B per event period before permanently and fully resolving back to the initial level $\tilde{P}$, with probability π_R. The elasticity estimated from such a price shift, $\hat{\varepsilon}$, recovers the shifter-process elasticity η. For short trading periods (large λ), we obtain the following relation between the estimated elasticity $\hat{\varepsilon}$ and*

the true elasticity ε_D:

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = -\pi_B + O(\lambda^{-2}) \;\to\; -\pi_B \quad as \;\; \lambda \to \infty. \tag{21}$$

Proof. See Appendix A.4. □

Proposition 4 reveals the severe biases that variation triggering momentum dynamics introduce into elasticity estimates. Momentum causes investors to anticipate further price increases, thereby reversing their portfolio response compared to a pure contemporaneous price increase. Consequently, estimated elasticities are opposite in sign to the true elasticity. The more likely the subsequent build-up the more extreme is the bias.

3.3 Equilibrium Pricing and Asset Purchase Programs

In this section, we first close the model by characterizing the supply-shock dynamics that sustain the postulated price dynamics and portfolio problem described in Sections 2 and 3.1. The construction is deliberately minimal and ensures that no new state variables beyond s_t are introduced into CRRA investors' maximization problem. We then characterize how price elasticities and shifter-process elasticities relate to the multipliers discussed in the literature and derive the dynamic asset-purchase programs required to induce a targeted stochastic price-support process in equilibrium.

In a general-equilibrium environment with heterogeneous agents, as in the asset-pricing literature reviewed in the introduction, the cross-sectional allocation of assets, consumption, or institutional capital is typically a state variable and equilibrium prices are functions of it. Our approach effectively inverts this mapping: the supply dynamics of an outside bloc are pinned down by, rather than pinning down, the postulated price process. Although this residual supply process is generally complex, by design it preserves the tractable, low-dimensional state space introduced in earlier sections, which is pedagogically helpful for examining any Markovian price-shift process, including momentum dynamics. Beyond this tractability advantage, our approach lends itself directly to our analysis of asset purchase program interventions: we can ask what dynamic properties a purchase program must have to sustain a given targeted Markovian price-support process.

3.3.1 Supply Behavior of the Outside Bloc

Let x_j denote a baseline number of shares outstanding of asset j. At any time t, denote by θ_j^* the optimal portfolio share invested in asset j by the unit mass of CRRA investors, and let W_t be their

aggregate wealth. The outside group chooses its unit position according to

$$z_{j,t} = x_j - \frac{\theta^*_{j,t} W_t}{P_{j,t}}, \qquad j = 0, \ldots, J, \tag{22}$$

which implies unit market clearing at every instant:

$$n_{j,t} + z_{j,t} = x_j, \quad \text{with} \quad n_{j,t} \equiv \frac{\theta^*_{j,t} W_t}{P_{j,t}} \qquad j = 0, \ldots, J. \tag{23}$$

Consequently, the law of motion of s_t and the transition matrix $q(s'|s)$ specified in Section 2 fully describe the evolution of equilibrium price shifts relative to the price level $\tilde{P}_j$. No separate tracking of the outside group's wealth or portfolio is necessary.

The idea that stochastic supply shocks generate price movements dates back at least to Grossman and Stiglitz (1980). In a dynamic context, Greenwood and Vayanos (2014) examine a term-structure model in which risk-averse arbitrageurs absorb shocks to the demand and supply for bonds of different maturities. Our model similarly captures such shocks but reverse engineers their magnitudes so as to match the equilibrium price dynamics postulated in Section 2. Several familiar economic mechanisms can rationalize the sources of such supply shocks. In particular, the outside group of agents affecting supply may consist of noise traders or central banks.

3.3.2 Elasticities, Multipliers, and Dynamic Asset Purchase Programs

In this section, we investigate how elasticities, which are partial-equilibrium objects pertaining to specific investors, relate to multipliers, which quantify the impact of purchases on equilibrium market prices. In its simplest form, a multiplier measures how the market capitalization of an asset changes when a particular type of investor engages in a purchase of that asset. In dynamic contexts, however, this concept requires further refinement.

As we explain below, once we move to dynamic purchase programs, the idea of a multiplier measuring the initial price impact of a purchase becomes problematic. The reason is that the entire stochastic path of future purchase commitments, rather than the initial transaction, drives the price adjustment at the time of the program's announcement. We then provide a key result of our paper: a characterization of the dynamic purchase programs required to induce any given marginal Markovian price-shift process for a particular asset.

Price elasticities and one-period purchases. Suppose the outside block purchases an additional marginal unit of asset j. Starting from the market-clearing condition in changes:

$$dz_{j,t} = -\frac{dn_{j,t}}{dP_{j,t}} dP_{j,t}, \tag{24}$$

we obtain the following relation by multiplying both sides by x_j and $P_{j,t}$ and rearranging terms:

$$\underbrace{\frac{dP_{j,t}}{P_{j,t}}}_{\text{Market cap change}} = \underbrace{\frac{dz_{j,t}}{x_j}}_{\text{Investment}} \cdot \underbrace{\left[-\frac{d\ln n_{j,t}}{d\ln P_{j,t}} \cdot \frac{n_{j,t}}{x_j}\right]^{-1}}_{\text{Multiplier}}. \tag{25}$$

That is, the relative change in the equilibrium price is given by the product of the fraction of shares outstanding purchased and the multiplier. The multiplier is the inverse of the product of CRRA investors' price elasticity of unit demand, $-\frac{d\ln n_{j,t}}{d\ln P_{j,t}}$, and their existing ownership share, $\frac{n_{j,t}}{x_j}$. The latter term measures the size of market participants that can take the other side of the outside block's trades.

Note that a change in the price $P_{j,t}$ also induces adjustments in CRRA investors' holdings of other assets $i \neq j$. Therefore, the outside block must compensate for these spillover adjustments in order to keep the prices of other assets unaffected. We further elaborate on this issue below, in the context of Proposition 5.

It is useful to characterize the relationship between the price elasticity of unit demand and the price elasticity of the portfolio weight ε. Using the identity $n_{j,t} = \frac{\theta^*_{j,t} W_t}{P_{j,t}}$, we obtain:

$$\begin{aligned} -\frac{d\ln n_{j,t}}{d\ln P_{j,t}} &= -\left(\frac{d\theta^*_{j,t}}{dP_{j,t}}\frac{W_t}{P_{j,t}} + \frac{dW_t}{dP_{j,t}}\frac{\theta^*_{j,t}}{P_{j,t}} - \frac{W_t}{P_{j,t}}\frac{\theta^*_{j,t}}{P_{j,t}}\right)\frac{P_{j,t}}{n_{j,t}} \\ &= \varepsilon_{j,t} + 1 - \theta^*_{j,t}. \end{aligned} \tag{26}$$

The difference between the two elasticity concepts arises from the fact that the impact of a price increase on unit demand can be decomposed into three channels: its effect on the portfolio weight, its effect on the investor's wealth level, and the direct mechanical effect of making the asset more expensive.

Dynamic purchase programs and pseudo multipliers. We now characterize the dynamic purchase programs that are required to induce targeted changes in the price dynamics of a given asset j.

Proposition 5. *The unanticipated initiation and public announcement of the following processes*

of incremental holdings by the outside block generate a stochastic log-price path differential $\{d\beta_j s_\tau\}_{\tau=t}^{\infty}$ *for asset* j*:*

$$\left\{\frac{dz_{j,\tau}}{x_j}\right\}_{\tau=t}^{\infty} = \left\{d\beta_j s_\tau \cdot [\eta_{j,\tau} + (1-\theta^*_{j,\tau})] \cdot \frac{n_{j,\tau}}{x_j}\right\}_{\tau=t}^{\infty}, \tag{27}$$

$$\left\{\frac{dz_{i,\tau}}{x_i}\right\}_{\tau=t}^{\infty} = \left\{d\beta_j s_\tau \cdot [\eta_{ij,\tau} - \theta^*_{j,\tau}] \cdot \frac{n_{i,\tau}}{x_i}\right\}_{\tau=t}^{\infty} \quad \forall i \neq j, \tag{28}$$

where we define the cross-shifter-process elasticity

$$\eta_{ij,\tau} \equiv -\frac{1}{s_t}\frac{d\ln\theta^*_i(s_t)}{d\beta_j}\bigg|_{\beta_i=0,\forall i}. \tag{29}$$

Proof. See Appendix A.5. □

While asset j's market price appreciates by $s_t d\beta_j$ upon announcement of the purchase program described in Proposition 5, this effect is sustained by the full future processes of incremental unit holdings, as detailed in equations (27) and (28). Among other factors, these depend on the dynamics of shifter-process elasticities $\{\eta_{j,\tau}\}_{\tau=t}^{\infty}$, the ownership shares of other market participants $\{n_{j,\tau}/x_\tau\}_{\tau=t}^{\infty}$, and their portfolio shares $\{\theta^*_{j,\tau}\}_{\tau=t}^{\infty}$.

Notably, while the shifter-process elasticity $\eta_{j,\tau}$ at time τ is specific to the targeted price-shift process, it can be traced back to a deeper structural object: the true price elasticity ε_j. In particular, as shown previously in Propositions 2 and 3 and Corollary 1, if $\pi_{R,\tau}$ reflects the expected fraction of the price shift that resolves over the next trading period starting at time τ, then $\eta_{j,\tau} \approx \pi_{R,\tau} \times \varepsilon_j$. Thus, once we know the true price elasticity ε_j, we can construct the approximate path of shifter-process elasticities $\{\eta_{j,\tau}\}_{\tau=t}^{\infty}$ implied by the targeted price-shift process.

Proposition 5 also shows that the *pseudo multiplier*, defined as the apparent impact of the date-t purchase quantity on the date-t price,

$$\frac{s_t d\beta_j}{dz_{j,t}/x_j} = \left[(\eta_{j,t} + (1-\theta^*_{j,t})) \cdot n_{j,t}/x_t\right]^{-1}, \tag{30}$$

does not actually measure the source of the initial price reaction. The reason is that it is the *future* changes in holdings that ultimately induce the initial price adjustment. To illustrate this point, we examine in Case Study 1 below an example of purchase program announcements, such as those sometimes initiated by central banks.

It is important to note that because the program described in Proposition 5 is intended to affect only the price dynamics of asset j, cross-shifter-process elasticities $\eta_{ij}(s_\tau) \neq 0$ imply that the outside block must generally also adjust its holdings in other assets $i \neq j$ to compensate for substitution effects by CRRA investors, as shown in equation (28). We will examine this issue in greater

detail in Case Study 2 below.

Case Study 1: Purchase Program Announcements

Central bank announcements of asset purchase programs provide a useful and relevant example illustrating the limited informational content of multipliers. Suppose the Federal Reserve announces that it will purchase an asset in future states of the world at prices higher than would prevail absent intervention. Anticipating this higher future resale value, the equilibrium price increases immediately at time t, even though the Federal Reserve does not purchase any assets at that time.

As a result, the multiplier measured at inception of the program (30) becomes infinite for any such announcement that promises *future* asset purchases. Yet, asset purchases are clearly not infinitely impactful; the multiplier is misleading and does not quantify the causal impact of date-t purchases on date-t equilibrium prices.

For CRRA investors to be content with a higher portfolio share in the asset (which results after the price increase), the expected return must rise. Consequently, momentum dynamics arise endogenously from such policy announcements, generally causing the shifter-process elasticity $\eta_{j,t}$ to become negative. This is not a special knife-edge case of merely theoretical relevance but rather a generic outcome for program announcements that do not immediately involve purchases upon inception, or feature dynamic position build-up. These issues underscore that estimating multipliers associated with dynamic residual supply shock processes is generally problematic and yields limited insights.

Case Study 2: Permanent Price Shifts

We now consider the special theoretical case of a constant and permanent log price shift affecting the entire stock market, an example also discussed in Gabaix and Koijen (2023). Our objective is to characterize the dynamics of trades required to induce such a shift and to provide a simple calibration to gauge magnitudes. In this scenario, the shifter state is constant; that is, $s_\tau = 1$ for all $\tau \geq t$.

Suppose that the investors taking the other side of the purchases are CRRA investors with a relative risk aversion of 2 and initial holdings equal to 95% of shares outstanding, i.e., $n_{1,t}/x_{1,t} = 0.95$. Assume that, prior to the announcement of the dynamic purchase program, the market risk premium is 6%, volatility is 19%, and the dividend yield is 1.7%. In this setting, CRRA investors' risky portfolio share before the intervention is then given by:

$$\theta_1^* = \frac{\mu_1 - r_f}{\gamma \sigma_1^2} = \frac{0.06}{2 \cdot 0.19^2} = 0.83. \tag{31}$$

Immediately after the onset of the asset purchase program, the expected return on the stock market declines due to a reduction in the dividend yield component of the expected return:

$$\frac{d\mu_1}{d\beta_1} = \frac{d\frac{D_1}{\tilde{P}_1 e^{\beta_1 \cdot 1}}}{d\beta_1} = -\frac{D_1}{P_1}. \tag{32}$$

Correspondingly, we obtain the following sensitivity of the portfolio share θ_1^* with respect to the exposure parameter β_1:

$$\frac{d\theta_1^*}{d\beta_1} = \frac{-D_1/P_1}{\gamma \cdot \sigma_1^2} = -0.2354. \tag{33}$$

That is, for a 1% price shift, $\Delta\beta_1 = 0.01$, the risky portfolio share decreases from 0.83 to 0.828. We then obtain the following shifter-process elasticity:

$$\eta_1 = -\frac{1}{s}\frac{d\ln\theta_1^*}{d\beta_1}\bigg|_{\beta=0} = \frac{D_1/P_1}{\mu_1 - r_f} = \frac{0.017}{0.06} = 0.28, \tag{34}$$

that is, a number close to zero. This result contrasts with Gabaix and Koijen's (2023) argument that this object is 20 in frictionless models and that, since this number is orders of magnitude larger than empirical estimates, significant frictions must be present in practice.[9] Converting (34) to the corresponding shifter-process elasticity in terms of unit demand gives $\eta_1 + (1 - \theta_1^*) = 0.45$. Substituting this value into equation (30) yields a pseudo multiplier equal to:

$$\frac{d\beta_1}{dz_{1,t}/x_1} = \left[(\eta_{1,t} + (1 - \theta_{1,t}^*)) \cdot n_{1,t}/x_t\right]^{-1} = 2.34. \tag{35}$$

However, as emphasized earlier, it would be incorrect to suggest that a one-time flow or purchase achieves these price effects. Rather, it is the entire process of future holdings that induces the price increase at the inception of the program. In particular, equation (27) in Proposition 5

[9] Gabaix and Koijen (2023) obtain their substantially higher frictionless benchmark value of 20 by considering a setting in which the elasticity with respect to a permanent price shift is indeterminate. In their model, identical claims to aggregate consumption are nominally separated into two categories, "equity" and "human wealth," both of which can be traded frictionlessly. Perfect substitutability between these shares implies that any intended quantity adjustment in response to a uniform price shift can be implemented through any combination of adjustments in these two classes of shares. Gabaix and Koijen then select the case in which the total quantity adjustments for all assets in the economy are made exclusively in those shares labeled equity. Relative to a pro-rata adjustment, this assumption substantially increases the elasticity value, having the same quantitative effect as setting the dividend yield to 80% in equation (34). We note that this approach of leveraging the indeterminacy can justify any value between $-\infty$ and $+\infty$, as shorting is also feasible absent frictions.

simplifies to the following dynamics for the incremental holdings of the intervening party:

$$\left\{ \frac{dz_{1,\tau}}{x_1} \right\}_{\tau=t}^{\infty} = \{ d\beta_1 \cdot 0.45 \cdot n_{1,\tau}/x_1 \}_{\tau=t}^{\infty}, \tag{36}$$

To raise the total market capitalization by $\Delta\beta_1 = 1\%$, the outside block must purchase $1\%/2.34 = 0.43\%$ of shares outstanding upon impact. Thereafter, it must adjust this incremental exposure dynamically, depending on how other investors' footprint, $n_{1,\tau}/x_1$, evolves over time as a consequence of their risky portfolio share θ_1^* and the stochastic realization of future stock market returns.

Moreover, to ensure that the purchase program does not also affect the price path of the risk-free asset, the outside block must adjust its holdings as follows (see details in Appendix A.5):

$$\{ dz_{0,\tau} \}_{\tau=t}^{\infty} = \left\{ -dz_{1,\tau} \frac{P_{1,\tau}}{P_{0,\tau}} \right\}_{\tau=t}^{\infty}. \tag{37}$$

Initially, at time t, the bond adjustment by the outside block merely funds its incremental stock market holding. However, for $\tau > t$, the outside block must adjust this position. As a result, we obtain a stochastic process for bond holdings adjustments, $\{ dz_{0,\tau} \}_{\tau=t}^{\infty}$, which depends on how the stock market performs relative to the risk-free asset, $\{\frac{P_{1,\tau}}{P_{0,\tau}}\}_{\tau=t}^{\infty}$, as well as on the process of stock market holdings adjustments, $\{ dz_{1,\tau} \}_{\tau=t}^{\infty}$.

In sum, even in this stylized case of a permanent price shift, the required asset-purchase program is far more complex than a single multiplier would suggest, as it requires managing an entire stochastic path of holdings in both the risky and risk-free assets.

Case Study 3: Constant Per-Period Resolution Probabilities

As a final case study, we examine price shifter dynamics characterized by a constant per-period resolution probability. As shown in Section 3, the more persistent the price shifts are, the less CRRA investors optimally adjust their portfolio shares. Hence, by equation (22), persistent price movements can be supported by smaller variations in the outsiders' net supply, $z_{j,t}$, upon impact. In contrast, the transitory price dynamics required for identification under Lemma 2 induce larger swings in θ_j^*. To sustain such dynamics, the outside block must inject (or withdraw) correspondingly larger quantities of the asset. This leads to the following corollary.

