How (Not) to Identify Demand Elasticities in Dynamic Asset Markets*

Jules H. van Binsbergen[†] Benjamin David[‡] Christian C. Opp[§]

February 25, 2026

Abstract

We evaluate approaches to estimating demand elasticities in dynamic asset markets, both theoretically and empirically. We establish strict, necessary conditions that the dynamics of instrumented asset price variation must satisfy for valid identification. We illustrate these insights in a general equilibrium model of dynamic trade and derive the magnitude of biases that arise when these conditions are violated. Estimates based on static IO models are severely biased when the instrumented price variation is persistent or predictable. Empirically, we show that commonly used instruments yield elasticity estimates that are off by orders of magnitude, or even have the wrong sign. In contrast to standard multiplier calculations, our theory characterizes the dynamic asset market interventions required to sustain a given price path support process, with direct implications for policies such as Quantitative Easing (QE).

*The authors thank Willie Fuchs (discussant) and seminar participants at HEC Montréal, University of North Carolina, and University of Rochester, as well as participants at the Financial Research Association Conference, the Finance Theory Group Meeting at UT Austin, the GEA Annual Conference, and the Macro-Finance Tepper-LAEF Conference for helpful comments.

[†]University of Pennsylvania and NBER. Email: julesv@wharton.upenn.edu.
[‡]University of Pennsylvania. Email: bdavid97@wharton.upenn.edu.
[§]University of Rochester and NBER. Email: 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, 2021). 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 when agents trade assets dynamically. We show that the instruments used in the existing literature to estimate demand elasticities generate large and systematic biases. Building on these insights, we characterize the relation between quantities and equilibrium prices in a generic general equilibrium setting with dynamic trade. This analysis characterizes the dynamic asset market interventions required to implement a targeted price-path support process, revealing the shortcomings of the standard multiplier calculations used in the literature.

The biases in price elasticity estimates of 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 resell securities in the future, so the entire path 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, any candidate instrument that shifts the contemporaneous price at time t while also affecting future resale values violates the exclusion restriction. 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 an asset. These concerns are not merely theoretical: in practice, financial institutions actively engage in dynamic trading strategies. Yet such trading behavior is absent from static IO models.

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 ex ante. 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 simple formulas that quantify the biases in elasticity estimates when these conditions are violated. In particular, we obtain such formulas for three stylized forms of price 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). In the first case, we show that the bias in the estimated elasticity is proportional to the probability that a price shift resolves within an agent's portfolio adjustment interval. For example, if institutions can rebalance daily and a price shift resolves with a constant per-period probability corresponding to an expected duration of two years, the estimated elasticity has a magnitude of only 1/730 of the true elasticity. Second, if resolution does not occur at all in the first trading period, the estimated elasticity is approximately zero, even if the true elasticity is infinite. Finally, in the momentum case, the bias is even more severe: the estimated elasticity takes the opposite sign of the true elasticity. The intuition is straightforward. When variation inducing a positive price shift also predicts higher future returns, agents rationally tilt *toward* the

[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).

security, creating the false appearance of a negative price elasticity.

While we provide bias corrections under stylized dynamics for price path shifts, such as constant per-period resolution or build-up probabilities, we emphasize that no general adjustment factor can correct the distortions inherent in static cross-sectional identification and event-study time-series approaches.

A simple example illustrates this central point. 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? The answer is: 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-based 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? Our answer is: not at all. The crux is that the shock upon impact alters the asset's intrinsic characteristics, specifically its future resale values; 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

[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.

Reserve must purchase. In this case, both static, cross-sectional identification and event-study time-series identification fail.

Note also that the conditions for identification we characterize cannot be relaxed simply by the econometrician choosing a lower data sampling frequency, as this frequency is distinct from the one at which agents can adjust their portfolios. Contemporaneous variation in portfolio weights reflects return opportunities over the next trading period, not over the interval until the econometrician next observes or samples holdings. In the above example, investors' portfolio shares are constant throughout the whole year and only briefly shift for one week before reverting back to their original level.[3]

Our empirical results are closely connected to these issues. In particular, we find that cross-sectional variation in Koijen and Yogo's (2019) (KY19 hereafter) instrumented price variable, $\widehat{me}$ (instrumented market equity), does not satisfy the conditions required for valid price elasticity identification in dynamic trading environments. In fact, cross-sectionally, higher values of $\widehat{me}$ are either unrelated, or even positively related, to average returns over the following month. Rational investors, therefore, should not tilt away from high-$\widehat{me}$ stocks but, if anything, overweight them. According to our model, using such variation as an instrument mechanically yields elasticity estimates that are close to zero or even negative, the opposite sign of the true price elasticity.