Corollary 2 (Persistence and magnitudes of supply shocks)**.** *Suppose the conditions laid out in Proposition 2 are satisfied, except that we consider continuous trading. For a given marginal price shift that resolves with the per-event resolution probability π_R, the initial unit risky-asset demand*

change by outsiders satisfies

$$\left|dz_1(\pi_R')\right| < \left|dz_1(\pi_R'')\right| \quad \textit{for all} \quad 0 \le \pi_R' < \pi_R'' \le 1.$$

Proof. See Appendix A.6. □

While Corollary 2 pertains to the single-asset case, we illustrate the same mechanics for a multi-asset setting in Figure 4, with parameter values reported in Table 1. The shifter state s follows the resolution Markov chain described in Proposition 2. The figure plots the magnitude of supply shocks required upon impact to induce a given log price shift.

The results show that shock persistence, measured by the expected time to resolution, has a dramatic effect on the size of the required initial supply shock. The baseline "one-period" shock corresponds to a weekly event interval and assumes full resolution within that period. If, instead, the price shock persists for a quarter on average, the required supply shock upon impact falls to approximately 7%–8% of the one-period baseline. With an average duration of one year, the magnitude further declines to just 2%, and for a two-year resolution window, to less than 1%.

Table 1: Calibration and Relevant Statistics

Panel A. Investor preferences and timing

Parameter	Symbol	Value	Unit / Note
Risk-free rate	r_f	0.04	annual
Subjective discount rate	ρ	0.02	annual
Relative risk aversion	γ	2	–
Event intensity	λ	52	weekly rebalancing opportunities
Consumption probability	π_C	0.002	per event date
Number of risky assets	J	3	–

Panel B. Asset characteristics (per risky asset j)

Parameter	Symbol	Value	Unit / Note
Price drift (all j)	μ_j	0.09	$= r_f + 0.05$
Common-factor volatility	$\sigma_{A,j}$	$(0.238, 0.190, 0.190)$	per asset
Idiosyncratic volatility	$\sigma_{I,j}$	$(0.300, 0.300, 0.300)$	per asset
Log price shift	$\Delta\beta_1$	0.0011	–

Simply put, more persistent shocks require much smaller supply shocks upon impact to induce a given price shift, as CRRA investors' optimal demand response is substantially muted when price changes are expected to persist. Importantly, these patterns do not reflect frictions that impede portfolio adjustments; rather, they arise from optimal dynamic trading behavior in response

to changes in asset characteristics, specifically the dynamics of expected resale values. This implies that persistent supply shocks generated by liquidity traders or irrational traders can generate substantial price-level shifts even when the rest of the market is populated by fully rational traders who can trade frictionlessly: if such shocks are expected to shift prices persistently, they have little effect on the return distribution and therefore give rational investors little incentive to trade against them. This observation is economically important because price levels are the object through which financial market prices affect real allocations (see Binsbergen and Opp, 2019).

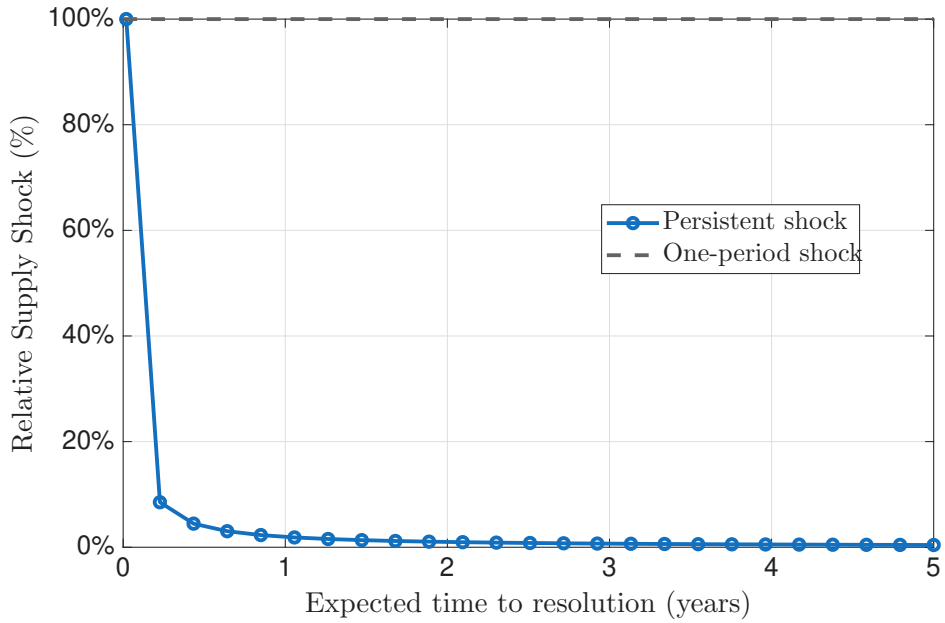


Figure 4: Supply Shock Magnitudes Upon Impact. The figure illustrates the magnitude of supply shocks required to induce a given log-price shift upon impact, as a function of the shock's persistence, measured by the expected time to resolution, $\frac{1}{\lambda \pi_R}$. Throughout, the inception of supply shocks is assumed to be unanticipated. We plot the relative supply shock, $\frac{|\Delta z_j(\pi_R)|}{|\Delta z_j(1)|}$, where varying π_R adjusts the expected time to resolution. The shifter state s follows the resolution Markov chain described in Proposition 2. See Appendix A.7 for further details.

4 Direct Statistical Tests of Instrument Validity

In this section, we build on the insights from our theoretical analysis in Sections 2 and 3 and conduct a series of empirical tests to examine the characteristics and validity of the statistical instruments used in the literature—particularly those proposed by KY19. KY19 recover the slope of an investor's demand schedule from a purely cross-sectional estimation. They define the instrument

for the log market equity of asset n at time t from the perspective of institution i as

$$\widehat{me}_{i,t}(n) = \log\left[\sum_{j\neq i} \text{AUM}_{j,t} \frac{\mathbf{1}_j(n)}{1+\sum_{m=1}^{N_{j,t}} \mathbf{1}_j(m)}\right], \tag{38}$$

that is, the log market value of firm n at the market-clearing price under the counterfactual where all institutions j other than i allocate their assets under management, $\text{AUM}_{j,t}$, as an equal-weighted portfolio within their investment universe. An institution's investment universe is the set of stocks it currently holds or has held in the past 11 quarters (see Appendix B for details).

As we explore the implications of time-series variation and predictability, we index the instrument by calendar time t and write $\widehat{me}_{i,t}(n)$. Throughout, we denote by $me_t(n)$ the actual log market value of asset n at time t, and by $\theta_{i,t}(n)$ the portfolio weight of institution i in asset n. As in KY19, the "outside asset" ($n = 0$) is a residual category of holdings, as defined in Appendix B.

In KY19, the demand equation takes the form of a cross-sectional logit specification:

$$\frac{\theta_{i,t}(n)}{\theta_{i,t}(0)} = \exp\{\beta_{0,i,t} me_t(n) + \sum_{k=1}^{K-1} \beta_{k,i,t} x_{k,t}(n) + \beta_{K,i,t}\} \cdot \nu_{i,t}(n), \tag{39}$$

where $x_{k,t}(n)$ denotes observable fundamental k for asset n at time t, and $\nu_{i,t}(n)$ is an institution-asset-specific demand shifter (an unobserved taste shock). We stack the characteristics in the vector $\mathbf{x}_t(n)$.

Under KY19's identifying assumptions, $\widehat{me}_{i,t}(n)$ serves as a valid instrument for $me_t(n)$: it shifts the price but is assumed to be orthogonal to $\nu_{i,t}(n)$, that is,

$$\mathbb{E}\left[\nu_{i,t}(n) | \widehat{me}_{i,t}(n), \mathbf{x}_t(n)\right] = 1. \tag{40}$$

4.1 Descriptive Statistics

In constructing the dataset for our empirical analysis, we follow KY19 both in terms of data sources and data selection criteria. After excluding observations where size falls below the NYSE 10th percentile breakpoint, we compute time-series averages of cross-sectional descriptive statistics (including percentiles) for various variables of interest. Tables 6 and 7 in Appendix C report these statistics for the stationary and non-stationary variables in our dataset, respectively (the latter computed over successive five-year sub-periods). As in KY19, our sample period is 1980Q1-2017Q4, and all characteristics represent signals measured at quarter-end. Accounting variables are lagged by six months relative to market variables.

The reported variables are computed as follows. Profitability (*profit*) is defined as operating profits divided by book equity, where operating profits equal revenues (REVTS) minus cost of

goods sold (COGS), SG&A expenses (XSGA), and interest expenses (XINT). Investment (*Gat*) is the annual log growth rate of book assets. The numerator in the dividend-to-book ratio (*divA_be*) represents the trailing-12-month dividends per current split-adjusted share multiplied by shares outstanding.

Market betas (*beta*) are estimated using 60-month rolling windows, retaining quarter-end values. The book-to-market ratio (BtM) is computed as six-month-lagged book equity divided by end-of-quarter market capitalization. Book equity is calculated as stockholders' equity (SEQ), plus deferred taxes and investment tax credit (TXDITC), minus the redemption value of preferred stock (PSTKRV), and set to missing if the result is negative.[10] The momentum return (MOM) at the end of quarter t is the compounded return over the past three quarters ($t-1$, $t-2$, $t-3$) plus the first two months of quarter t.

The component of the log market capitalization attributed to variation in $\widehat{me}$ is constructed as follows. For each quarter t, we run a cross-sectional regression of the form

$$me_t(n) = a_t + b_t \times \widehat{me}_t(n) + \mathbf{c}_t' \mathbf{x}_t(n) + e_t(n), \tag{41}$$

where $\mathbf{x}_t(n)$ contains the four characteristics described above: *profit*, *Gat*, *divA_be*, and *beta*. We then compute $\hat{b}_t \times \widehat{me}_t(n)$ to extract the component of $me_t(n)$ attributable to variation in the instrument, controlling for the observables $\mathbf{x}_t(n)$.

For all three non-stationary variables, the average cross-sectional mean trends upward over time, as expected. Interestingly, while the cross-sectional dispersion of log size (*me*) remains roughly constant throughout the sample period, the cross-sectional spread (in log terms) of $\widehat{me}$ declines over time. This decline results in a narrowing cross-sectional spread of $\hat{b} \times \widehat{me}$.

4.2 Instrumented Price Shifts and Future Returns

As highlighted in our theoretical analysis in Sections 2 and 3, investors' portfolio weights respond to price level shifts only to the extent that these shifts affect the return distribution going forward. In particular, for price elasticity estimation, the most relevant object is how the instrument affects prices, and thus expected returns, up to and including the next trading date of the institutional investors examined.

Correspondingly, we now examine whether cross-sectional variation in instrumented price shifts $\hat{b}_t \times \widehat{me}_{i,t}(n)$ predicts cross-sectional variation in future returns, controlling for the KY19 asset pricing factors. If these price shifts do not translate into resolution returns by the next trading

[10] If SEQ is missing, we use the sum of common equity and preferred stock (CEQ + PSTK). If this is also missing, we use the difference between book assets (AT) and total liabilities (LT). If PSTKRV is missing, we use the liquidating value (PSTKL). If this is also missing, we use the total value of preferred stock (PSTK).

date, a mean-variance investor's optimal portfolio weights will remain unresponsive to this variation. Consider, for example, two stocks with identical observable characteristics $\mathbf{x}_t(n)$, except that one is priced 1% higher, both today and at the next trading date. Absent dividends before the next trading date, both stocks offer identical return opportunities. As a result, a mean-variance investor will have no incentive to tilt her portfolio weights based on the price difference.

Moreover, dividends have quantitatively little impact on this baseline result. Consider, for example, a monthly rebalancing frequency. Suppose the "cheaper" stock has a monthly dividend yield of $\frac{D}{\tilde{P}} = 20$ basis points. Then, the dividend yield for the more expensive stock is $\frac{D}{\tilde{P}\times 1.01} = 19.8$ basis points. If the dividend yield accounts for 20% of the expected return, the two stocks would have expected returns of 100 basis points and 99.8 basis points, respectively, a difference of just 0.2 basis points. In other words, there is hardly any incentive to tilt portfolio weights.

This scenario is entirely distinct from the case where the full cross-sectional price difference of 1% resolves by the next rebalancing date. In that case, the expensive asset's expected return would be 1% lower than that of the other asset, giving the investor a strong incentive to tilt away from the higher-priced stock. However, only this latter scenario identifies investors' true demand elasticity. By contrast, in the case of persistent price level shifts, we are comparing two assets that, for all intents and purposes, appear identical to the investor.

This example motivates our empirical tests of identification condition 1 of Lemma 2, specifically, whether instrumented shifts in prices are associated with corresponding (full) resolution returns of the same magnitude over the next trading period. To begin, we examine the extent to which cross-sectional variation in instrumented shifts predicts future returns. Specifically, we estimate panel regressions of cumulative individual stock returns over various horizons H. Each excess-return regression takes the following generic form:[11]

$$R_{(t,H)}(n) = A_H + B_H \times [\hat{b}_t \times \widehat{me}_t(n)] + C_H \times \tilde{\mathbf{X}}_{n,t} + E_{n,t,H}$$

where $\tilde{\mathbf{X}}$ contains cross-sectionally standardized stock characteristics. In Table 2, each regressor is standardized as a within-period z-score. In Table 3, we alternatively map cross-sectional ranks onto the interval $[-1,1]$ (see, e.g., Gu et al., 2020). Moreover, in the full-resolution tests below, we consider a specification that does not apply any standardization. We separately examine the extent to which cross-sectional variation in instrumented price shifts at time t predicts returns over the first day, one week, and one month. Moreover, longer horizons (one quarter, two quarters, and four quarters) are examined in Appendix C. To avoid overlapping data at the two- and four-quarter horizons, we retain only every H-th observation within each PERMNO-specific time series. Standard errors are two-way clustered. Further details on the data are provided in Appendix B.

[11] For these return predictability tests, we use KY19's instrumented market equity applied to the household sector.

Table 2: Return predictability regressions with z-score standardized regressors. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are standardized in each period (z-scored). Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

(a) One Day

Regressor	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	0.0006	—	0.0008
	(0.0005)		(0.0005)
PROF	0.0004	0.0005	0.0004
	(0.0004)	(0.0004)	(0.0004)
INV	-0.0008***	-0.0008***	-0.0008***
	(0.0002)	(0.0002)	(0.0002)
DIV	0.0004**	0.0005***	0.0004**
	(0.0002)	(0.0002)	(0.0002)
BETA	-0.0003	-0.0003	-0.0004
	(0.0004)	(0.0004)	(0.0004)
$me - \hat{b} \times \widehat{me}$	—	-0.0003	-0.0004
		(0.0003)	(0.0004)

(b) One Week

(a)	(b)	(c)
0.0005	—	0.0006
(0.0010)		(0.0010)
0.0003	0.0003	0.0003
(0.0008)	(0.0008)	(0.0008)
-0.0024***	-0.0025***	-0.0024***
(0.0006)	(0.0006)	(0.0006)
0.0007	0.0007*	0.0007
(0.0005)	(0.0005)	(0.0005)
0.0002	0.0002	0.0002
(0.0011)	(0.0011)	(0.0011)
—	-0.0000	-0.0001
	(0.0007)	(0.0007)

(c) One Month

(a)	(b)	(c)
0.0011	—	0.0018
(0.0021)		(0.0022)
0.0041**	0.0044**	0.0043**
(0.0020)	(0.0020)	(0.0020)
-0.0033***	-0.0032***	-0.0031***
(0.0011)	(0.0011)	(0.0011)
-0.0003	0.0001	-0.0001
(0.0009)	(0.0008)	(0.0008)
0.0019	0.0018	0.0017
(0.0022)	(0.0022)	(0.0022)
—	-0.0016	-0.0020
	(0.0017)	(0.0018)

Table 3: Return predictability regressions with rank-standardized regressors. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are rank-standardized in each period. Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

	(a) One Day			**(b)** One Week			**(c)** One Month		
Regressor	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	0.0010	—	0.0015*	0.0008	—	0.0013	0.0022	—	0.0041
	(0.0007)		(0.0008)	(0.0014)		(0.0015)	(0.0030)		(0.0033)
PROF	0.0004	0.0006	0.0005	0.0005	0.0007	0.0007	0.0047**	0.0054**	0.0053**
	(0.0005)	(0.0005)	(0.0005)	(0.0010)	(0.0010)	(0.0010)	(0.0021)	(0.0023)	(0.0022)
INV	-0.0013***	-0.0012***	-0.0011***	-0.0039***	-0.0038***	-0.0037***	-0.0045***	-0.0041***	-0.0039**
	(0.0003)	(0.0003)	(0.0003)	(0.0009)	(0.0008)	(0.0008)	(0.0015)	(0.0016)	(0.0016)
DIV	0.0012**	0.0016***	0.0014***	0.0017	0.0021*	0.0019	-0.0001	0.0011	0.0006
	(0.0005)	(0.0006)	(0.0005)	(0.0013)	(0.0012)	(0.0013)	(0.0024)	(0.0024)	(0.0024)
BETA	-0.0003	-0.0003	-0.0005	0.0004	0.0004	0.0002	0.0027	0.0027	0.0021
	(0.0006)	(0.0006)	(0.0006)	(0.0016)	(0.0016)	(0.0016)	(0.0033)	(0.0032)	(0.0032)
$me - \hat{b} \times \widehat{me}$	—	-0.0011*	-0.0015**	—	-0.0012	-0.0015	—	-0.0045	-0.0056
		(0.0007)	(0.0007)		(0.0012)	(0.0013)		(0.0033)	(0.0036)

The point estimate of the coefficient on the instrumented price shift is positive across specifications and return windows, although it is statistically insignificant throughout. Among the windows considered, the one-day window is the most relevant for institutional investors that can trade intra-day; the one-month window provides a conservative upper bound on the relevant trading period.[12] As noted in our theoretical analysis, a positive association between price shifts and future returns could justify tilts toward stocks with higher instrumented market equity, consistent with negative estimated elasticities when using this instrument. Indeed, in Section 4.3, we find that estimated elasticities are negative for a substantial fraction of stocks and institutions.

These findings are robust to trading frequency and hold in the micro-cap subsample, which is excluded from the baseline sample. Tables 8 and 9 in Appendix C report analogous results for lower trading frequencies (one quarter to one year), and Tables 11 and 12 report results for micro-cap stocks. The coefficient on $\hat{b} \times \widehat{me}$ remains statistically insignificant in column (a) throughout; some specifications that additionally control for the component of size that is not the price shift ($me - \hat{b} \times \widehat{me}$; see column (c)) yield significant coefficients, but still remain close to zero economically.

Full-resolution tests. The above tests reveal the relationship between instrumented price shifts and future returns, and find that the two are unrelated or even positively related across all specifications and return windows. We now turn to a more direct test, which evaluates the magnitude of this relationship against the quantitative benchmark implied by identification Condition 1 of Lemma 2, namely, that a price shift used for elasticity identification must resolve fully by the next trading date. To examine this restriction, we no longer normalize variables, and instead directly use $\hat{b} \times \widehat{me}$ as the regressor, which, under the assumption that the instrument shifts prices for non-fundamental reasons (controlling for the observables $\mathbf{x}$), captures the corresponding log price shift. Thus, if institutions can trade daily, full resolution within that period would correspond to a coefficient of -1 on $(\hat{b} \times \widehat{me})$ when the KY19 demand-equation variables are included as controls (column (a)).

The results are reported in Table 4. At the one-day window most relevant for institutional investors trading intra-day, the point estimate for the relationship between the instrumented price shift and subsequent returns is positive and close to zero, far from the full-resolution benchmark of -1 (column (a)). We can reject the null hypothesis that the coefficient on $(\hat{b} \times \widehat{me})$ equals -1, with a p-value of 0.000 across all specifications. This provides direct evidence that the instrument fails to satisfy the full-resolution requirement for elasticity identification stated in Lemma 2. Table 10 in Appendix C reports analogous results for lower trading frequencies (one quarter to one year), and Table 13 reports results for micro-cap stocks; in all cases, the coefficient on $(\hat{b} \times \widehat{me})$ remains

[12] Davis et al. (2025) show in their Appendix Table B.1 that, in univariate Fama-MacBeth regressions without controls for the characteristics in the KY19 demand function, KY19's instrumented market equity positively predicts returns over the next month and over horizons up to one year.

Table 4: Price shift resolution regressions. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, beta, and size) as well as $\hat{b} \times \widehat{me}$. Variables are not standardized. Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). Across all specifications in the table, we can reject the null hypothesis that the coefficient on $(\hat{b} \times \widehat{me})$ is -1, with a p-value of 0.000. The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

	(a) One Day			**(b)** One Week			**(c)** One Month		
Regressor	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
$\hat{b} \times \widehat{me}$	0.0002	—	0.0002	-0.0006	—	-0.0000	0.0001	—	-0.0013
	(0.0002)		(0.0002)	(0.0006)		(0.0005)	(0.0012)		(0.0012)
PROF	0.0016	0.0017	0.0016	0.0022	0.0017	0.0017	0.0135**	0.0137**	0.0148**
	(0.0014)	(0.0014)	(0.0014)	(0.0026)	(0.0025)	(0.0025)	(0.0064)	(0.0064)	(0.0063)
INV	-0.0069***	-0.0069***	-0.0069***	-0.0163***	-0.0165***	-0.0165***	-0.0160*	-0.0155*	-0.0155*
	(0.0020)	(0.0020)	(0.0020)	(0.0044)	(0.0044)	(0.0044)	(0.0087)	(0.0087)	(0.0087)
DIV	0.0240***	0.0258***	0.0238***	0.0188	0.0093	0.0095	0.0076	0.0138	0.0303
	(0.0084)	(0.0087)	(0.0092)	(0.0206)	(0.0219)	(0.0220)	(0.0374)	(0.0410)	(0.0436)
BETA	0.0003	0.0003	0.0003	0.0013	0.0013	0.0013	0.0039	0.0038	0.0038
	(0.0005)	(0.0005)	(0.0005)	(0.0017)	(0.0017)	(0.0017)	(0.0031)	(0.0031)	(0.0031)
$me - \hat{b} \times \widehat{me}$	—	-0.0002	0.0000	—	0.0008	0.0008	—	-0.0004	-0.0019
		(0.0003)	(0.0003)		(0.0007)	(0.0007)		(0.0015)	(0.0019)

close to zero and we reject the null of -1 at $p = 0.000$.

4.3 Unrestricted Elasticity Estimates

The results in the previous section show that KY19's instrumented price shifts and future returns are unrelated or even positively related in the data. When positively related, our Proposition 4 indicates that we should expect negative elasticity estimates, opposite in sign to the true demand elasticity. However, the standard KY19 estimation rules out this possibility by assumption: it requires $\beta_{0,i,t} < 1$, which forces estimated own-price elasticities to be positive (i.e., demand curves to be downward sloping). In this section, we reestimate their demand equations (39) without imposing this restriction to assess how this assumption affects the estimation results.

We consider the sample of all 13F institutions, following the two-tier estimation procedure of KY19. For each quarter, institutions with at least 1,000 stock holdings are estimated individually. Institutions with fewer holdings are pooled into bins of approximately 2,000 holdings each, formed within each institution type and sorted by AUM, with a common slope coefficient estimated per bin (allowing institution-specific intercepts). The sample period is 1980Q1–2017Q4. We follow KY19's classification of institutions into banks, insurance companies, investment advisors, mutual funds, and pension funds. The "banks" category includes large institutional asset managers such as State Street and Norges Bank, as well as commercial banks such as JP Morgan and Bank of America that exert investment discretion over client equity holdings.

According to KY19, for each date-manager pair (i,t), the point estimate of $\hat{\beta}_{0,i,t}$ can be used to compute the estimated own-price elasticities of asset n as follows:

$$\hat{\varepsilon}_{i,t}(n) = 1 - \hat{\beta}_{0,i,t} \cdot (1 - \theta_{i,t}(n)), \tag{42}$$

where $\theta_{i,t}(n)$ is the portfolio weight of asset n. To weight institutions by their economic exposure within an institution type (or group, indexed by g), we define institution i's AUM share within group g as $a_{i,t} \equiv \text{AUM}_{i,t} / \sum_{j \in \mathscr{I}_{g,t}} \text{AUM}_{j,t}$, where $\mathscr{I}_{g,t}$ denotes group-g institutions active in quarter t and $\text{AUM}_{i,t}$ is institution i's AUM at time t. Throughout this section, AUM-weighted statistics of institution-level objects (e.g., $\hat{\beta}_{0,i,t}$) weight institution i by $a_{i,t}$, and value-weighted statistics of position-level objects (e.g., $\hat{\varepsilon}_{i,t}(n)$) weight position (i,n) by $a_{i,t}\,\theta_{i,t}(n)$, the dollar share of position (i,n) in total group-g AUM.

Table 5 reports these statistics for both $\hat{\beta}$ and $\hat{\varepsilon}$. Panel A shows the AUM-weighted distribution of institution-level $\hat{\beta}_{0,i,t}$: the AUM-weighted share of institutions with $\hat{\beta}_{0,i,t} > 1$ ranges from 18% (investment advisors) to 52% (banks). Panel B shows the value-weighted distribution of position-level $\hat{\varepsilon}_{i,t}(n)$: the value-weighted share of holdings with $\hat{\varepsilon}_{i,t}(n) < 0$ ranges from 17% (investment advisors) to 51% (banks). For banks, mutual funds, insurance companies, and pension

funds, the value-weighted 25$^{\text{th}}$ percentile of estimated elasticities lies at or below zero, indicating that negative estimates arise not only at the tails but throughout a substantial central portion of the value-weighted distribution. To gauge the price elasticity of the market overall, value-weighting is the relevant aggregation: each holding is weighted by its dollar magnitude rather than treated symmetrically with a much smaller one. Table 14 in Appendix C reports the corresponding unweighted statistics, where the equally-weighted fractions are much smaller, reflecting the dominance of the many small institutions and small positions in the sample.

Table 5: Estimated Betas and Demand Elasticities. The table reports value-weighted summary statistics for estimated betas and demand elasticities by institution type, based on unconstrained estimation (i.e., without imposing $\beta_{0,i,t} < 1$). Panel A reports the AUM-weighted distribution of institution-level $\hat{\beta}_{0,i,t}$, weighting institution i by its AUM share within group g, $a_{i,t}$. Panel B reports the value-weighted distribution of position-level $\hat{\varepsilon}_{i,t}(n)$, weighting position (i,n) by $a_{i,t}\,\theta_{i,t}(n)$, its dollar share in total group-g AUM. The sample consists of all 13F institutions over 1980Q1–2017Q4 (31 evenly-spaced quarters): institutions with at least 1,000 stock holdings are estimated individually, and institutions with fewer holdings are pooled into bins of approximately 2,000 holdings each (formed within each type-quarter and sorted by AUM, with a common slope coefficient estimated per bin and institution-specific intercepts). For each quarter, distribution moments are computed within each institution type; cells report cross-quarter averages of these moments.

Statistic	Banks	Insurance Companies	Investment Advisors	Mutual Funds	Pension Funds
Panel A: Estimated $\hat{\beta}_{0,i,t}$					
Mean	0.91	0.79	0.60	0.98	0.90
Std	0.65	0.51	0.46	0.46	0.39
p_{25}	0.51	0.43	0.28	0.69	0.67
p_{50}	0.97	0.87	0.58	0.96	0.91
p_{75}	1.35	1.18	0.91	1.25	1.11
Value share, $\hat{\beta} > 1$	0.52	0.36	0.18	0.47	0.39
Mean $\hat{\beta}$ if $\hat{\beta} > 1$	1.48	1.40	1.39	1.37	1.35
Panel B: Estimated Elasticities $\hat{\varepsilon}_{i,t}(n)$					
Mean	0.10	0.21	0.42	0.04	0.11
Std	0.64	0.50	0.46	0.46	0.39
p_{25}	−0.32	−0.17	0.11	−0.23	−0.10
p_{50}	0.05	0.14	0.44	0.05	0.10
p_{75}	0.50	0.57	0.74	0.32	0.34
Value share, $\hat{\varepsilon}_{i,t}(n) < 0$	0.51	0.34	0.17	0.46	0.38

Time-series of estimates by institution type. Figure 5 reports quarterly time series of these statistics for each 13F institution type. Panel (a) plots the value-weighted mean of $\hat{\varepsilon}_{i,t}(n)$ within each type, and Panel (b) plots the value-weighted share of (i,n) positions with $\hat{\varepsilon}_{i,t}(n) < 0$ within

each type. The two panels tell a consistent story. For four of the five institution types—banks, insurance companies, mutual funds, and pension funds—average estimated elasticities decline substantially over the sample, transitioning from positive averages near or above 0.5 in the early 1980s to averages near or below zero from the early 2000s onward. Correspondingly, the value-weighted share of negative-elasticity positions rises sharply: by the end of the sample it reaches roughly 85% for banks, around 70% for insurance companies and mutual funds, and 55% for pension funds. Investment advisors are the only group that retains a persistently positive value-weighted mean elasticity, with the share of negative-elasticity positions remaining around 25%. The prevalence of negative estimated elasticities under the unconstrained KY19 procedure is thus a robust feature of the data across major sectors of institutional ownership and rises over time. These patterns are consistent with the KY19 instrument recovering a shifter-process elasticity that is opposite in sign to the true structural demand elasticity for a substantial value share of institutional ownership.

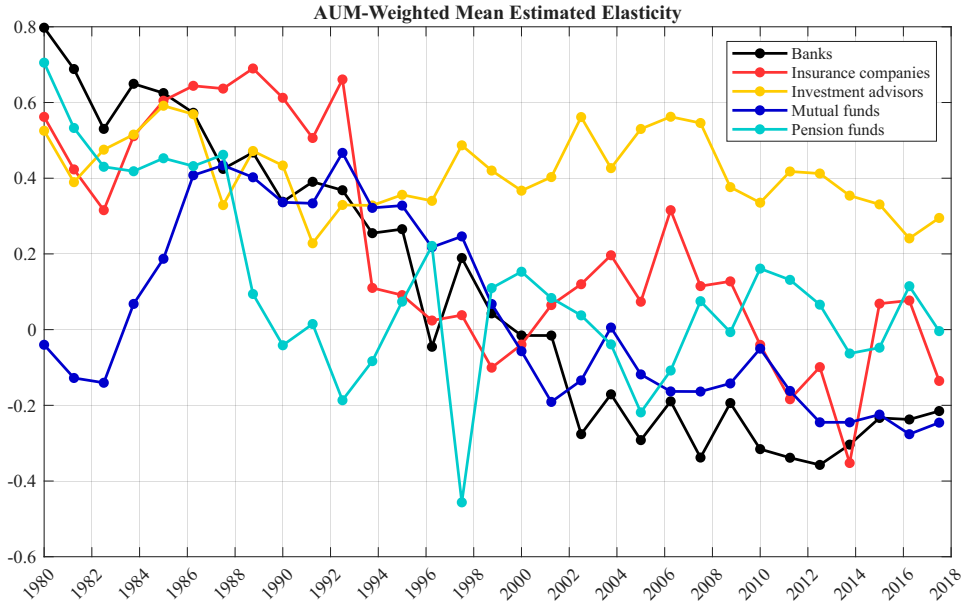
4.4 Discussion on Learning and Beliefs

The tests above take the view that investors would, at best, respond to return predictability associated with the instrument only to the extent that it materialized ex post over the full sample period. A simpler explanation for a lack of responsiveness is that investors may never have attempted to investigate a signal equivalent to the instrument during the sample period, and thus would naturally not react to this variation, implying an observed elasticity around zero. Moreover, even if investors had explored such a signal, detecting statistically significant patterns would have been even harder given the shorter samples they had access to within our overall sample period. This, again, would result in a lack of responsiveness.

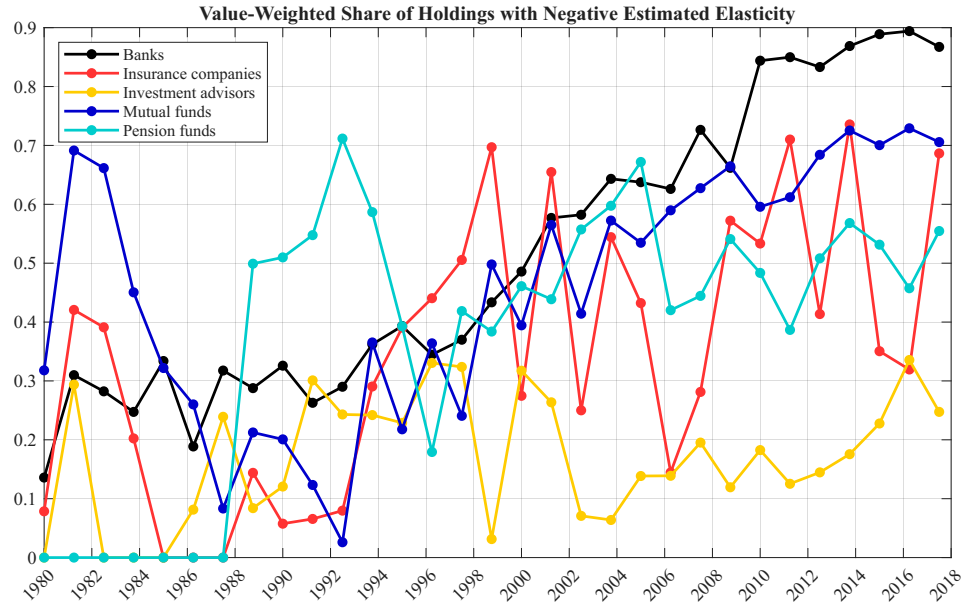
The only scenarios in which identification could succeed are ones in which investors held one of two forms of biased beliefs. First, period after period, they would have to believe that they would never trade again, effectively behaving as buy-and-hold investors, even though empirically they traded repeatedly. Second, they would have to believe that the cross-sectional price shifts induced by $\widehat{me}$ would fully resolve within each trading period, even though this consistently failed to occur throughout the sample period.

5 Conclusion

In this paper, we demonstrate that prevailing approaches to estimating asset demand elasticities can generate substantial and systematic bias when applied to financial markets, whose "products" are dynamically traded. When instruments induce price shifts that are persistent or gradually build up over time, the exclusion restrictions necessary for identification are violated, leading to severely



(a) Value-weighted mean of $\hat{\varepsilon}_{i,t}(n)$ by institution type.



(b) Value-weighted share of holdings with $\hat{\varepsilon}_{i,t}(n) < 0$ by institution type.

Figure 5: Estimated own-price elasticities by institution type, 1980Q1–2017Q4. Panel (a) plots, for each institution type and quarter, the value-weighted mean of $\hat{\varepsilon}_{i,t}(n)$, weighting each (i,n) position by $a_{i,t}\,\theta_{i,t}(n)$. Panel (b) plots the corresponding value-weighted share of (i,n) positions with $\hat{\varepsilon}_{i,t}(n) < 0$. The series consist of quarterly observations evenly spaced over the sample period.

biased price-elasticity estimates. Our theoretical framework clarifies the specific conditions under which price elasticities can be meaningfully estimated, quantifies the magnitude of the bias when these conditions fail, and shows when approximate bias corrections can be applied. Empirically, we show that widely used instruments, such as KY19's instrumented market equity, exhibit the types of dynamic patterns that violate the exclusion restriction.