Consistent with this prediction, 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] These results underscore that even when true demand elasticities are very large, using instruments that induce persistent or momentum-type price dynamics mechanically yields estimated price elasticities close to zero, or even negative. The estimates are orders of magnitude lower than the theoretical benchmarks to which they are compared,[5] because the instrument violates the exclusion

[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. Moreover, whereas a regular good can plausibly maintain its intrinsic characteristics in the face of persistent price changes, this is infeasible for financial assets that agents can trade dynamically.

[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.

[5]Koijen and Yogo (2021) and Gabaix and Koijen (2021), for example, compare their estimates

restriction required for price elasticity estimation. Technically, these estimates measure how demand responds to an entire 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 therefore 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 regarding the effectiveness of dynamic asset purchase programs and other supply-affecting policies. First, we show that the persistence of a policy's intended price path support has a dramatic impact on the size of the *initial* supply intervention required to achieve it. For instance, under a constant per-period resolution probability, increasing the expected duration of a 1% price shift from one week to one year reduces the required residual supply shock at inception by 98%. This reflects the fact that rational, optimizing traders lean far more aggressively against a price shift when it is expected to reverse quickly. The results also differ markedly depending on whether policy commitments have a fixed expiration date or resolve randomly.

Second, we characterize the path of dynamic asset market interventions required to sustain a given price path support process. This analysis reveals significant limitations of the typical multiplier calculations considered in the literature. Even in the case of 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

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.

response. In particular, purchase programs that solely promise future interventions naturally generate infinite announcement multipliers, mirroring the inverse of an erroneously inferred zero price elasticity. Overall, our findings have useful implications for the literature on quantitative easing and the dynamics of policy interventions (see, e.g., Haddad et al., 2025a).

Our insights speak to how elasticity estimates are commonly interpreted in the existing literature. Davis et al. (2025) argue that limited price pass-through and the presence of unspanned factors imply that price elasticities should be low, on the order of 5. We highlight that the object examined by Davis et al. (2025) is a shifter-process elasticity, not a price elasticity. A price elasticity is the partial derivative of demand at time t with respect to the contemporaneous price and reflects structural objects, such as frictions that inhibit trade. We show that, absent such frictions, price elasticities are in fact infinite. In contrast, a shifter-process elasticity reflects the demand response to a stochastic process that shifts both the contemporaneous price and future prices. We show that this latter elasticity can take any value, including negative values as in the case of the KY19 instrument, and is entirely specific to the price-shifter process at hand. Finally, our dynamic general equilibrium analysis characterizes how quantities and equilibrium prices are related in dynamic settings and shows which dynamic asset market interventions are required to implement a targeted price-path support process.

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 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 state that rational models 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, multipliers used to measure price impact at the inception of multi-period purchases are insufficient and generally misleading in characterizing the relationship between quantities and prices. This is because initial price movements are driven primarily by expectations about the future path of holdings rather than by the initial purchases themselves. In contrast, we characterize the entire path of incremental holdings required to sustain a given path of price shifts. Second, we show that permanent interventions can plausibly exhibit market multipliers of around 2 or higher without requiring additional frictions. Third, empirical estimates of multipliers should be expected to take high values (up to infinity) 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 exhibiting very high price elasticities of demand.

Finally, our key results pertaining to dynamics are distinct from the channels emphasized in the important literature on the implications and validity of different substitution structures assumed in static models (Fuchs et al., 2025a,b; Haddad et al., 2025b; Koijen and Yogo, 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 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 explicitly maximizing CRRA utility and an "outside group" introducing exogenous residual supply shocks. The latter group can be interpreted as noise traders, central banks, or corporations engaging in share issuances or repurchases. We start by detailing CRRA investors' trading and consumption opportunities conditional on postulated equilibrium asset price dynamics. Thereafter, we 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 dynamics of shocks that can identify demand elasticities.

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 over time (i.e., 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 corresponding to this 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

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

Then $Var\big[\mathrm{d}W_t/W_t\big] = \boldsymbol{\theta}^\top \boldsymbol{\Sigma} \boldsymbol{\theta} dt$. 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

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

where W' denotes post-jump wealth:

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

We conjecture and verify that the value function takes the form

$$V(W,s) = u(W) \cdot A(s). \quad (9)$$

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 $s \in \Omega$:

$$0 = -\rho A(s) + (1-\gamma) A(s) \Big[r_f + (\boldsymbol{\mu} - r_f \mathbf{1})^\top \boldsymbol{\theta} - \tfrac{1}{2} \gamma \boldsymbol{\theta}^\top \boldsymbol{\Sigma} \boldsymbol{\theta} \Big] \\ + \lambda \Big[\sum_{s'} q(s'|s) \big[\pi_C + (1-\pi_C) A(s') \big] \cdot \big(1 + \boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'} \big)^{1-\gamma} - A(s) \Big]. \quad (10)$$

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

$$0 = A(s) \Big[(\mu_j - r_f) - \gamma (\boldsymbol{\Sigma} \boldsymbol{\theta})_j \Big] \\ + \lambda \sum_{s'} q(s'|s) \big[\pi_C + (1-\pi_C) A(s') \big] \cdot \kappa^{(j)}_{s,s'} \cdot \big(1 + \boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'} \big)^{-\gamma}. \quad (11)$$

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.