A central force underlying our analysis is the disconnect between return opportunities and price levels: price-level shifts that last for multiple trading dates may generate little to no change in the relevant return distribution, thereby giving even highly elastic, frictionless investors little reason to trade. This fact implies that persistent supply shocks generated by liquidity traders, irrational traders, or policy interventions can generate substantial price-level shifts even when the rest of the market is populated by highly elastic investors. Using a dynamic general-equilibrium framework, we characterize the dynamic asset-market interventions required to induce targeted stochastic price-support processes in equilibrium, with direct implications for the design of central-bank asset-purchase programs. These results reveal significant shortcomings in standard multiplier calculations, even in the most stylized cases. More broadly, our results call for a reassessment of how existing elasticity estimates should be used to inform asset-pricing models, evaluate policy interventions, and draw conclusions about financial-market frictions.

References

An, Yu and Amy W. Huber, "Demand Propagation Through Traded Risk Factors," Technical Report September 2025. Working Paper.

Basak, Suleyman and Domenico Cuoco, "An Equilibrium Model with Restricted Stock Market Participation," *The Review of Financial Studies*, 1998, *11* (2), 309–341.

Beneish, Messod D. and Robert E. Whaley, "An Anatomy of the "S&P Game": The Effects of Changing the Rules," *The Journal of Finance*, 1996, *51* (5), 1909–1930.

Bhamra, Harjoat S. and Raman Uppal, "Asset Prices with Heterogeneity in Preferences and Beliefs," *The Review of Financial Studies*, 2014, *27* (2), 519–580.

Binsbergen, Jules H. van and Christian C. Opp, "Real Anomalies," *The Journal of Finance*, August 2019, *74* (4), 1659–1706.

_ , Martijn Boons, Christian C. Opp, and Andrea Tamoni, "Dynamic Asset (Mis)Pricing: Build-up versus Resolution Anomalies," *Journal of Financial Economics*, February 2023, *147* (2), 406–431.

Borovička, Jaroslav, "Survival and Long-Run Dynamics with Heterogeneous Beliefs under Recursive Preferences," *Journal of Political Economy*, 2020, *128* (1), 206–251.

Brunnermeier, Markus K. and Yuliy Sannikov, "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 2014, *104* (2), 379–421.

Chang, Yen-cheng, Harrison Hong, and Inessa Liskovich, "Regression Discontinuity and the Price Effects of Stock Market Indexing," *The Review of Financial Studies*, 2015, *28* (1), 212–246.

Chen, Jiawei, Susanna Esteban, and Matthew Shum, "Demand and Supply Estimation Biases Due to Omission of Durability," *Journal of Econometrics*, 2008, *147* (2), 247–257.

Davis, Carter, Mahyar Kargar, and Jiacui Li, "Why Do Portfolio Choice Models Predict Inelastic Demand?," May 2025. Unpublished manuscript.

Dumas, Bernard, "Two-Person Dynamic Equilibrium in the Capital Market," *The Review of Financial Studies*, 1989, *2* (2), 157–188.

Dynkin, E. B., *Markov Processes*, Berlin, Heidelberg: Springer-Verlag, 1965. Two volumes; Die Grundlehren der Mathematischen Wissenschaften, Vols. 121–122.

Fuchs, William, Satoshi Fukuda, and Daniel Neuhann, "Demand-System Asset Pricing: Theoretical Foundations," 2025. Working paper.

_ , _ , and _ , "A Trilemma for Asset Demand Estimation," 2025. Working paper.

Fukasawa, Takeshi, "The Biases in Applying Static Demand Models Under Dynamic Demand," *Review of Industrial Organization*, 2024, *65* (2), 561–594.

Gabaix, Xavier and Ralph S.J. Koijen, "In Search of the Origins of Financial Fluctuations: The

Inelastic Markets Hypothesis," Working Paper w28967, National Bureau of Economic Research December 2023.

Gârleanu, Nicolae and Stavros Panageas, "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing," *Journal of Political Economy*, 2015, *123* (3), 670–685.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Gowrisankaran, Gautam and Marc Rysman, "Dynamics of consumer demand for new durable goods," *Journal of Political Economy*, 2012, *120* (6), 1173–1219.

Greenwood, Robin and Dimitri Vayanos, "Bond Supply and Excess Bond Returns," *Review of Financial Studies*, mar 2014, *27* (3), 663–713.

_ **and Marco Sammon**, "The Disappearing Index Effect," *Journal of Finance*, April 2025, *80* (2).

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, June 1980, *70* (3), 393–408.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, "Empirical Asset Pricing via Machine Learning," *The Review of Financial Studies*, 02 2020, *33* (5), 2223–2273.

Haddad, Valentin, Zhiguo He, Paul Huebner, Peter Kondor, and Erik Loualiche, "Causal Inference for Asset Pricing," 2025. Working paper.

Harris, Lawrence E. and Eitan Gurel, "Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures," *The Journal of Finance*, September 1986, *41* (4), 815–829.

Hasbrouck, Joel, "Measuring the Information Content of Stock Trades," *The Journal of Finance*, 1991, *46* (1), 179–207.

He, Zhiguo and Arvind Krishnamurthy, "Intermediary Asset Pricing," *American Economic Review*, 2013, *103* (2), 732–770.

_ **, Peter Kondor, and Jessica Shi Li**, "Demand Elasticity in Dynamic Asset Pricing," June 2025. Unpublished manuscript.

Heaton, John and Deborah J. Lucas, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 1996, *104* (3), 443–487.

Hendel, Igal and Aviv Nevo, "Measuring the implications of sales and consumer inventory behavior," *Econometrica*, 2006, *74* (6), 1637–1673.

Kocherlakota, Narayana R., "Log-Linear Relative Asset Demand," NBER Working Paper 34395, National Bureau of Economic Research October 2025.

Koijen, Ralph S. J. and Motohiro Yogo, "A Demand System Approach to Asset Pricing," *Journal*

of Political Economy, apr 2019, *127* (4), 1475–1515.

_ **and** _ , "Asset Demand Systems in Macro-Finance," *NBER Reporter*, 2021, *4*, 21–24.

Koijen, Ralph S.J. and Motohiro Yogo, "On the Theory and Econometrics of (Demand System) Asset Pricing," May 30 2025. Working paper.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

Lynch, Anthony W. and Richard R. Mendenhall, "New Evidence on Stock Price Effects Associated with Changes in the S&P 500 Index," *The Journal of Business*, July 1997, *70* (3), 351–383.

Merton, Robert C., "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, August 1969, *51* (3), 247–257.

Milgrom, Paul and Nancy Stokey, "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 1982, *26* (1), 17–27.

Øksendal, Bernt, *Stochastic Differential Equations: An Introduction with Applications* Universitext, 6 ed., Berlin, Heidelberg: Springer, 2003.

Petajisto, Antti, "Why Do Demand Curves for Stocks Slope Down?," *Journal of Financial and Quantitative Analysis*, 2009, *44* (5), 1013–1044.

Schiraldi, Pasquale, "Automobile Replacement: A Dynamic Structural Approach," *RAND Journal of Economics*, 2011, *42* (2), 266–291.

Shleifer, Andrei, "Do Demand Curves for Stocks Slope Down?," *The Journal of Finance*, July 1986, *41* (3), 579–590.

Wurgler, Jeffrey and Ekaterina Zhuravskaya, "Does Arbitrage Flatten Demand Curves for Stocks?," *The Journal of Business*, 2002, *75* (4), 583–608.

Appendix

A Additional Analyses and Proofs

A.1 The Case of Log Utility

For the case of log utility we conjecture and verify that the value function takes the form

$$V(W,s) = \upsilon \log W + A(s), \tag{43}$$

with $\upsilon \in (0,1)$ chosen so that the HJB has no residual dependence on $\log W$. Substituting $V_W = \upsilon/W$ and $V_{WW} = -\upsilon/W^2$ into (7), using $W' = W\big(1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}\big)$ and $u(W') = \log W'$, gives

$$\begin{aligned} 0 = \max_{\boldsymbol{\theta}} \Big\{ & -\rho\big[\upsilon \log W + A(s)\big] + \upsilon\Big[r_f + (\boldsymbol{\mu} - r_f \mathbf{1})^\top \boldsymbol{\theta} - \tfrac{1}{2}\boldsymbol{\theta}^\top \boldsymbol{\Sigma} \boldsymbol{\theta}\Big] \\ & + \lambda \sum_{s'} q(s'|s)\Big[\pi_C\big(\log W + \log(1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}) - \upsilon \log W - A(s)\big) \\ & + (1-\pi_C)\big(\upsilon \log W + \upsilon \log(1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}) + A(s') - \upsilon \log W - A(s)\big)\Big]\Big\}. \end{aligned}$$

Collecting the $\log W$ terms yields the coefficient

$$-\rho\,\upsilon + \lambda\,\pi_C(1-\upsilon) = 0 \quad \Rightarrow \quad \upsilon = \frac{\lambda\,\pi_C}{\rho + \lambda\,\pi_C}.$$

The remaining terms (independent of $\log W$) give the system solved by $A(s)$:

$$\begin{aligned} 0 \;=\; & -\rho A(s) \;+\; \upsilon\Big[r_f + (\boldsymbol{\mu} - r_f\mathbf{1})^\top \boldsymbol{\theta} - \tfrac{1}{2}\boldsymbol{\theta}^\top \boldsymbol{\Sigma}\boldsymbol{\theta}\Big] \\ & + \lambda \sum_{s'} q(s'|s)\Big[(1-\pi_C)A(s') - A(s) + \widetilde{C}\cdot \log\big(1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}\big)\Big], \end{aligned} \tag{44}$$

where we define:

$$\widetilde{C} \equiv \pi_C + (1-\pi_C)\,\upsilon. \tag{45}$$

Differentiating (44) with respect to each θ_j yields the first-order conditions

$$0 \;=\; \upsilon\Big[(\mu_j - r_f) - (\boldsymbol{\Sigma}\boldsymbol{\theta})_j\Big] + \lambda\,\widetilde{C}\sum_{s'} q(s'|s)\,\frac{\kappa^{(j)}_{s,s'}}{1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}}, \qquad j = 1,\ldots,J. \tag{46}$$

A.2 Proof of Propositions 1 and 2

Suppose that θ_C^* is the optimal portfolio share in the economy with continuous trade and θ_D^* the optimal portfolio share when agents can trade only upon the discrete event dates τ_n. Below we derive the results of Propositions 1 and 2 jointly. As stated in the Propositions, we consider one risky asset. We first consider the case of continuous trading. Thereafter, we analyze the setting where trading can only occur at the event dates τ_n.

A.2.1 Continuous Trading

We first consider the continuous-trade economy and derive the sensitivity of the optimal risky-asset share $\theta_C^*(s)$ with respect to the exposure to the shifter state

$$\left.\frac{d\theta_C^*(s)}{d\beta}\right|_{\beta=0} \tag{47}$$

under the assumption of three shifter states $\{-1,0,+1\}$ with transition probabilities

$$q(0\,|1) = q(0\,|-1) = \pi_R, \quad q(1,|1) = q(-1\,|-1) = 1-\pi_R, \quad q(0\,|0) = 1,$$

where all other Markov transition rates are zero. We consider only one risky asset. Thus, we drop the subscript "1" on $\mu_1, \sigma_{1,A}, \sigma_{1,I}$. Define the total variance

$$\sigma^2 \equiv \sigma_A^2 + \sigma_I^2.$$

For any shifter state $s \in \{-1,+1\}$, the return if the shifter state jumps to $s' = 0$ is:

$$\kappa(s) \equiv \kappa_{s,0} = e^{\beta\,(0-s)} - 1 = e^{-\beta s} - 1.$$

Note that $\kappa(0) = 0$. Under the stated assumptions, the first-order condition (11) can be written as

$$0 = A(s)\big[\mu - r_f - \gamma\sigma^2\theta_C(s)\big] + \lambda\,\pi_R C(0)\,\kappa(s)\big[1+\theta_C(s)\,\kappa(s)\big]^{-\gamma}, \tag{48}$$

where we define:

$$C(s) \equiv \pi_C + (1-\pi_C)A(s). \tag{49}$$

Note that $\kappa_{s,s} = e^{\beta(s-s)} - 1 = 0$, so all terms in equation (11) where the state does not change ($s' = s$) drop out.

No shift $s = 0$. Because the state $s = 0$ is absorbing, the FOC at $s = 0$ contains no jump term and yields the standard Merton share

$$\theta_C^*(0) = \frac{\mu - r_f}{\gamma\sigma^2}.$$

Further, $A(0)$ is independent of π_R because it is determined entirely in the absorbing state $s = 0$.

State-independence at $\beta = 0$. When $\beta = 0$ we have $\kappa_{s,s'} = 0$ for all states. Then the FOC reduces to

$$\mu - r_f - \gamma\sigma^2\,\theta_C^*(s) = 0 \tag{50}$$

so $\theta_C^*(s) = \theta_C^*(0)$ is independent of s. Substituting this θ_C^* into (10) shows that a constant function $A(s) \equiv A(0)$ solves the corresponding linear system. Evaluating it at the absorbing state $s = 0$ pins down $A(0)$ via (10) (see the expression derived below).

Shifter states $s \in \{-1, +1\}$. Define

$$f(\theta, \beta; s) \equiv A(s)\big[\mu - r_f - \gamma\sigma^2\theta\big] + \lambda\,\pi_R C(0)\,\kappa(s)\,[1 + \theta\,\kappa(s)]^{-\gamma}, \tag{51}$$

The FOC (48) is $f(\theta_C(s), \beta; s) = 0$. At $\beta = 0$, we have $\kappa(s) = 0$, so

$$f(\theta, 0; s) = A(s)\big[\mu - r_f - \gamma\sigma^2\,\theta\big], \tag{52}$$

yields $\theta_0 \equiv \theta_C^*(0) = \frac{\mu - r_f}{\gamma\sigma^2}$. Note that whereas $A(s)$ can depend on β through the equilibrium system (for $s \in \{-1, +1\}$), $C(0)$ is independent of β since $s = 0$ is an absorbing state. To apply the implicit function theorem, we compute the partial derivatives $\partial_\theta f$ and $\partial_\beta f$ at $(\theta_0, 0)$:

$$\partial_\theta f(\theta_0, 0; s) = -A(s)\gamma\sigma^2. \tag{53}$$

Moreover, as shown above, at $\beta = 0$ we have $A(s) = A(0)$. For $\partial_\beta f$, the first term in (51) contributes:

$$\frac{\partial A(s)}{\partial \beta}\Big|_{\beta=0} \underbrace{\big[\mu - r_f - \gamma\sigma^2\theta_0\big]}_{=0} = 0, \tag{54}$$

because the FOC implies $\mu - r_f - \gamma\sigma^2\theta_0 = 0$. The derivative of the second term in (51) is:

$$\lambda\,\pi_R C(0)\left[\frac{\partial \kappa}{\partial \beta} \cdot [1 + \theta_0\kappa]^{-\gamma} + \kappa \cdot (-\gamma)[1 + \theta_0\kappa]^{-\gamma-1} \cdot \theta_0 \cdot \frac{\partial \kappa}{\partial \beta}\right]. \tag{55}$$

Since $\kappa(s) = e^{-\beta s} - 1$ and $\frac{\partial \kappa}{\partial \beta} = -s \cdot e^{-\beta s}$, we obtain $\kappa|_{\beta=0} = 0$, $\frac{\partial \kappa}{\partial \beta}|_{\beta=0} = -s$, and $[1+\theta_0 \kappa]|_{\beta=0} = 1$. Overall, this yields:

$$\partial_\beta f(\theta_0, 0; s) = \lambda \pi_R C(0) \cdot (-s) \cdot 1 = -\lambda \pi_R C(0)\, s. \tag{56}$$

By the implicit function theorem we then obtain:

$$\left.\frac{d\theta_C^*(s)}{d\beta}\right|_{\beta=0} = -\left.\frac{\partial_\beta f}{\partial_\theta f}\right|_{(\theta_0,0)} = -\frac{\lambda\, \pi_R C(0)}{A(0)\, \gamma \sigma^2}\, s. \tag{57}$$

It follows that the shifter-process elasticity is

$$\eta_C(s) = -\frac{1}{s} \left.\frac{\frac{d\theta_C^*(s)}{d\beta}}{\theta_C^*(0)}\right|_{\beta=0} = \pi_R \frac{\lambda\, C(0)}{(\mu - r_f) A(0)}. \tag{58}$$

With continuous trade, the shifter-process elasticity η_C is equal to the price elasticity ε_C when the shifter process feature instantaneous and deterministic resolution over the first instant, which obtains in the limit when $\lambda \to \infty$ and $\pi_R = 1$. In this limit, the formula for $A(0)$ derived below gives $A(0) \to 1$, and hence $C(0) = \pi_C + (1 - \pi_C) A(0) \to 1$. Therefore:

$$\eta_C \to \infty \quad \text{as } \lambda \to \infty. \tag{59}$$

The true price elasticity under continuous trade is the value of that limit, that is, $\varepsilon_C = \infty$.

Deriving $A(0)$. When state $s = 0$ is an absorbing state, equation (10) simplifies to:

$$\begin{aligned} 0 = &-\rho A(0) + (1-\gamma) A(0) \left[r_f + (\mu - r_f)\theta(0) - \tfrac{1}{2}\gamma \theta(0)^2 \sigma^2 \right] \\ &+ \lambda \Big[\big[\pi_C + (1-\pi_C) A(0) \big] - A(0) \Big], \end{aligned} \tag{60}$$

and solving for $A(0)$ yields:

$$A(0) = \frac{\lambda \pi_C}{\lambda \pi_C + \rho - (1-\gamma)\left[r_f + (\mu - r_f)\theta(0) - \tfrac{1}{2}\gamma \theta(0)^2 \sigma^2 \right]}. \tag{61}$$

Further, the optimal portfolio share is given by the Merton demand $\theta_C^*(0) = \frac{\mu - r_f}{\gamma\sigma^2}$. Substituting this result yields:

$$A(0) = \frac{\lambda \pi_C}{\lambda \pi_C + \rho - (1-\gamma)\left[r_f + \frac{(\mu - r_f)^2}{2\gamma\sigma^2}\right]}. \tag{62}$$

A.2.2 Discrete Trading

To keep CRRA utility well-defined under discrete trading, we restrict attention to policies that keep wealth strictly positive, which corresponds to the restriction that at each event date a risky asset weight $\theta_D \in [0,1]$ is chosen. Let the risk-free asset price be $P_{0,t}$ with $dP_{0,t}/P_{0,t} = r_f dt$. At an event time τ, a CRRA investor chooses a weight $\theta_D \in [0,1]$ that corresponds to fixed units

$$n_1 \;=\; \frac{\theta_D W_\tau}{P_\tau}, \qquad n_0 \;=\; \frac{(1-\theta_D) W_\tau}{P_{0,\tau}}. \tag{63}$$

For $t \in [0, \Delta]$, we obtain:

$$\begin{aligned} W_t \;&= n_1 P_t \;+\; n_0 P_{0,t} \;=\; \theta_D W_\tau \frac{P_t}{P_\tau} \;+\; (1-\theta_D) W_\tau \frac{P_{0,t}}{P_{0,\tau}} \\ &= \theta_D W_\tau \frac{P_t}{P_\tau} \;+\; (1-\theta_D) W_\tau e^{r_f(t-\tau)} \;>\; 0 \quad \text{a.s.} \end{aligned} \tag{64}$$

since $P_t > 0$ and $P_{0,t} > 0$ almost surely. Let

$$\theta_C(0) \;=\; \frac{\mu - r_f}{\gamma\sigma^2} \in (0,1) \tag{65}$$

be the Merton share under continuous trade and $s = 0$. We assume there exist $\underline{\lambda} < \infty$ and $\bar{\beta} > 0$ such that for all $\lambda \geq \underline{\lambda}$ and all $|\beta| \leq \bar{\beta}$ the optimal discrete-trade choice remains interior:

$$\theta_D^*(s; \lambda, \beta) \in (0,1) \qquad \text{for all } s \in \{-1, 0, +1\}. \tag{66}$$

Note that the perturbation we study below is local in β and our results below show that:

$$\theta_D^*(s; \lambda, \beta) = \theta_C(0) \;+\; O\left(\frac{1}{\lambda}\right) \;+\; O(\beta), \tag{67}$$

so for large enough λ and small enough $|\beta|$ the constraint $\theta_D \in [0,1]$ is slack. Correspondingly, the unconstrained FOC and the elasticity results derived below apply.

Given a weight θ_D is chosen at the last event date, the post-event weight after time t has passed

is given by:

$$\theta_t \equiv \frac{\theta_D P_t}{\theta_D P_t + (1-\theta_D)e^{r_f t}}.$$

Then $\{\theta_t\}_{t\geq 0}$ follows

$$d\theta_t = \theta_t(1-\theta_t)\big(\mu - r_f - \sigma^2\theta_t\big)\,dt + \sigma\,\theta_t(1-\theta_t)\,dB_t, \qquad \theta_0 = \theta_D. \tag{68}$$

Given the above assumptions, the post-event weight satisfies $\theta_t \in (0,1)$ for $t \in [\,\tau, \tau+\Delta]$, that is, it is bounded, implying that

$$\mathbb{E}\left[\int_0^\Delta \theta_t^2\,dt\right] \;\leq\; \mathbb{E}[\Delta] \;=\; \frac{1}{\lambda} \;<\infty, \tag{69}$$

and the Itô integral against θ_t is square-integrable.

The discrete-trade HJB differs from the continuous-trade case because the portfolio weight drifts between event times. We derive two components: (i) expected utility growth over the inter-event period $[0,\Delta]$, accounting for the drifting weight θ_t; and (ii) a jump term at event arrival, which depends on θ_Δ, the weight at the moment of the jump (after a time Δ has passed), not the initially chosen weight θ_D. We again start with the value function taking the form $V(W,s) = u(W)A(s)$.

Expected utility growth. We first derive an expression for $\mathbb{E}\left[\frac{u(W_\Delta)-u(W_0)}{u(W_0)}\right]$. Itô's lemma gives

$$\frac{du(W_t)}{u(W_t)} = \alpha_t dt + \zeta_t dB_t. \tag{70}$$

where we define:

$$\alpha_t \equiv (1-\gamma)\Big[r_f + \theta_t(\mu - r_f) - \tfrac{1}{2}\,\gamma\sigma^2\theta_t^2\Big], \tag{71}$$

$$\zeta_t \equiv (1-\gamma)\sigma\,\theta_t. \tag{72}$$

Using the SDE for $u(W_t)$, dividing by $u(W_0)$, and integrating over $[0,\Delta]$ gives

$$\frac{u(W_\Delta)-u(W_0)}{u(W_0)} = \int_0^\Delta \frac{u(W_t)}{u(W_0)}\alpha_t dt + \int_0^\Delta \frac{u(W_t)}{u(W_0)}\zeta_t dB_t. \tag{73}$$

Let $Z_t \equiv u(W_t)/u(W_0)$. The Itô integral $\int_0^\tau Z_t \zeta_t \, dB_t$ is a square-integrable martingale with mean zero. With $\Delta \sim \text{Exp}(\lambda)$ independent of the Brownian shocks, we obtain:

$$\mathbb{E}\left[\int_0^\Delta \frac{u(W_t)}{u(W_0)} \zeta_t dB_t\right] = \mathbb{E}\left[\int_0^\Delta Z_t \, \zeta_t \, dB_t\right] = \mathbb{E}\left[\, \mathbb{E}\left[\int_0^\Delta Z_t \, \zeta_t \, dB_t \,\middle|\, \Delta\right] \,\right] = 0. \tag{74}$$

Then

$$\lambda \, \mathbb{E}\left[\frac{u(W_\Delta) - u(W_0)}{u(W_0)}\right] = \lambda \, \mathbb{E}\int_0^\Delta Z_t \, \alpha_t \, dt = \lambda \int_0^\infty e^{-\lambda t} \, \mathbb{E}[Z_t \alpha_t] \, dt. \tag{75}$$

Next, we split $\mathbb{E}[Z_t \alpha_t] = \mathbb{E}[\alpha_t] + R(t)$, where we define the remainder

$$R(t) \equiv \mathbb{E}[(Z_t - 1)\alpha_t], \tag{76}$$

and evaluate the order of this remainder. Since $\theta_t \in (0,1)$, there exists a scalar $F < \infty$ such that $|\alpha_t| \leq F$. Define $m(t) \equiv \mathbb{E}[Z_t]$. Since

$$dZ_t = Z_t \alpha_t \, dt + Z_t \zeta_t \, dB_t \tag{77}$$

and $\mathbb{E}\int_0^\tau (Z_t \zeta_t)^2 dt < \infty$, we obtain:

$$m(t) = 1 + \int_0^t \mathbb{E}[Z_i \alpha_i] \, di \;\leq\; 1 + F \int_0^t m(i) \, di, \tag{78}$$

so by Grönwall, $m(t) \leq e^{Ft}$ (see Øksendal, 2003). Consequently,

$$|R(t)| = \left|\mathbb{E}[(Z_t - 1)\alpha_t]\right| \leq F \, \mathbb{E}|Z_t - 1| \leq F \, (e^{F \cdot t} - 1). \tag{79}$$

As a result, we obtain:

$$\lambda \int_0^\infty e^{-\lambda t} \, |R(t)| \, dt \;\leq\; F \, \lambda \int_0^\infty e^{-\lambda t} \left(e^{F \cdot t} - 1\right) dt \;=\; \frac{F^2}{\lambda - F} \;=\; O\Big(\frac{1}{\lambda}\Big), \tag{80}$$

where $\lambda > F$ is required to ensure that the integral is finite. Combining (75) with the remainder bound yields

$$\lambda \int_0^\infty e^{-\lambda t} \, \mathbb{E}[Z_t \alpha_t] \, dt = \lambda \int_0^\infty e^{-\lambda t} \, \mathbb{E}[\alpha_t] \, dt \;+\; O(\lambda^{-1}) = \lambda \, \mathbb{E}\int_0^\Delta \alpha_t \, dt \;+\; O(\lambda^{-1}).$$

Substituting the definition of α_t and recalling that $\theta_0 = \theta_D$ is the weight set at the last event, we obtain

$$\begin{aligned}
&\lambda\, \mathbb{E}\left[\frac{u(W_\Delta) - u(W_0)}{u(W_0)}\right] \\
=&\lambda(1-\gamma)\, \mathbb{E}\left[r_f \Delta + (\mu - r_f)\int_0^\Delta \theta_t\, dt \;-\; \tfrac{1}{2}\gamma\sigma^2 \int_0^\Delta \theta_t^2\, dt \;\middle|\; \theta_0 = \theta_D\right] + O(\lambda^{-1}).
\end{aligned} \tag{81}$$

We now evaluate the expectations of the integrals involving θ_s and θ_s^2. To do so, we utilize Dynkin's formula (see Dynkin, 1965; Øksendal, 2003). For the Markov process (θ_t) with generator $\mathscr{L}$, we obtain for any $\varphi \in C^2([0,1])$

$$\mathbb{E}_\theta[\varphi(\theta_t)] = \varphi(\theta) + \mathbb{E}_\theta \int_0^t \mathscr{L}\varphi(\theta_r)\, dr, \tag{82}$$

where the dynamics of θ given in (68) imply the generator

$$\mathscr{L}\varphi(x) = x(1-x)\big(\mu - r_f - \sigma^2 x\big)\, \varphi'(x) + \tfrac{1}{2}\sigma^2 x^2 (1-x)^2\, \varphi''(x). \tag{83}$$

The Dynkin formula yields the following expansions for small deterministic t:

$$\mathbb{E}_\theta[\varphi(\theta_t)] \;= \varphi(\theta) + t\,\mathscr{L}\varphi(\theta) + O(t^2) \tag{84}$$

$$\mathbb{E}_\theta \int_0^t \varphi(\theta_r) dr \;= t\,\varphi(\theta) + \tfrac{1}{2} t^2 \mathscr{L}\varphi(\theta) + O(t^3) \tag{85}$$

Integrating the first expansion (84) over $s \in [0,t]$ yields the second (85). Conditioning on $\theta_0 = \theta_D$ and Δ and taking expectations yields:

$$\begin{aligned}
\mathbb{E}\left[\int_0^\Delta \varphi(\theta_r)\, dr\right] &= \mathbb{E}\big[\Delta\,\varphi(\theta_D) + \tfrac{1}{2}\Delta^2\, \mathscr{L}\varphi(\theta_D) + O(\Delta^3)\big] \\
&= \frac{\varphi(\theta_D)}{\lambda} + \frac{\mathscr{L}\varphi(\theta_D)}{\lambda^2} + O(\lambda^{-3}),
\end{aligned} \tag{86}$$

where we use $\mathbb{E}[\Delta] = 1/\lambda$, $\mathbb{E}[\Delta^2] = 2/\lambda^2$. With $\varphi(x) = x$ and $\varphi(x) = x^2$ we obtain:

$$\mathscr{L}(x)(\theta_D) = \theta_D(1-\theta_D)(\mu - r_f - \sigma^2\theta_D) \tag{87}$$

$$\mathscr{L}(x^2)(\theta_D) = 2\theta_D^2(1-\theta_D)\big(\mu - r_f - \sigma^2\theta_D\big) + \sigma^2\theta_D^2(1-\theta_D)^2 \tag{88}$$

and we define $b(\theta_D) \equiv \mathscr{L}(x)(\theta_D)$. Conditioning on $\theta_0 = \theta_D$, we obtain:

$$\mathbb{E}\left[\int_0^\Delta \theta_t \, dt\right] = \frac{\theta_D}{\lambda} + \frac{b(\theta_D)}{\lambda^2} + O(\lambda^{-3}), \tag{89}$$

$$\mathbb{E}\left[\int_0^\Delta \theta_t^2 \, dt\right] = \frac{\theta_D^2}{\lambda} + \frac{2\theta_D^2(1-\theta_D)\big(\mu - r_f - \sigma^2\theta_D\big) + \sigma^2\theta_D^2(1-\theta_D)^2}{\lambda^2} + O(\lambda^{-3}). \tag{90}$$

Substituting (89), (90), and $\mathbb{E}[\Delta] = 1/\lambda$ into (81) yields:

$$\begin{aligned}\lambda\,\mathbb{E}\left[\frac{u(W_\Delta) - u(W_0)}{u(W_0)}\right] &= (1-\gamma)\Big[r_f + \theta_D \cdot (\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\Big] \\ &\quad + \frac{1-\gamma}{\lambda}\, g(\theta_D) + O(\lambda^{-1}),\end{aligned} \tag{91}$$

where we define the correction polynomial relative to continuous rebalancing:

$$\begin{aligned}g(\theta_D) \equiv{}& (\mu - r_f)\,\theta_D(1-\theta_D)\big((\mu - r_f) - \sigma^2\theta_D\big) \\ &- \gamma\sigma^2\theta_D^2(1-\theta_D)\big((\mu - r_f) - \sigma^2\theta_D\big) \;-\; \tfrac{1}{2}\,\gamma\sigma^4\theta_D^2(1-\theta_D)^2.\end{aligned} \tag{92}$$

Jump terms and marginal exposure to the shifter state. The second component of the HJB is the jump term. At an event arrival, in the case of shift resolution, wealth changes by factor $(1 + \theta_\Delta \kappa(s))$, where θ_Δ is the portfolio weight at the jump time, not the initially chosen θ_D. We now expand this term for small β. For a marginal exposure $d\beta$ with $s \in \{-1, 0, +1\}$, we can write:

$$\kappa(s) = e^{-\beta s} - 1 = -s\,d\beta + O(d\beta^2). \tag{93}$$

Using $(1+x)^{1-\gamma} = 1 + (1-\gamma)x + O(x^2)$ and independence of Δ, we obtain

$$\mathbb{E}\big[(1 + \theta_\Delta\kappa(s))^{1-\gamma} \mid \theta_0 = \theta_D\big] = 1 - (1-\gamma)\,\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D]\, s\,d\beta + O(d\beta^2), \tag{94}$$

and by the Dynkin expansion,

$$\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D] = \theta_D + \frac{b(\theta_D)}{\lambda} + O(\lambda^{-2}). \tag{95}$$

HJB and FOC. With $V(W,s) = u(W)A(s)$, the per-time HJB is

$$\begin{aligned} 0 = &-\rho A(s) + (1-\gamma)A(s)\left[r_f + \theta_D \cdot (\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\right] + \frac{1-\gamma}{\lambda}A(s)\,g(\theta_D) \\ &+ \lambda\Big\{\pi_R C(0)\,\mathbb{E}\big[(1+\theta_\Delta \kappa(s))^{1-\gamma} \mid \theta_0 = \theta_D\big] + (1-\pi_R)C(s) - A(s)\Big\} \\ &+ O(\lambda^{-1}) + O(\lambda\, d\beta^2), \end{aligned} \tag{96}$$

Substituting (94) and (95) into the HJB, we see that the jump term contributes a factor linear in $d\beta$ that depends on θ_D through $\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D]$. Differentiating w.r.t. θ_D and dividing by $(1-\gamma)A(s)$, yields the FOC:

$$\begin{aligned} 0 =&(\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} - \lambda\,\pi_R\,\frac{C(0)}{A(s)}\left(1 + \frac{b'(\theta_D)}{\lambda}\right) s\,d\beta \\ &+ O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1} d\beta). \end{aligned} \tag{97}$$

We compute the marginal elasticity by differentiating with respect to β and evaluating at $\beta = 0$. Note that the $O(\lambda \cdot d\beta^2)$ term vanishes under this operation: differentiation yields $O(\lambda \cdot d\beta)$, which is zero at $\beta = 0$.