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^*.[6] 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

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

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 introduces a price-shift process, rather than merely a contemporaneous price shift. Formally, we consider the case where the asset price becomes marginally exposed to the shifter-state process $\{s_\tau\}_{\tau=t}^{\infty}$. Because $\ln P_j = \beta_j s + \ln \tilde{P}_j$, 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. We then define the *shifter-process elasticity*:

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

This expression 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 $d\beta_j s_t$ (as seen in the denominator), the date-t portfolio share $\theta_{j,t}^*$ generally responds to the dynamics of the entire future path of price shifts $\{d\beta_j s_\tau\}_{\tau=t}^{\infty}$. A shifter-process elasticity is specific to the dynamic properties of the shifter process associated with a given instrument.

While empirical variation creating an exposure to such a state process can help identify the shifter-process elasticity (15), the requirement for own-price elasticity identification are significantly 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 1. *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.[7]

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 1).

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 $j = 0, \dots, J$ 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

[7] KY19's "identification comes from cross-sectional variation in the investment universe and not from time-series variation in assets moving in and out of the investment universe." KY19 acknowledge that they "maintain the assumption that the wealth distribution across other investors is predetermined and exogenous to current demand shocks." and that "this assumption ultimately appeals to a static view of portfolio choice."

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 1 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., 2025b; Koijen and Yogo, 2025). Rather than investigating the empirical plausibility of such structural assumptions, our paper focuses on identification issues arising in dynamic environments in which long-lived assets such as equities are traded over time. Given this focus, our analysis is particularly concerned with the theoretical and empirical implications of violations of Condition 1 of Lemma 1.

First, we theoretically investigate how violations of Condition 1 of Lemma 1 affect elasticity estimates. The following proposition compares the true demand 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 1. The price-shift dynamics are illustrated in Figure 1.

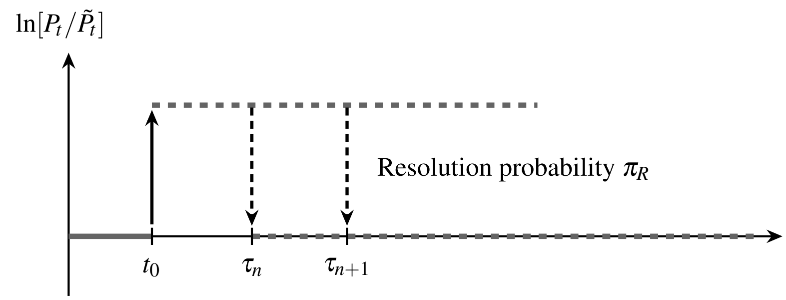

Figure 1: Stochastic price-shift resolution. The figure illustrates the price-shift dynamics considered in Proposition 2. At time t_0, the price level P_t is shifted up relative to the baseline level $\tilde{P}_t$, and at each Poisson arrival date $\tau_n > t_0$, this price shift is permanently resolved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

Proposition 2 (Stochastic price-shift resolution)**.** *Suppose conditions 2 and 3 of Lemma 1 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 1 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:*

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

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 dramatically underestimate the true elasticity, by a factor of $1/730 = 0.14\%$.

The first-order effect on the relative bias in the elasticity estimates is governed by the ratio of the expected per-period resolution return to the hypothetical return that would arise if the shift were to resolve fully in the first trading period. The resolution parameter π_R exactly represents this ratio; it measures the extent to which price-shift resolution is spread out over time rather than occurring entirely in the first trading period.

While this result suggests that a bias-adjustment factor of $1/\pi_R$ could, in principle, be applied to debias elasticity estimates, the following proposition highlights the limitations of such an approach when price shifts do not resolve in the first trading period. This case is illustrated in Figure 2.

Proposition 3 (No resolution in the first trading period)**.** *Suppose conditions 2 and 3 of Lemma 1 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 being 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 in the first trading period, a plausible scenario for a variety of candidate instruments. This occurs because investors set portfolio weights primarily based on the return distribution over the 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.

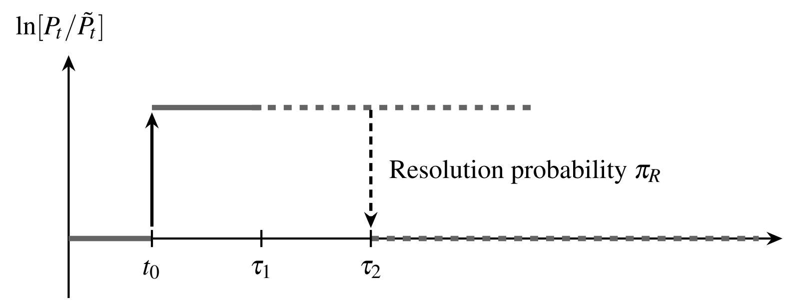

Figure 2: No resolution in the first trading period. The figure illustrates the price-shift dynamics considered in Proposition 3. At time t_0, the price level P_t is shifted upward relative to the baseline level $\tilde{P}_t$. This shift is not resolved in the first trading period, implying that it is 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.

Next, 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 instrumented prices.

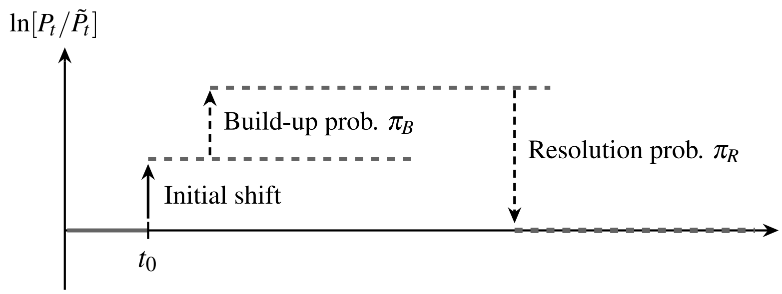

Figure 3: Price-shift build-up (momentum). The figure illustrates the price-shift dynamics considered in Proposition 4. Initially, at time $t = 0$, the price level P_t is shifted up relative to the baseline level $\tilde{P}_t$. Thereafter, at each subsequent Poisson arrival date $\tau_n > 0$, the shift 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 1 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*:*

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = -\,\pi_B + O(\lambda^{-2}) \;\to\; -\,\pi_B \quad as \quad \lambda \to \infty. \tag{19}$$

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.[8] We then proceed to 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 given stochastic price path differential in equilibrium.

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

$$m_{j,t} = x_j - \frac{\theta_{j,t}^* W_t}{P_{j,t}}, \qquad j = 0, \ldots, J, \tag{20}$$

which implies unit market clearing at every instant:

$$n_{j,t} + m_{j,t} = x_j, \quad \text{with} \quad n_{j,t} \equiv \frac{\theta_{j,t}^* W_t}{P_{j,t}} \qquad j = 0, \ldots, J. \tag{21}$$

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. The present model similarly captures such shocks but reverse engineers their magnitudes so as to

[8]In general, in an economy with heterogeneous agents, the wealth distribution emerges as a relevant state variable.

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, central banks, or corporations engaging in state-contingent share issuance or repurchases.

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:

$$dm_{j,t} = -\frac{dn_{j,t}}{dP_{j,t}} dP_{j,t}, \tag{22}$$

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

$$\underbrace{\frac{dP_{j,t}}{P_{j,t}}}_{\text{Market cap change}} = \underbrace{\frac{dm_{j,t}}{x_j}}_{\text{Investment}} \cdot \underbrace{\left[-\frac{d\ln n_{j,t}}{d\ln P_{j,t}} \cdot \frac{n_{j,t}}{x_j}\right]^{-1}}_{\text{Multiplier}}. \tag{23}$$

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:

$$\begin{aligned} -\frac{d\ln n_{j,t}}{d\ln P_{j,t}} &= -\left(\frac{d\theta^*_{j,t}}{dP_{j,t}}\frac{W_t}{P_{j,t}} + \frac{dW_t}{dP_{j,t}}\frac{\theta^*_{j,t}}{P_{j,t}} - \frac{W_t}{P_{j,t}}\frac{\theta^*_{j,t}}{P_{j,t}}\right)\frac{P_{j,t}}{n_{j,t}} \\ &= \varepsilon_{j,t} + 1 - \theta^*_{j,t}. \end{aligned} \tag{24}$$

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:*

$$\left\{\frac{dm_{j,\tau}}{x_j}\right\}_{\tau=t}^{\infty} = \left\{d\beta_j s_\tau \cdot [\eta_{j,\tau} + (1-\theta^*_{j,\tau})] \cdot \frac{n_{j,\tau}}{x_j}\right\}_{\tau=t}^{\infty}, \tag{25}$$

$$\left\{\frac{dm_{i,\tau}}{x_i}\right\}_{\tau=t}^{\infty} = \left\{d\beta_j s_\tau \cdot [\eta_{ij,\tau} - \theta^*_{j,\tau}] \cdot \frac{n_{i,\tau}}{x_i}\right\}_{\tau=t}^{\infty} \quad \forall i \neq j, \tag{26}$$

where we define the cross-shifter-process elasticity

$$\eta_{ij,\tau} \equiv -\frac{1}{s_t} \frac{d \ln \theta_i^*(s_t)}{d\beta_j}\bigg|_{\beta_i=0,\forall i}. \tag{27}$$