Sensitivity to β at $\beta = 0$. Let θ_D^* denote the solution to the FOC. Applying the implicit function theorem (differentiating the FOC with respect to β and evaluating at $\beta = 0$) yields:

$$f_\theta \left.\frac{d\theta_D^*}{d\beta}\right|_0 + f_\beta = 0, \tag{98}$$

with

$$\begin{aligned} f_\theta &= -\gamma\sigma^2 + \frac{g''(\theta_D^*)}{\lambda} + O(\lambda^{-1}), \\ f_\beta &= -\lambda\,\pi_R\,\frac{C(0)}{A(0)}\left(1 + \frac{b'(\theta_D^*)}{\lambda}\right) s + O(\lambda^{-1}). \end{aligned} \tag{99}$$

Evaluating at the optimal solution at $\beta = 0$, denoted by $\theta_{D,0}^*$, gives

$$\left.\frac{d\theta_D^*}{d\beta}\right|_{\beta=0} = -\frac{f_\beta}{f_\theta} = -\frac{\lambda\,\pi_R C(0)}{A(0)\,\gamma\sigma^2}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right) s + O(\lambda^{-1}). \tag{100}$$

Normalizing by $\theta^*_{D,0}$ yields the shifter-process elasticity under discrete trade (for large but finite λ):

$$\begin{aligned}\eta_D(s) &= -\frac{1}{s}\left(\frac{1}{\theta^*_D}\frac{d\theta^*_D}{d\beta}\right)\Bigg|_{\beta=0} \\ &= \frac{\pi_R C(0)}{A(0)\,\gamma\sigma^2\theta^*_{D,0}}\left(\lambda + b'(\theta^*_{D,0}) + \frac{g''(\theta^*_{D,0})}{\gamma\sigma^2}\right) + O(\lambda^{-1}). \end{aligned} \tag{101}$$

By setting $\pi_R = 1$ we obtain the true price elasticity under discrete trade (for large but finite λ):

$$\varepsilon_D = \frac{C(0)}{A(0)\,\gamma\sigma^2\theta^*_{D,0}}\left(\lambda + b'(\theta^*_{D,0}) + \frac{g''(\theta^*_{D,0})}{\gamma\sigma^2}\right) + O(\lambda^{-1}). \tag{102}$$

Comparing (101) and (102), we see that $\eta_D = \pi_R\varepsilon_D + O(\lambda^{-1})$. Since $\varepsilon_D = O(\lambda)$, dividing yields:

$$\frac{\eta_D}{\varepsilon_D} = \pi_R + \frac{O(\lambda^{-1})}{\varepsilon_D} = \pi_R + O(\lambda^{-2}). \tag{103}$$

Recognizing that the estimated elasticity $\hat{\varepsilon}_D$ uncovers η_D, we obtain

$$\frac{\hat{\varepsilon}_D}{\varepsilon_D} = \frac{\eta_D}{\varepsilon_D} = \pi_R + O(\lambda^{-2}) \to \pi_R \quad \text{as } \lambda \to \infty. \tag{104}$$

A.3 Proof of Proposition 3

Because the price shift stays in place for one trading period before it can resolve, we must distinguish states that share the same shifter value but differ in their transition dynamics. Let the state variable be $s \in \{0,1,2\}$ with shifter values

$$v(0) = 1, \qquad v(1) = 1, \qquad v(2) = 0.$$

States $s = 0$ and $s = 1$ both represent the shifted price level ($v = 1$), whereas $s = 2$ represents the resolved baseline ($v = 0$). The distinction is that $s = 0$ is the initial post-shift state, which transitions deterministically to $s = 1$ in the first trading period. Thereafter, resolution occurs stochastically. The Markov transition probabilities are:

$$q(1|0) = 1, \qquad q(0|0) = 0, \tag{105}$$

$$q(2|1) = \pi_R, \qquad q(1|1) = 1 - \pi_R, \tag{106}$$

$$q(2|2) = 1. \tag{107}$$

All other transition probabilities are zero. As stated in Proposition 3, we consider one risky asset ($J = 1$) and discrete trade at event dates, for the case with large λ. Unless stated otherwise, we use the same notation as in the proof of Propositions 1 and 2. In particular, the diffusive terms (91), the Dynkin expansion (95), and the functions $g(\theta)$ and $b(\theta)$ depend only on the parameters governing wealth dynamics between event times (r_f, μ, σ) and the drifting portfolio weight θ_t. They are independent of the shifter-process structure, so they apply unchanged here.

Returns $\kappa_{s,s'}$. The return from transitioning from state s to state s' is

$$\kappa_{s,s'} = e^{\beta(v(s')-v(s))} - 1. \tag{108}$$

Since $v(0) = v(1) = 1$, the initial transition yields

$$\kappa_{0,1} = e^{\beta(v(1)-v(0))} - 1 = e^0 - 1 = 0 \qquad \text{for all } \beta. \tag{109}$$

That is, the state transition from $s = 0$ to $s = 1$ generates no trading return, regardless of the magnitude of β.

HJB and FOC at state $s = 0$. At state $s = 0$, the only transition with positive probability is to $s = 1$ (with $q(1|0) = 1$). Since $\kappa_{0,1} = 0$, the wealth factor at the jump simplifies to

$$\mathbb{E}\big[(1 + \theta_\Delta \cdot \kappa_{0,1})^{1-\gamma} \mid \theta_0 = \theta_D\big] = 1,$$

and the jump contribution to the HJB reduces to $\lambda\big[C(1) - A(0)\big]$, which is independent of the portfolio choice θ_D. The full HJB at state $s = 0$ is therefore:

$$\begin{aligned} 0 = &-\rho A(0) + (1-\gamma)A(0)\Big[r_f + \theta_D(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\Big] \\ &+ \frac{1-\gamma}{\lambda} A(0)\, g(\theta_D) + \lambda\big[C(1) - A(0)\big] + O(\lambda^{-1}). \end{aligned} \tag{110}$$

Differentiating with respect to θ_D and dividing by $(1-\gamma)A(0)$ yields the FOC:

$$0 = (\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} + O(\lambda^{-1}). \tag{111}$$

The only β-dependence in this FOC is confined to the $O(\lambda^{-1})$ remainder.

Sensitivity to β at $\beta = 0$. When $\beta = 0$, all returns $\kappa_{s,s'} = 0$, so the FOC is state-independent and yields a common optimal weight $\theta^*_{D,0}$ across all states, with $A(0) = A(1) = A(2)$ and $C(0) =$

$C(1) = C(2)$. Define $f(\theta, \beta)$ as the right-hand side of the FOC (111), so that $f(\theta^*_D(0;\beta), \beta) = 0$. Evaluating the partial derivatives at $(\theta^*_{D,0}, 0)$:

$$f_\theta(\theta^*_{D,0}, 0) = -\gamma\sigma^2 + \frac{g''(\theta^*_{D,0})}{\lambda} + O(\lambda^{-1}), \tag{112}$$

$$f_\beta(\theta^*_{D,0}, 0) = O(\lambda^{-1}). \tag{113}$$

Applying the implicit function theorem yields:

$$\left.\frac{d\theta^*_D(0)}{d\beta}\right|_{\beta=0} = -\left.\frac{f_\beta}{f_\theta}\right|_{(\theta^*_{D,0},0)} = -\frac{O(\lambda^{-1})}{-\gamma\sigma^2 + O(\lambda^{-1})} = O(\lambda^{-1}). \tag{114}$$

Shifter-process elasticity at $s = 0$. In the notation of equation (15), the current shifter value $v(0) = 1$ plays the role of s_t. The shifter-process elasticity at state $s = 0$ is therefore:

$$\eta(0) = -\frac{1}{v(0)} \cdot \frac{1}{\theta^*_{D,0}} \cdot \left.\frac{d\theta^*_D(0)}{d\beta}\right|_{\beta=0} = O(\lambda^{-1}). \tag{115}$$

Comparison with the price elasticity. All elasticity calculations evaluate derivatives at $\beta = 0$. At $\beta = 0$, the portfolio problem is independent of the Markov process, so the true own-price elasticity ε_D and the shifter-process elasticity η are both marginal deviations from a common baseline, evaluated for different shifter processes. Therefore, the expression for the true price elasticity ε_D derived in (102) applies here unchanged, giving $\varepsilon_D = O(\lambda)$ as $\lambda \to \infty$. Since $\eta(0) = O(\lambda^{-1})$, we obtain:

$$\frac{\eta(0)}{\varepsilon_D} = \frac{O(\lambda^{-1})}{O(\lambda)} = O(\lambda^{-2}). \tag{116}$$

Recognizing that the estimated elasticity $\hat{\varepsilon}$ recovers the shifter-process elasticity η, we obtain

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = \frac{\eta(0)}{\varepsilon_D} = O(\lambda^{-2}) \to 0 \quad \text{as } \lambda \to \infty. \tag{117}$$

Log utility and continuous trade. Under log utility with continuous trading, the value function takes the form $V(W,s) = \upsilon \log W + A(s)$ and the FOC is given by (46) (see Appendix A.1). Applied to the current setting with one risky asset at state $s = 0$, where the only transition is to $s = 1$ with $q(1|0) = 1$, the FOC becomes:

$$0 = \upsilon\left[(\mu - r_f) - \sigma^2\theta\right] + \lambda\,\widetilde{C}\,\frac{\kappa_{0,1}}{1 + \theta\,\kappa_{0,1}}. \tag{118}$$

Since $\kappa_{0,1} = 0$ exactly for all β, the jump term vanishes, and the FOC reduces to $\theta^* = (\mu - r_f)/\sigma^2$ for all β. Therefore $d\theta^*/d\beta = 0$ and $\hat{\varepsilon} = \eta = 0$ exactly. □

A.4 Proof of Proposition 4

We consider a Markov process for s with three states, $\Omega = \{0, 1, 2\}$, and the Markov transition probabilities:

$$q(2|1) = \pi_B, \qquad q(1|1) = 1 - \pi_B, \tag{119}$$

$$q(0|2) = \pi_R, \qquad q(2|2) = 1 - \pi_R, \tag{120}$$

$$q(0|0) = 1. \tag{121}$$

All other Markov transition probabilities are zero. As stated in Proposition 4, we consider one risky asset ($J = 1$) and discrete trade at event dates, for the case with large λ. Unless stated otherwise, we use the same notation as in the proof of Propositions 1 and 2. In particular, the diffusive terms (91), the Dynkin expansion (95), and the functions $g(\theta)$ and $b(\theta)$ depend only on the parameters governing wealth dynamics between event times (r_f, μ, σ) and the drifting portfolio weight θ_t. They are independent of the shifter-process structure, so they apply unchanged here.

Returns $\kappa_{s,s'}$. The return from transitioning from state s to state s' is $\kappa_{s,s'} = e^{\beta(s'-s)} - 1$. At state $s = 1$, the two transitions with positive probability are to $s' = 2$ (build-up, probability π_B) and $s' = 1$ (stay, probability $1 - \pi_B$). For a marginal exposure $d\beta$, the corresponding returns are

$$\kappa_{1,2} = d\beta + O(d\beta^2), \tag{122}$$

$$\kappa_{1,1} = 0. \tag{123}$$

Importantly, the build-up return $\kappa_{1,2} \approx +d\beta$ is positive for $d\beta > 0$.

HJB and FOC at state $s = 1$. At state $s = 1$, the jump terms of the HJB are

$$\lambda \Big[\pi_B C(2)\, \mathbb{E}\big[(1 + \theta_\Delta \kappa_{1,2})^{1-\gamma} \mid \theta_0 = \theta_D\big] + (1 - \pi_B) C(1) - A(1) \Big]. \tag{124}$$

Expanding $(1+\theta_\Delta \kappa_{1,2})^{1-\gamma} \approx 1+(1-\gamma)\theta_\Delta \kappa_{1,2}$ for small $d\beta$ and using the Dynkin expansion (95) for $\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D]$, we obtain:

$$\begin{aligned}&\mathbb{E}\left[(1+\theta_\Delta \kappa_{1,2})^{1-\gamma} \mid \theta_0 = \theta_D\right] \\ =&1+(1-\gamma)\left(\theta_D + \frac{b(\theta_D)}{\lambda}\right)\kappa_{1,2} + O(d\beta^2) + O(\lambda^{-2} d\beta).\end{aligned} \tag{125}$$

Combining with the diffusive terms (91), the full HJB at state $s=1$ takes the form

$$\begin{aligned}0 = &-\rho A(1) + (1-\gamma)A(1)\left[r_f + \theta_D(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\right] + \frac{1-\gamma}{\lambda}A(1)\, g(\theta_D) \\ &+ \lambda\left[\pi_B C(2)\left(1+(1-\gamma)\left(\theta_D + \frac{b(\theta_D)}{\lambda}\right)\kappa_{1,2}\right) + (1-\pi_B)C(1) - A(1)\right] \\ &+ O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1} d\beta).\end{aligned} \tag{126}$$

Differentiating with respect to θ_D and dividing by $(1-\gamma)A(1)$ yields the FOC:

$$\begin{aligned}0 = &(\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} + \frac{\lambda}{A(1)}\pi_B C(2)\,\kappa_{1,2}\left(1 + \frac{b'(\theta_D)}{\lambda}\right) \\ &+ O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1} d\beta),\end{aligned} \tag{127}$$

with $\kappa_{1,2} = d\beta + O(d\beta^2)$.

State-independence at $\beta = 0$. When $\beta = 0$, all returns $\kappa_{s,s'} = 0$, so the FOC is state-independent and yields a common optimal weight $\theta^*_{D,0}$ across all states, with $A(0) = A(1) = A(2)$ and $C(0) = C(1) = C(2)$. Substituting these common values and $\kappa_{1,2} = d\beta + O(d\beta^2)$ into (127), the FOC at state $s=1$ becomes

$$\begin{aligned}0 = &(\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} + \frac{\lambda}{A(0)}\pi_B C(0)\, d\beta\left(1 + \frac{b'(\theta_D)}{\lambda}\right) \\ &+ O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1} d\beta).\end{aligned} \tag{128}$$

At $\beta = 0$, the substitutions $A(1) = A(0)$ and $C(2) = C(0)$ are exact. In the next step, we differentiate with respect to β and evaluate at $\beta = 0$. Although $A(1;\beta)$ and $C(2;\beta)$ depend on β, their derivatives enter only multiplied by $\kappa_{1,2}\big|_{\beta=0} = 0$, so those terms vanish. We use that fact when computing the partial derivatives below.

Sensitivity to β at $\beta = 0$. Define $f(\theta, \beta)$ as the right-hand side of (128), so that $f(\theta_D^*(1;\beta), \beta) = 0$. Evaluating the partial derivatives at $(\theta_{D,0}^*, 0)$:

$$f_\theta(\theta_{D,0}^*, 0) = -\gamma\sigma^2 + \frac{g''(\theta_{D,0}^*)}{\lambda} + O(\lambda^{-1}), \tag{129}$$

$$f_\beta(\theta_{D,0}^*, 0) = \frac{\lambda}{A(0)}\,\pi_B C(0)\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda}\right) + O(\lambda^{-1}). \tag{130}$$

Applying the implicit function theorem yields:

$$\left.\frac{d\theta_D^*(1)}{d\beta}\right|_{\beta=0} = -\left.\frac{f_\beta}{f_\theta}\right|_{(\theta_{D,0}^*,0)} = \frac{\lambda\,\pi_B C(0)}{A(0)\,\gamma\sigma^2}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right) + O(\lambda^{-1}). \tag{131}$$

Shifter-process elasticity at $s = 1$. The shifter-process elasticity at state $s = 1$ is

$$\begin{aligned}\eta(1) &= -\frac{1}{1}\,\frac{1}{\theta_{D,0}^*}\left.\frac{d\theta_D^*(1)}{d\beta}\right|_{\beta=0} \\ &= -\frac{\lambda\,\pi_B C(0)}{A(0)\,\gamma\sigma^2\,\theta_{D,0}^*}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right) + O(\lambda^{-1}).\end{aligned} \tag{132}$$

Comparison with the price elasticity. As noted previously, in the proof of Proposition 3, when $\beta = 0$, the true price elasticity derived in (102) applies independent of the properties of the Markov process for the state s. Thus, the expression derived in (102) for the price elasticity ε_D applies here unchanged. Dividing $\eta(1)$ by (102), yields

$$\frac{\eta(1)}{\varepsilon_D} = -\pi_B + O(\lambda^{-2}). \tag{133}$$

Recognizing that the estimated elasticity $\hat{\varepsilon}$ recovers the shifter-process elasticity η (by conditions 2 and 3 of Lemma 2), we obtain

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = \frac{\eta(1)}{\varepsilon_D} = -\pi_B + O(\lambda^{-2}) \to -\pi_B \quad \text{as } \lambda \to \infty. \tag{134}$$

□

A.5 Proof of Proposition 5 and Holdings Dynamics in Section 3.3

Starting from the market-clearing condition in changes

$$dz_{j,t} = -\frac{dn_{j,t}}{d\beta_j} d\beta_j \tag{135}$$

we obtain the relation

$$x_j P_{j,t} d\beta_j = P_{j,t} dz_{j,t} \cdot \left(-\frac{dn_{j,t}/n_{j,t}}{d\beta_j} \frac{n_{j,t}}{x_j} \right)^{-1}. \tag{136}$$

Next, we rewrite the shifter-process elasticity of the unit demand as a function of the portfolio share demand

$$\begin{aligned} -\frac{dn_{j,t}/n_{j,t}}{d\beta_j} &= -\frac{d\left(\frac{\theta^*_{j,t} W_t}{P_{j,t}}\right)/d\beta_j}{n_{j,t}} = -\left(\frac{d\theta^*_{j,t}}{d\beta_j} \frac{W_t}{P_{j,t}} + \frac{dW_t}{d\beta_j} \frac{\theta^*_{j,t}}{P_{j,t}} - \frac{s_t W_t \theta^*_{j,t}}{P_{j,t}} \right) \frac{1}{n_{j,t}} \\ &= -\left(\frac{1}{s_t} \frac{d\theta^*_{j,t}/\theta^*_{j,t}}{d\beta_j} n_{j,t} s_t + P_{j,t} n_{j,t} s_t \frac{\theta^*_{j,t}}{P_{j,t}} - s_t n_{j,t} \right) \frac{1}{n_{j,t}} \\ &= [\eta_{j,t} + (1 - \theta^*_{j,t})] \cdot s_t, \end{aligned} \tag{137}$$

where we use the relations:

$$\frac{dW}{d\beta_j} = \frac{d(\sum_{j=0}^{J} n_{j,t} \cdot e^{\beta_j s_t} \tilde{P}_{j,t})}{d\beta_j} = n_{j,t} e^{\beta_j s_t} \tilde{P}_{j,t} s_t = P_{j,t} n_{j,t} s_t, \tag{138}$$

$$\frac{dP^{-1}_{j,t}}{d\beta_j} = \frac{d(e^{\beta_j s_t} \tilde{P}_{j,t})^{-1}}{d\beta_j} = -(e^{\beta_j s_t} \tilde{P}_{j,t})^{-2} s_t e^{\beta_j s_t} \tilde{P}_{j,t} = -\frac{s_t}{P_{j,t}}. \tag{139}$$

Combining (137) with market clearing in changes (135) yields the stated result (27) that to induce a stochastic log price path differential $\{d\beta_j s_\tau\}_{\tau=t}^{\infty}$ for asset j the following process of holdings changes in asset j is required:

$$\left\{ \frac{dz_{j,\tau}}{x_j} \right\}_{\tau=t}^{\infty} = \left\{ d\beta_j s_\tau \cdot [\eta_{j,\tau} + (1 - \theta^*_{j,\tau})] \cdot \frac{n_{j,\tau}}{x_j} \right\}_{\tau=t}^{\infty}. \tag{140}$$

Holdings dynamics for all other assets $i \neq j$. We now derive the intervention requirements for all other assets $i \neq j$ when the price of asset j is supposed to be shifted by $\{d\beta_j s_\tau\}_{\tau=t}^{\infty}$. For any other asset $i \neq j$ (where $i \in \{0, 1, \ldots, J\}$), the market-clearing condition gives:

$$dz_{i,t} = -\frac{dn_{i,t}}{d\beta_j} d\beta_j. \tag{141}$$

Following the same approach as for asset j, we rewrite the unit demand elasticity as a function of the portfolio share demand. Since $i \neq j$, we have $P_{i,t} = e^{\beta_i s_t} \tilde{P}_{i,t}$, which implies $dP_{i,t}/d\beta_j = 0$. Therefore:

$$\begin{aligned}\frac{dn_{i,t}}{d\beta_j} &= \frac{d}{d\beta_j}\left(\frac{\theta^*_{i,t} W_t}{P_{i,t}}\right) = \frac{1}{P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t}\frac{dW_t}{d\beta_j}\right] + \theta^*_{i,t} W_t \cdot \frac{d(P^{-1}_{i,t})}{d\beta_j} \\ &= \frac{1}{P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t}\frac{dW_t}{d\beta_j}\right] + 0. \end{aligned} \tag{142}$$

Dividing by $n_{i,t}$ and using $\frac{dW_t}{d\beta_j} = P_{j,t} n_{j,t} s_t$:

$$\begin{aligned}\frac{dn_{i,t}/n_{i,t}}{d\beta_j} &= \frac{1}{n_{i,t}P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t} P_{j,t} n_{j,t} s_t\right] \\ &= \frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{n_{i,t}P_{i,t}} + \theta^*_{i,t}\frac{P_{j,t} n_{j,t} s_t}{n_{i,t}P_{i,t}}. \end{aligned} \tag{143}$$

Now using $n_{i,t} = \frac{\theta^*_{i,t} W_t}{P_{i,t}}$, the first term becomes:

$$\frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{n_{i,t}P_{i,t}} = \frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{\theta^*_{i,t} W_t} = \frac{1}{\theta^*_{i,t}}\frac{d\theta^*_{i,t}}{d\beta_j} = \frac{1}{s_t}\frac{d\ln\theta^*_{i,t}}{d\beta_j} s_t = -\eta_{ij,t} s_t. \tag{144}$$

For the second term, we use $n_{i,t}P_{i,t} = \theta^*_{i,t} W_t$ and $n_{j,t}P_{j,t} = \theta^*_{j,t} W_t$:

$$\theta^*_{i,t}\frac{P_{j,t} n_{j,t} s_t}{n_{i,t}P_{i,t}} = \theta^*_{i,t}\frac{\theta^*_{j,t} W_t \cdot s_t}{\theta^*_{i,t} W_t} = \theta^*_{j,t} s_t. \tag{145}$$

Consequently:

$$-\frac{dn_{i,t}/n_{i,t}}{d\beta_j} = [\eta_{ij,t} - \theta^*_{j,t}] s_t. \tag{146}$$

Using the market-clearing condition $dz_{i,t} = -\frac{dn_{i,t}}{d\beta_j} d\beta_j$, we obtain:

$$\frac{dz_{i,t}}{x_i} = -\frac{dn_{i,t}}{d\beta_j} d\beta_j \frac{1}{x_i} = [\eta_{ij,t} - \theta^*_{j,t}] s_t \cdot d\beta_j \cdot \frac{n_{i,t}}{x_i}. \tag{147}$$

This gives the final result

$$\left\{\frac{dz_{i,\tau}}{x_i}\right\}^{\infty}_{\tau=t} = \left\{d\beta_j s_\tau \cdot [\eta_{ij,\tau} - \theta^*_{j,\tau}] \cdot \frac{n_{i,\tau}}{x_i}\right\}^{\infty}_{\tau=t}, \tag{148}$$

as stated in (28).

Safe-asset holdings dynamics in the case of one risky asset. We now consider the case of one risk asset ($J = 1$) and examine the required holdings dynamics for the safe asset $i = 0$, applying our general formula (28). For the case $i = 0$ (safe asset), $j = 1$ (risky asset), we have:

$$\frac{dz_{0,t}}{x_0} = d\beta_1 \cdot s_t \cdot [\eta_{01,t} - \theta^*_{1,t}] \cdot \frac{n_{0,t}}{x_0}. \tag{149}$$

We compute the cross-elasticity using $\theta^*_{0,t} = 1 - \theta^*_{1,t}$:

$$\begin{aligned}
\eta_{01,t} &= -\frac{1}{s_t}\frac{d\ln\theta^*_{0,t}}{d\beta_1}\bigg|_{\beta_1=0} = -\frac{1}{s_t}\frac{d\ln(1-\theta^*_{1,t})}{d\beta_1}\bigg|_{\beta_1=0} \\
&= -\frac{1}{s_t}\cdot\frac{1}{1-\theta^*_{1,t}}\cdot\left(-\frac{d\theta^*_{1,t}}{d\beta_1}\right)\bigg|_{\beta_1=0} \\
&= \frac{1}{1-\theta^*_{1,t}}\cdot\left(-\frac{1}{s_t}\frac{d\ln\theta^*_{1,t}}{d\beta_1}\right)\cdot\theta^*_{1,t}\bigg|_{\beta_1=0} \\
&= -\eta_{1,t}\cdot\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}}.
\end{aligned} \tag{150}$$

Substituting this into (149) yields:

$$\begin{aligned}
\frac{dz_{0,t}}{x_0} &= d\beta_1\cdot s_t\cdot\left[-\eta_{1,t}\cdot\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}} - \theta^*_{1,t}\right]\cdot\frac{n_{0,t}}{x_0} \\
&= d\beta_1\cdot s_t\cdot\left[-\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}}\right]\cdot\left[\eta_{1,t} + (1-\theta^*_{1,t})\right]\cdot\frac{n_{0,t}}{x_0} \\
&= -d\beta_1\cdot s_t\cdot[\eta_{1,t} + (1-\theta^*_{1,t})]\cdot\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}}\cdot\frac{n_{0,t}}{x_0}.
\end{aligned} \tag{151}$$

Using the relation $\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}} = \frac{n_{1,t}P_{1,t}}{n_{0,t}P_{0,t}}$ we obtain:

$$\begin{aligned}
\frac{dz_{0,t}}{x_0} &= -d\beta_1\cdot s_t\cdot[\eta_{1,t} + (1-\theta^*_{1,t})]\cdot\frac{n_{1,t}P_{1,t}}{n_{0,t}P_{0,t}}\cdot\frac{n_{0,t}}{x_0} \\
&= -d\beta_1\cdot s_t\cdot[\eta_{1,t} + (1-\theta^*_{1,t})]\cdot\frac{n_{1,t}}{x_0}\cdot\frac{P_{1,t}}{P_{0,t}}.
\end{aligned} \tag{152}$$

Multiplying both sides by x_0:

$$dz_{0,t} = -d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1 - \theta^*_{1,t})] \cdot n_{1,t} \cdot \frac{P_{1,t}}{P_{0,t}}. \tag{153}$$

But we also know from the result (27) for asset $j = 1$ that:

$$dz_{1,t} = d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1 - \theta^*_{1,t})] \cdot n_{1,t}. \tag{154}$$

Therefore:

$$dz_{0,t} = -dz_{1,t} \cdot \frac{P_{1,t}}{P_{0,t}}, \tag{155}$$

or in process form:

$$\{dz_{0,\tau}\}_{\tau=t}^{\infty} = \left\{ -dz_{1,\tau} \cdot \frac{P_{1,\tau}}{P_{0,\tau}} \right\}_{\tau=t}^{\infty}. \tag{156}$$

A.6 Proof of Corollary 2

The assumptions in Corollary 2 are the same as those in Proposition 2, except that we consider continuous trade. Further, equation (27) implies the following relation for the incremental holding upon inception of a dynamic the program:

$$dz_{1,t} = s_t d\beta_1 \cdot [\eta_{1,t} + (1 - \theta^*_{1,t})] \cdot n_{1,t}, \tag{157}$$

where $\eta_{1,t}$ is the shifter-process elasticity corresponding to the postulated price dynamics. For the resolution state process, we derived in the proof of Propositions 1 and 2 the following shifter-process elasticity in states $s \in \{-1, +1\}$:

$$\eta_C(s) = \pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)}. \tag{158}$$

Moreover, given the considered marginal exposure we can write:

$$\theta^*_C(s) = \theta^*_C(0) + \theta_1 s_t d\beta_1, \tag{159}$$

for some constant θ_1. Plugging these results into (157) yields

$$dz_{1,t} = s_t d\beta_1 \cdot [\pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)} + (1 - (\theta^*(0) + \theta_1 s_t d\beta_1))] \cdot n_{1,t}, \tag{160}$$

$$= s_t d\beta_1 \cdot [\pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)} + (1 - \theta^*(0))] \cdot n_{1,t}. \tag{161}$$

As a result, the initial unit change in the risky asset declines linearly in the per-event resolution probability π_R.

A.7 Measuring Outsider Flows

Consider a single jump from the pre-state s to the post-state s', keeping the notation introduced in Section 3. Let $\theta_j^- = \theta_j^*(s)$ and $\theta_j^+ = \theta_j^*(s')$ be the optimal shares before and after the jump, and denote the pre-jump price by P_j^-. The price and wealth immediately after the jump are given by:

$$P_j^+ = P_j^- (1 + \kappa), \qquad \kappa \equiv e^{\beta_j (s' - s)} - 1,$$

$$W^+ = W^- \left[1 + \theta_j^- \kappa \right]$$

The unit positions of insiders, n_j, and outsiders, z_j, are given by:

$$n_j^- = \frac{\theta_j^- W^-}{P_j^-},$$

$$n_j^+ = \frac{\theta_j^+ W^+}{P_j^+} = \frac{\theta_j^+ W^-}{P_j^-} \frac{1 + \theta_j^- \kappa}{1 + \kappa}.$$

Because the outsiders clear the market in units, $\Delta z_j = -(n_j^+ - n_j^-)$, implying

$$\Delta z_j = \frac{W^-}{P_j^-} \left[\theta_j^- - \theta_j^+ \frac{1 + \theta_j^- \kappa}{1 + \kappa} \right].$$

B Data and Measurement

B.1 Predictive Regressions

The main dataset used in this analysis is the dataset "IV.dta" from the replication files of KY19. This dataset consists of a quarterly (FDATE) panel dataset on stocks (PERMNO). The sample period is 1980Q1-2017Q4. For each PERMNO/FDATE, the dataset provides information on the characteristics of that stock as well as the value of the instrument for that stock's market cap (*IVme*). The construction of the latter variable is explained in a later section. All characteristics represent signals computed at quarter-end. Accounting variables are lagged by 6 months compared to market variables. Profitability (*profit*) is the ratio of operating profits to book equity (operating profits are defined as revenue REVTS minus cost of goods sold COGS, SG&A expenses XSGA and interest expenses XINT) to book equity. Investment (*Gat*) is the annual log growth rate of book assets. The numerator in the dividend-to-book ratio (*divA_be*) represents the trailing-12-month dividends per current split-adjusted share multiplied by shares outstanding. Market betas (*beta*) are computed with 60-month rolling windows. This dataset has 615,789 observations.

To this baseline dataset, information on monthly stock returns are added as follows. For each PERMNO/FDATE that appears in IV.dta, three columns *retm1*, *retm2* and *retm3* are added and filled with the three monthly returns for the corresponding stock in the corresponding quarter. For instance, for PERMNO 10107 (MSFT) and FDATE 2000Q1, those three columns take on values -0.161670, -0.086845 and 0.188811 respectively. Monthly stock returns from CRSP are retrieved from common stocks (codes 10 and 11) and adjusted for potential delisting. NYSE size breakpoints at the end of each quarter are retrieved from Kenneth French's website and added to the dataset.

For a given PERMNO, quarterly observations may not be consecutive. This may be due to missing data on stock characteristics or returns for some quarters or months. While not frequent, this phenomenon affects 4,219 rows that represent time gaps of 2 quarters or more, with most occurrences pertaining to microcap stocks. KY define the universe of stocks as common stocks with nonmissing data on characteristics and return. To avoid incomplete time series, the longest sequence of quarterly observations for each PERMNO is retained. The resulting dataset has 574,541 observations.

For each PERMNO/FDATE, LNme is the log of the market capitalization of the stock, BtM is its book-to-market ratio. Momentum returns (as of the end of the quarter, i.e its last month) are computed as the compounded return of that stock over the past 3 quarters and the first two months of the considered quarter.

B.2 Elasticity Estimates

A panel of stock characteristics at each quarter-end is constructed as in KY19, using both CRSP-Compustat Quarterly and CRSP-Compustat Quarterly files. Characteristics are constructed for observations whose share code is 10, 11, 12 or 18, or whose exchange code is 1, 2 or 3. Observations with a missing stock price (PRC) or a missing value of shares outstanding (SHROUT) are discarded. Profitability, investment, and market beta are winsorized at the 2.5th and 97.5th percentiles for each date, while dividends to book equity are winsorized at the 97.5th percentile.

Data on reported 13F holdings are from Thomson Reuters' S34 dataset and start in 1980Q1. In this dataset, one row identifies a number of shares (SHARES) of a given security (CUSIP) held by a given institution (MGRNO) at a given date (FDATE, in quarters). Information on 13F filers are provided by the file "Manager.dta" retrieved from KY19's replication file. Institutions in the S34 dataset that can't be matched to an institution in "Manager.dta" or for whom RDATE and FDATE do not coincide are discarded. S34 identifies positions with stock CUSIPs. It is possible to use stock CUSIPs to add information on stock characteristics. We use a right merge with the panel of stock characteristics, so that stocks held by no institutions at a given quarter are still retained and attributed entirely to the HH sector.

The "household sector" (MGRNO = 0, "HH") is then defined following KY19 as holding residual shares outstanding (if any) for all securities in the universe. If a position in S34 has a value for SHARES above SHROUT, SHARES is set equal to SHROUT. If the sum of shares of a given security held by 13F filers at a given fdate exceeds SHROUT, all positions are scaled down by the corresponding ratio.

For any given MGRNO/FDATE, rows with zero values for SHARES are added for PERMNOs that have been held during any of the past 11 quarters by this MGRNO (but are not held during the quarter in consideration). This set of PERMNOs (held currently or in the past 11 quarters) is called the MGRNO's investment universe at that date. The number of stocks held in nonzero amount is labeled NHOLDING.

Following KY19, all positions of a given MGRNO at a given FDATE in a stock with share code 12, 18 or with missing return or characteristics are aggregated and labelled as the "outside asset" (PERMNO = 0). Dollar holdings of individual securities (HOLDING) as well as dollar holdings of the outside asset (OUTAUM) and total dollar holdings (AUM) are computed for each MGRNO/FDATE. If a given MGRNO has assets under management (AUM) below 10M at any given time or if it has no holding in the outside asset (or only holds the outside asset), its holdings are rebated to the HH sector. Once these modifications are made, AUMs and dollar holdings of individual securities are recomputed.

After those steps, portfolio weights (WEIGHT) and relative portfolio weights (RWEIGHT) are computed for each position. The logarithm of relative weights (LNRWEIGHT) is also defined for

rows where $RWEIGHT > 0$. For each MGRNO/FDATE, the average of LNRWEIGHT is stored in a variable CONS.

Any position held by a MGRNO at a given FDATE is labelled persistent ($IHOLDING = 100$) if and only if the stock has been held by the MGRNO at least once in the past 11 quarters. Otherwise, $IHOLDING$ receives a value of 0. The variable MHOLDING records the average value of IHOLDING for a given MGRNO/FDATE.

Finally, it is possible to construct $IVme_{it}(n)$ for a stock n and an institution i at time t as in KY19, formula (19). For any stock n at time t, the institutions over which the sum is taken are only MGRNOs that have $MHOLDING \geq 95$ at time t (the household sector is also excluded).

C Additional Empirical Results

Table 6: Summary statistics: stationary variables. The table reports time-series averages of within-quarter cross-sectional descriptive statistics for the stationary stock-level variables used in our analysis: profitability (PROF), investment (INV), dividend-to-book ratio (DIV), market beta (BETA), book-to-market ratio (BtM), and momentum (MOM). For each quarter, we compute the cross-sectional mean, standard deviation, and 25th, 50th, and 75th percentiles of each variable; cells report time-series averages of these moments. Sample period: 1980Q1–2017Q4, excluding observations with size below the NYSE 10th percentile breakpoint. Variable definitions are given in Section 4.

Variable	Mean	Std	p25	p50	p75
PROF	0.21	0.24	0.13	0.22	0.32
INV	0.12	0.21	0.01	0.08	0.19
DIV	0.02	0.03	0.00	0.01	0.04
BETA	1.14	0.62	0.70	1.05	1.49
BtM	0.78	0.92	0.35	0.59	0.95
MOM	0.21	0.52	−0.06	0.13	0.36

Table 7: Summary statistics: non-stationary variables. The table reports time-series averages of within-quarter cross-sectional descriptive statistics for three log market-equity variables, computed separately within seven successive five-year sub-periods. Panel A reports statistics for KY19's instrumented log market equity, $\widehat{me}$. Panel B reports statistics for actual log market equity, *me*. Panel C reports statistics for the instrumented price-shift component, $\hat{b} \times \widehat{me}$, defined in Section 4 as the projection of *me* onto $\widehat{me}$ controlling for the asset-pricing characteristics in **x**. Within each sub-period and quarter, we compute the cross-sectional mean, standard deviation, and 25th, 50th, and 75th percentiles; cells report time-series averages of these moments within the sub-period. Sample period: 1980Q1–2017Q4, excluding observations with size below the NYSE 10th percentile breakpoint.