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 (25) and (26). 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 $\{\theta_{j,\tau}^*\}_{\tau=t}^{\infty}$. In other words, the *pseudo multiplier*, defined as the apparent impact of the date-t purchase quantity on the date-t price,

$$\frac{s_t d\beta_j}{dm_{j,t}/x_j} = \left[(\eta_{j,t} + (1-\theta_{j,t}^*)) \cdot n_{j,t}/x_t\right]^{-1}, \tag{28}$$

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 (26). 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 (28) 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 purchase quantity dynamics 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 potential investors taking the other side of the purchases are CRRA investors with a relative risk aversion of 2 and initial holdings equal to 60% of shares outstanding, i.e., $n_{1,t}/x_{1,t} = 0.6$. Assume that, prior to the announcement of the dynamic purchase program, the market risk premium is 6%, volatility is 20%, and the dividend yield is 1.7%. In this setting, CRRA investors' risky portfolio share before the unanticipated intervention is then given by:

$$\theta_1^* = \frac{\mu_1 - r_f}{\gamma \sigma_1^2} = \frac{0.06}{2 \cdot 0.2^2} = 0.75. \tag{29}$$

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:

$$\frac{d\mu_1}{d\beta_1} = \frac{d\frac{D_1}{\tilde{P}_1 e^{\beta_1 \cdot 1}}}{d\beta_1} = -\frac{D_1}{P_1}. \tag{30}$$

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

$$\frac{d\theta_1^*}{d\beta_1} = \frac{-D_1/P_1}{\gamma \cdot \sigma_1^2} = -0.2125. \tag{31}$$

That is, for a 1% price shift, $\Delta\beta_1 = 0.01$, the risky portfolio share decreases from 0.75 to 0.748. We then obtain the following shifter-process elasticity:

$$\eta_1 = -\frac{1}{s}\frac{d\ln\theta_1^*}{d\beta_1}\bigg|_{\beta=0} = \frac{D_1/P_1}{\mu_1 - r_f} = \frac{0.017}{0.06} = 0.28, \tag{32}$$

that is, a number close to zero. Further, converting this to the corresponding shifter-process elasticity in terms of unit demand gives $\eta_1 + (1-\theta_1^*) = 0.53$. Substituting these values into equation (28) yields a pseudo multiplier equal to:

$$\frac{d\beta_1}{dm_{1,t}/x_1} = \left[(\eta_{1,t} + (1-\theta_{1,t}^*)) \cdot n_{1,t}/x_t\right]^{-1} = 3.14. \tag{33}$$

However, as emphasized earlier, it is the entire process of future holdings that induces the price increase at the inception of the program. In particular, equation (25) simplifies to the following dynamics for the incremental holdings of the intervening party:

$$\left\{\frac{dm_{1,\tau}}{x_1}\right\}_{\tau=t}^{\infty} = \{d\beta_1 \,\cdot\, 0.53 \,\cdot\, n_{1,\tau}/x_1\}_{\tau=t}^{\infty}, \tag{34}$$

To raise the total market capitalization by $\Delta\beta_1 = 1\%$, the outside block must purchase $1\%/3.14 = 0.32\%$ 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):

$$\left\{dm_{0,\tau}\right\}_{\tau=t}^{\infty} = \left\{-dm_{1,\tau}\frac{P_{1,\tau}}{P_{0,\tau}}\right\}_{\tau=t}^{\infty}. \tag{35}$$

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 step into any gap in risk-free asset holdings by CRRA investors, relative to the counterfactual. As a result, we obtain a stochastic process for bond holdings adjustments, $\left\{dm_{0,\tau}\right\}_{\tau=t}^{\infty}$, 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, $\left\{dm_{1,\tau}\right\}_{\tau=t}^{\infty}$. For example, consider a path where the stock market outperforms the risk-free asset. In the counterfactual without intervention, CRRA agents would have become wealthier and increased their bond holdings over time. With the intervention, the outside block must offset this effect by progressively increasing its bond holdings as the stock market's relative valuation rises.

Overall, even in this stylized case, the asset purchase program is far more complex than what the multiplier suggests, as it requires managing an entire stochastic path of holdings in both 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 (20), persistent price movements can be supported by smaller variations in the outsiders' net supply, $m_{j,t}$, upon impact. In contrast, the transitory price dynamics required for identification under Lemma 1 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 1 (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*

$$\left|dm_1(\pi_R')\right| < \left|dm_1(\pi_R'')\right| \quad \textit{for all} \quad 0 \le \pi_R' < \pi_R'' \le 1.$$

Proof. See Appendix A.6. □

While Corollary 1 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%.

Simply put, more persistent shocks require much smaller supply interventions upon impact to achieve 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 endogenously from optimal dynamic trading behavior in response to changes in asset characteristics—specifically, the dynamics of expected resale values.