Panel A: $\widehat{me}$							
Statistic	1980-85	1985-90	1990-95	1995-00	2000-05	2005-10	2010-17
Average CS Mean	4.07	5.08	5.75	6.47	7.27	7.83	8.36
Average CS P25	3.24	4.48	5.17	6.13	7.02	7.61	8.12
Average CS P50	4.15	5.09	5.75	6.50	7.30	7.79	8.31
Average CS P75	5.09	5.78	6.34	6.87	7.56	8.04	8.62
Average CS Std	1.32	0.91	0.78	0.63	0.48	0.38	0.48
Panel B: *me*							
Statistic	1980-85	1985-90	1990-95	1995-00	2000-05	2005-10	2010-17
Average CS Mean	5.50	5.93	6.02	6.52	6.90	7.44	7.74
Average CS P25	4.48	4.80	4.81	5.35	5.75	6.35	6.61
Average CS P50	5.28	5.70	5.76	6.21	6.60	7.15	7.51
Average CS P75	6.33	6.84	7.00	7.40	7.73	8.22	8.59
Average CS Std	1.24	1.36	1.48	1.46	1.47	1.38	1.43
Panel C: $\hat{b} \times \widehat{me}$							
Statistic	1980-85	1985-90	1990-95	1995-00	2000-05	2005-10	2010-17
Average CS Mean	3.23	5.39	7.33	8.05	10.48	13.50	14.43
Average CS P25	2.57	4.75	6.60	7.63	10.12	13.11	14.01
Average CS P50	3.29	5.39	7.33	8.09	10.52	13.42	14.34
Average CS P75	4.03	6.12	8.08	8.54	10.89	13.85	14.87
Average CS Std	1.04	0.95	0.99	0.78	0.70	0.66	0.82

Table 8: Return predictability regressions with z-score standardized regressors: lower trading frequencies. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are standardized in each period (z-scored). Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one quarter, two quarters, and four quarters.

(a) One Quarter

Regressor	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	-0.0033	—	-0.0013
	(0.0037)		(0.0039)
PROF	0.0087**	0.0092***	0.0092***
	(0.0034)	(0.0035)	(0.0035)
INV	-0.0107***	-0.0101***	-0.0101***
	(0.0019)	(0.0019)	(0.0019)
DIV	-0.0018	-0.0012	-0.0011
	(0.0017)	(0.0015)	(0.0017)
BETA	0.0007	0.0001	0.0002
	(0.0039)	(0.0039)	(0.0039)
$me - \hat{b} \times \widehat{me}$	—	-0.0060**	-0.0058*
		(0.0029)	(0.0030)

(b) Two Quarters

(a)	(b)	(c)
-0.0049	—	-0.0014
(0.0074)		(0.0076)
0.0172***	0.0180***	0.0181***
(0.0055)	(0.0058)	(0.0057)
-0.0209***	-0.0199***	-0.0199***
(0.0028)	(0.0028)	(0.0028)
-0.0032	-0.0022	-0.0021
(0.0027)	(0.0027)	(0.0028)
-0.0001	-0.0012	-0.0011
(0.0064)	(0.0064)	(0.0064)
—	-0.0105**	-0.0102**
	(0.0047)	(0.0048)

(c) Four Quarters

(a)	(b)	(c)
-0.0058	—	0.0019
(0.0146)		(0.0144)
0.0345***	0.0365***	0.0364***
(0.0086)	(0.0087)	(0.0087)
-0.0359***	-0.0339***	-0.0338***
(0.0047)	(0.0048)	(0.0045)
-0.0059	-0.0032	-0.0034
(0.0057)	(0.0055)	(0.0058)
-0.0016	-0.0036	-0.0037
(0.0110)	(0.0107)	(0.0108)
—	-0.0225**	-0.0228***
	(0.0089)	(0.0084)

Table 9: Return predictability regressions with rank-standardized regressors: lower trading frequencies. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are rank-standardized in each period. Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one quarter, two quarters, and four quarters.

	(a) One Quarter			**(b)** Two Quarters			**(c)** Four Quarters		
Regressor	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	-0.0038	—	0.0000	-0.0061	—	0.0007	-0.0057	—	0.0101
	(0.0055)		(0.0058)	(0.0105)		(0.0110)	(0.0207)		(0.0205)
PROF	0.0119***	0.0130***	0.0130***	0.0226***	0.0245***	0.0245***	0.0419***	0.0466***	0.0463***
	(0.0038)	(0.0040)	(0.0040)	(0.0064)	(0.0067)	(0.0067)	(0.0103)	(0.0103)	(0.0104)
INV	-0.0154***	-0.0142***	-0.0142***	-0.0303***	-0.0282***	-0.0281***	-0.0479***	-0.0434***	-0.0427***
	(0.0029)	(0.0028)	(0.0029)	(0.0045)	(0.0045)	(0.0044)	(0.0080)	(0.0084)	(0.0078)
DIV	-0.0026	-0.0012	-0.0012	-0.0030	-0.0004	-0.0005	-0.0023	0.0051	0.0036
	(0.0047)	(0.0045)	(0.0047)	(0.0073)	(0.0075)	(0.0076)	(0.0144)	(0.0144)	(0.0149)
BETA	0.0012	-0.0001	-0.0001	0.0005	-0.0017	-0.0018	-0.0009	-0.0049	-0.0061
	(0.0057)	(0.0057)	(0.0056)	(0.0098)	(0.0097)	(0.0096)	(0.0168)	(0.0160)	(0.0162)
$me - \hat{b} \times \widehat{me}$	—	-0.0115**	-0.0115**	—	-0.0204**	-0.0206**	—	-0.0464***	-0.0490***
		(0.0053)	(0.0056)		(0.0091)	(0.0095)		(0.0170)	(0.0163)

Table 10: Price shift resolution regressions: lower trading frequencies. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, beta, and size) as well as $\hat{b} \times \widehat{me}$. Variables are not standardized. Coefficient estimates and standard errors (in parentheses) are presented. The sample excludes micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). Across all specifications in the table, we can reject the null hypothesis that the coefficient on $(\hat{b} \times \widehat{me})$ is -1, with a p-value of 0.000. The panels represent the following horizons: one quarter, two quarters, and four quarters.

(a) One Quarter

Regressor	(a)	(b)	(c)
$\hat{b} \times \widehat{me}$	-0.0007	—	-0.0047**
	(0.0019)		(0.0020)
PROF	0.0230**	0.0228**	0.0266**
	(0.0112)	(0.0110)	(0.0111)
INV	-0.0556***	-0.0542***	-0.0542***
	(0.0145)	(0.0145)	(0.0145)
DIV	-0.0875	-0.0823	-0.0249
	(0.0775)	(0.0799)	(0.0848)
BETA	0.0015	0.0013	0.0013
	(0.0053)	(0.0052)	(0.0052)
$me - \hat{b} \times \widehat{me}$	—	-0.0000	-0.0053*
		(0.0024)	(0.0029)

(b) Two Quarters

(a)	(b)	(c)
-0.0017	—	-0.0088**
(0.0037)		(0.0038)
0.0484**	0.0476**	0.0544***
(0.0189)	(0.0186)	(0.0192)
-0.1092***	-0.1067***	-0.1068***
(0.0165)	(0.0166)	(0.0164)
-0.1225	-0.1168	-0.0103
(0.1232)	(0.1241)	(0.1432)
0.0016	0.0012	0.0010
(0.0094)	(0.0094)	(0.0093)
—	0.0004	-0.0095*
	(0.0047)	(0.0049)

(c) Four Quarters

(a)	(b)	(c)
-0.0041	—	-0.0170**
(0.0069)		(0.0075)
0.1001***	0.0978***	0.1111***
(0.0323)	(0.0311)	(0.0332)
-0.1890***	-0.1843***	-0.1851***
(0.0252)	(0.0252)	(0.0250)
-0.1457	-0.1424	0.0583
(0.2648)	(0.2514)	(0.2872)
0.0065	0.0058	0.0054
(0.0176)	(0.0175)	(0.0173)
—	0.0018	-0.0173**
	(0.0083)	(0.0077)

Table 11: Return predictability regressions with z-score standardized regressors: micro-cap stocks. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are standardized in each period (z-scored). Coefficient estimates and standard errors (in parentheses) are presented. The sample consists of micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

	(a) One Day			**(b)** One Week			**(c)** One Month		
Regressor	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	0.0004	—	-0.0041***	-0.0003	—	-0.0096***	-0.0024*	—	-0.0210***
	(0.0003)		(0.0009)	(0.0008)		(0.0018)	(0.0014)		(0.0031)
PROF	-0.0011***	-0.0009***	-0.0005*	-0.0022**	-0.0019**	-0.0011	-0.0014	-0.0011	0.0006
	(0.0004)	(0.0003)	(0.0003)	(0.0009)	(0.0008)	(0.0008)	(0.0018)	(0.0017)	(0.0016)
INV	-0.0007***	-0.0002	0.0002	-0.0018***	-0.0008*	-0.0001	-0.0071***	-0.0052***	-0.0036***
	(0.0002)	(0.0002)	(0.0002)	(0.0005)	(0.0005)	(0.0005)	(0.0011)	(0.0011)	(0.0011)
DIV	-0.0001	0.0003	0.0008***	-0.0008	-0.0002	0.0010	-0.0005	0.0004	0.0030*
	(0.0003)	(0.0003)	(0.0003)	(0.0006)	(0.0006)	(0.0006)	(0.0016)	(0.0016)	(0.0015)
BETA	0.0004	0.0000	0.0005	0.0018*	0.0010	0.0020*	0.0032	0.0013	0.0037
	(0.0003)	(0.0004)	(0.0003)	(0.0011)	(0.0011)	(0.0011)	(0.0023)	(0.0023)	(0.0023)
$me - \hat{b} \times \widehat{me}$	—	-0.0031***	-0.0058***	—	-0.0055***	-0.0118***	—	-0.0096***	-0.0236***
		(0.0005)	(0.0010)		(0.0010)	(0.0019)		(0.0016)	(0.0031)

Table 12: Return predictability regressions with rank-standardized regressors: micro-cap stocks. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, and beta) as well as $\widehat{b} \times \widehat{me}$ and the component of firm value (size) unrelated to $\widehat{me}$, that is, $me - \hat{b} \times \widehat{me}$. The regressors are rank-standardized in each period. Coefficient estimates and standard errors (in parentheses) are presented. The sample consists of micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

	(a) One Day			**(b)** One Week			**(c)** One Month		
Regressor	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
$\widehat{b} \times \widehat{me}$	0.0020*	—	-0.0033**	0.0007	—	-0.0090***	-0.0065	—	-0.0281***
	(0.0011)		(0.0016)	(0.0024)		(0.0032)	(0.0049)		(0.0065)
PROF	-0.0019***	-0.0015***	-0.0013**	-0.0031**	-0.0025*	-0.0021	-0.0009	-0.0001	0.0012
	(0.0006)	(0.0006)	(0.0005)	(0.0014)	(0.0014)	(0.0013)	(0.0032)	(0.0031)	(0.0030)
INV	-0.0013***	-0.0006**	-0.0005*	-0.0032***	-0.0021***	-0.0019***	-0.0124***	-0.0101***	-0.0093***
	(0.0003)	(0.0003)	(0.0003)	(0.0007)	(0.0007)	(0.0007)	(0.0017)	(0.0017)	(0.0017)
DIV	0.0001	0.0006	0.0009	-0.0015	-0.0008	-0.0001	-0.0009	-0.0001	0.0021
	(0.0006)	(0.0006)	(0.0006)	(0.0013)	(0.0013)	(0.0013)	(0.0031)	(0.0032)	(0.0031)
BETA	0.0008	0.0004	0.0007	0.0032*	0.0022	0.0031*	0.0059	0.0030	0.0059
	(0.0006)	(0.0006)	(0.0006)	(0.0018)	(0.0019)	(0.0018)	(0.0037)	(0.0040)	(0.0037)
$me - \hat{b} \times \widehat{me}$	—	-0.0044***	-0.0056***	—	-0.0072***	-0.0104***	—	-0.0127***	-0.0229***
		(0.0007)	(0.0010)		(0.0014)	(0.0019)		(0.0024)	(0.0034)

Table 13: Price shift resolution regressions: micro-cap stocks. The table reports panel regressions of excess returns (in decimal form) on standard asset pricing characteristics (profitability, investment, dividends, beta, and size) as well as $\hat{b} \times \widehat{me}$. Variables are not standardized. Coefficient estimates and standard errors (in parentheses) are presented. The sample consists of micro-cap stocks (those below the 10th percentile of NYSE market cap breakpoints). Across all specifications in the table, we can reject the null hypothesis that the coefficient on $(\hat{b} \times \widehat{me})$ is -1, with a p-value of 0.000. The panels represent the following horizons: one day (return over the first day of quarter $t+1$), one week, and one month.

(a) One Day

Regressor	(a)	(b)	(c)
$\hat{b} \times \widehat{me}$	0.0000	—	-0.0038***
	(0.0002)		(0.0008)
PROF	-0.0032***	-0.0030***	-0.0013
	(0.0011)	(0.0010)	(0.0009)
INV	-0.0031***	-0.0026**	0.0004
	(0.0011)	(0.0011)	(0.0014)
DIV	0.0012	-0.0022	0.0337***
	(0.0107)	(0.0114)	(0.0101)
BETA	0.0004	0.0002	0.0010*
	(0.0005)	(0.0005)	(0.0005)
$me - \hat{b} \times \widehat{me}$	—	-0.0004	-0.0047***
		(0.0003)	(0.0010)

(b) One Week

(a)	(b)	(c)
-0.0005	—	-0.0084***
(0.0006)		(0.0021)
-0.0066**	-0.0064**	-0.0028
(0.0026)	(0.0025)	(0.0021)
-0.0087**	-0.0080**	-0.0016
(0.0035)	(0.0037)	(0.0045)
-0.0322	-0.0443	0.0344
(0.0265)	(0.0278)	(0.0271)
0.0019	0.0014	0.0031*
(0.0015)	(0.0015)	(0.0016)
—	-0.0001	-0.0097***
	(0.0008)	(0.0028)

(c) One Month

(a)	(b)	(c)
0.0000	—	-0.0203***
(0.0013)		(0.0043)
-0.0049	-0.0039	0.0038
(0.0054)	(0.0053)	(0.0045)
-0.0311***	-0.0282***	-0.0127
(0.0081)	(0.0085)	(0.0100)
-0.0629	-0.0798	0.1101*
(0.0640)	(0.0673)	(0.0585)
0.0023	0.0015	0.0054*
(0.0030)	(0.0030)	(0.0033)
—	-0.0020	-0.0253***
	(0.0017)	(0.0056)

C.1 Unweighted Beta and Elasticity Statistics by Institution Type

Table 14 reports the unweighted analogue of Table 5. To gauge the price elasticity of the market overall, the value-weighted statistics in the main text are the relevant ones, since they aggregate institutions and holdings by their economic magnitude. The unweighted statistics treat all institutions, bins, and stock positions symmetrically, and are therefore dominated by the many small institutions and small holdings in the sample, for which estimated betas are typically below one and elasticities positive. The unweighted fraction of holdings with $\hat{\varepsilon}_{i,t}(n) < 0$ ranges from only 9% (investment advisors) to 33% (pension funds), smaller than the value shares in Table 5.

Table 14: Estimated Betas and Demand Elasticities — Unweighted. The table reports the unweighted analogue of Table 5. Panel A reports the unweighted distribution of institution-level $\hat{\beta}_{0,i,t}$, treating each institution (or bin) symmetrically. Panel B reports the unweighted distribution of position-level $\hat{\varepsilon}_{i,t}(n)$, treating each stock position symmetrically. Within-quarter distribution moments are computed within each institution type; cells report cross-quarter averages of these moments. Sample: all 13F institutions, 1980Q1–2017Q4 (31 evenly-spaced quarters), unconstrained KY19 estimation.

Statistic	Banks	Insurance Companies	Investment Advisors	Mutual Funds	Pension Funds
Panel A: Estimated $\hat{\beta}_{0,i,t}$					
Mean	−0.15	0.18	0.20	0.48	0.58
Std	0.64	0.52	0.33	0.43	0.39
p_{25}	−0.58	−0.16	−0.02	0.16	0.32
p_{50}	−0.21	0.11	0.19	0.42	0.45
p_{75}	0.24	0.46	0.37	0.73	0.80
Fraction with $\hat{\beta}_{0,i,t} > 1$	0.06	0.09	0.02	0.13	0.16
Mean $\hat{\beta}_{0,i,t}$ if $\hat{\beta}_{0,i,t} > 1$	1.38	1.34	1.37	1.32	1.33
Panel B: Estimated Elasticities $\hat{\varepsilon}_{i,t}(n)$					
Mean	0.60	0.41	0.61	0.24	0.22
Std	0.77	0.60	0.45	0.50	0.42
p_{25}	0.05	-0.03	0.38	-0.09	-0.06
p_{50}	0.63	0.43	0.69	0.25	0.22
p_{75}	1.14	0.83	0.91	0.63	0.56
Fraction with $\hat{\varepsilon}_{i,t}(n) < 0$	0.24	0.28	0.09	0.31	0.33