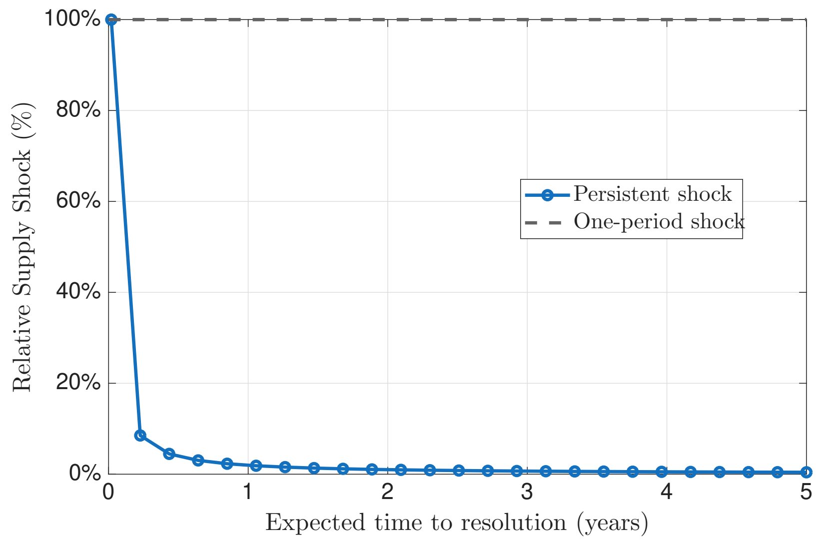

Figure 4: Supply Shock Magnitudes Upon Impact. The figure illustrates the magnitude of supply shocks required to induce a given log-price shift upon impact, as a function of the shock's persistence, measured by the expected time to resolution, $\frac{1}{\lambda \pi_R}$. Throughout, the inception of supply shocks is assumed to be unanticipated. We plot the relative supply shock, $\frac{|\Delta m_j(\pi_R)|}{|\Delta m_j(1)|}$, 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.

Table 1: Calibration and Relevant Statistics

Parameter	Symbol	Value	Unit / Note
Panel A. Investor preferences and timing			
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	–

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 from the perspective of institution i as

$$\widehat{me}_i(n) = \log\left[\sum_{j\neq i} A_j \frac{\mathbf{1}_j(n)}{1+\sum_{m=1}^{N}\mathbf{1}_j(m)}\right], \tag{36}$$

that is, the log market value of firm n under the counterfactual where all institutions other than i hold an equal-weighted portfolio within their investment universe.

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.

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

$$\frac{\theta_{i,t}(n)}{\theta_{i,t}(0)} = \exp\{\beta_{0,i,t} me_t(n) + \sum_{k=1}^{K-1} \beta_{k,i,t} x_{k,t}(n) + \beta_{K,i,t}\} \cdot \varepsilon_{i,t}(n), \tag{37}$$

where $x_{k,t}(n)$ denotes observable fundamental k for asset n at time t, and $\varepsilon_{i,t}(n)$ 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 supply curve but is assumed to be orthogonal to $\varepsilon_{i,t}(n)$, that is,

$$\mathbb{E}\left[\varepsilon_{i,t}(n) | \widehat{me}_{i,t}(n), \mathbf{x}_t(n)\right] = 1. \tag{38}$$

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. Table 2 reports these statistics for the stationary variables in our dataset. For non-stationary variables, we present corresponding statistics for successive five-year periods in Table 3. 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*) is the sum of dividend payouts over the past 12 months.

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.[9] 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

$$me_t(n) = a_t + b_t \times \widehat{me}_t(n) + \mathbf{c}_t'\mathbf{x}_t(n) + \varepsilon_t(n), \tag{39}$$

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 $\mathbf{x}_t(n)$.

[9] 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).

Table 2: Summary Statistics: Stationary Variables

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 3: Summary Statistics: Non-Stationary Variables

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

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 Market Equity 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 a myopic investor, the relevant return distribution is that up to the next trading date.

Correspondingly, we now examine how cross-sectional variation in instrumented market equity $\widehat{me}_{i,t}(n)$ induces cross-sectional variation in expected returns over different horizons, controlling for standard asset pricing factors commonly used by institutional investors.

If cross-sectional variation in $\widehat{me}_{i,t}(n)$ does not translate into resolution returns by the next trading date, a myopic 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 myopic investor will have no incentive to tilt her portfolio weights away from the higher-priced stock.

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, provided that both stocks are exposed to idiosyncratic risk.

In contrast, in the former case of persistent mispricing, we compare two assets that for all intents and purposes look identical from the perspective of a myopic investor. This example motivates empirically testing our key conditions for identification, in particular, whether instrumented variation in market equity is associated with corresponding (full) resolution returns of the same magnitude over the next trading horizon.

Importantly, this response 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 100.2 basis points 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 a myopic investor.

This example motivates our empirical tests of identification condition 1 of Lemma 1 —specifically, whether instrumented variation in market equity is associated with corresponding (full) resolution returns of the same magnitude over the next trading horizon.

To begin, we examine the extent to which cross-sectional variation in instrumented market equity 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:[10]

$$R_{(t,H)}(n) = a_H + b_H \times \widehat{me}_t(n) + c_H \times \tilde{\mathbf{X}}_{n,t} + e_{n,t,H}$$

where $\tilde{\mathbf{X}}$ contains cross-sectionally standardized stock characteristics: log market equity, book-to-market ratio, momentum, investment, profitability, dividend-to-book ratio, and market beta. As the tables will show, some specifications include all components of $\tilde{\mathbf{X}}$, while others use only a subset.

In Table 4, each regressor is standardized as a within-period z-score. In Table 5, we alternatively map cross-sectional ranks onto the interval $[-1,1]$ (see, e.g., Gu et al., 2020). Moreover, in Section 4.3, we consider a specification that does not apply any standardization.

We separately examine the extent to which cross-sectional dispersion at time t predicts returns over the first month (reflecting the initial trading period) and over subsequent periods. Specifically, we define the following four horizons: $H = 1$ month

[10]For these return predictability tests, we use the pooled version of instrumented market equity rather than the institution-specific version.

(return over the first month of quarter $t+1$), $H=1$ quarter (return over months 2 and 3 of quarter $t+1$), $H=4$ quarters (return over months 2 through 12), and $H=6$ quarters (return over months 2 through 18).

To avoid overlapping data, we retain only every H-th observation within each PERMNO-specific time series. Standard errors are two-way clustered. The tables report coefficient estimates with p-values shown in parentheses. Further details on the data are provided in Appendix B.

The coefficient on instrumented market equity varies in sign depending on the control variables included and the return horizon considered. However, at the one-month horizon, the relationship between instrumented market equity and subsequent returns is consistently positive. Moreover, at longer horizons, the relationship remains positive whenever size (uninstrumented market equity), a standard asset pricing (or risk) factor widely used by institutional investors, is included as a control. In other words, a cross-sectional increase in a firm's instrumented market equity is associated with *higher* average future returns.

While many of the coefficients on $\widehat{me}$ are positive, only some are statistically significant. In particular, coefficient estimates are significant when considering rank-standardized regressors (see Table 5) and including size as a control variable at the 1-month horizon. This pattern could justify tilts toward stocks with higher instrumented market equity, consistent with negative estimated elasticities when using this instrument. Indeed, in Section 4.4, we find that estimated elasticities are negative for a large fraction of stocks and institutions.

The coefficients on profitability (investment) are strongly positive (negative), consistent with the documented profitability premium and investment premium in the literature. Both are highly statistically significant. The coefficient on size is negative and statistically significant at horizons longer than one month. The coefficients on CAPM beta and dividends are generally insignificant. Finally, the coefficient on value is typically positive, consistent with the value premium, while the coefficient on momentum tends to be negative in these multivariate specifications.

Table 4: Return predictability regressions with z-score standardized regressors. The table reports panel regressions of excess returns on standard asset pricing characteristics (profitability, investment, dividends, beta, size, value, and momentum) as well as the instrumented log market equity $\widehat{me}$. The regressors are standardized in each period (z-scored). The table reports the coefficient estimates as well as the p-values in parentheses. The panels represent the following horizons: 1 month (return over the first month of quarter $t+1$), one quarter (return over the last two months of quarter $t+1$, i.e., months 2 and 3), 4 quarters (i.e., months 2 through 12) and 6 quarters (i.e, months 2 through 18).

(a) One Month

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	0.22	—	0.56	0.21	0.19	0.45
	(0.28)		(0.19)	(0.30)	(0.36)	(0.23)
PROF	0.41	0.43	0.43	0.47	0.42	0.49
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)
INV	-0.31	-0.32	-0.29	-0.26	-0.31	-0.25
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.02)
DIV	-0.05	-0.02	-0.03	-0.02	-0.06	-0.01
	(0.57)	(0.77)	(0.75)	(0.83)	(0.53)	(0.94)
BETA	0.14	0.15	0.12	0.20	0.14	0.18
	(0.53)	(0.49)	(0.58)	(0.37)	(0.53)	(0.40)
SIZE	—	-0.00	-0.36	—	—	-0.28
		(0.99)	(0.33)			(0.39)
VALUE	—	—	—	0.57	—	0.52
				(0.01)		(0.01)
MOM	—	—	—	—	-0.19	-0.10
					(0.35)	(0.58)

(b) One Quarter

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	-0.35	—	0.43	-0.35	-0.28	0.73
	(0.17)		(0.40)	(0.17)	(0.28)	(0.15)
PROF	0.47	0.51	0.52	0.46	0.45	0.51
	(0.04)	(0.03)	(0.03)	(0.06)	(0.05)	(0.04)
INV	-0.72	-0.70	-0.68	-0.73	-0.72	-0.67
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-0.12	-0.06	-0.07	-0.13	-0.11	-0.04
	(0.35)	(0.61)	(0.60)	(0.36)	(0.40)	(0.77)
BETA	-0.21	-0.23	-0.25	-0.22	-0.21	-0.27
	(0.45)	(0.41)	(0.36)	(0.41)	(0.45)	(0.31)
SIZE	—	-0.56	-0.84	—	—	-1.06
		(0.00)	(0.03)			(0.01)
VALUE	—	—	—	-0.12	—	-0.01
				(0.68)		(0.96)
MOM	—	—	—	—	0.40	0.50
					(0.11)	(0.03)

(c) Four Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	-0.11	—	4.08	-0.21	-0.34	3.52
	(0.93)		(0.14)	(0.87)	(0.80)	(0.19)
PROF	3.06	3.27	3.30	3.36	3.10	3.55
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-2.82	-2.77	-2.57	-2.61	-2.80	-2.41
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-0.86	-0.55	-0.57	-0.72	-0.89	-0.48
	(0.13)	(0.29)	(0.27)	(0.21)	(0.12)	(0.36)
BETA	-0.81	-0.81	-1.05	-0.52	-0.79	-0.77
	(0.40)	(0.39)	(0.28)	(0.56)	(0.40)	(0.40)
SIZE	—	-1.80	-4.42	—	—	-4.04
		(0.02)	(0.02)			(0.03)
VALUE	—	—	—	2.33	—	2.01
				(0.01)		(0.02)
MOM	—	—	—	—	-1.19	-0.57
					(0.12)	(0.43)

(d) Six Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	0.10	—	6.13	-0.04	-0.33	5.04
	(0.96)		(0.08)	(0.98)	(0.86)	(0.12)
PROF	4.37	4.66	4.72	5.01	4.44	5.29
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-4.37	-4.31	-4.02	-3.96	-4.34	-3.68
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-0.65	-0.22	-0.27	-0.38	-0.71	-0.11
	(0.50)	(0.81)	(0.78)	(0.70)	(0.46)	(0.91)
BETA	-0.72	-0.71	-1.09	-0.16	-0.67	-0.52
	(0.56)	(0.58)	(0.37)	(0.89)	(0.59)	(0.66)
SIZE	—	-2.26	-6.22	—	—	-5.47
		(0.07)	(0.01)			(0.01)
VALUE	—	—	—	4.44	—	3.94
				(0.00)		(0.00)
MOM	—	—	—	—	-2.21	-1.25
					(0.02)	(0.15)

Table 5: Return predictability regressions with rank-standardized regressors. The table reports panel regressions of excess returns on standard asset pricing characteristics (profitability, investment, dividends, beta, size, value, and momentum) as well as the instrumented log market equity $\widehat{me}$. The regressors are rank-standardized in each period. The table reports the coefficient estimates as well as the p-values in parentheses. The panels represent the following horizons: 1 month (return over the first month of quarter $t+1$), one quarter (return over the last two months of quarter $t+1$, i.e., months 2 and 3), 4 quarters (i.e., months 2 through 12) and 6 quarters (i.e, months 2 through 18).

(a) One Month

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	0.36	—	1.35	0.33	0.31	1.08
	(0.22)		(0.06)	(0.25)	(0.27)	(0.07)
PROF	0.47	0.50	0.53	0.85	0.50	0.86
	(0.02)	(0.01)	(0.01)	(0.00)	(0.01)	(0.00)
INV	-0.41	-0.42	-0.33	-0.27	-0.42	-0.22
	(0.01)	(0.00)	(0.03)	(0.06)	(0.01)	(0.11)
DIV	-0.06	0.03	0.04	-0.17	-0.07	-0.08
	(0.80)	(0.89)	(0.86)	(0.48)	(0.78)	(0.71)
BETA	0.20	0.24	0.13	0.34	0.19	0.28
	(0.56)	(0.47)	(0.69)	(0.32)	(0.57)	(0.38)
SIZE	—	-0.14	-1.54	—	—	-1.16
		(0.73)	(0.11)			(0.16)
VALUE	—	—	—	0.91	—	0.85
				(0.00)		(0.00)
MOM	—	—	—	—	-0.32	0.04
					(0.02)	(0.89)

(b) One Quarter

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	-0.48	—	0.48	-0.49	-0.37	1.18
	(0.18)		(0.55)	(0.18)	(0.30)	(0.13)
PROF	0.68	0.73	0.74	0.81	0.59	0.87
	(0.01)	(0.00)	(0.00)	(0.04)	(0.02)	(0.03)
INV	-1.03	-0.99	-0.95	-0.99	-1.03	-0.81
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-0.16	-0.06	-0.06	-0.20	-0.14	-0.04
	(0.65)	(0.86)	(0.86)	(0.55)	(0.69)	(0.89)
BETA	-0.29	-0.32	-0.36	-0.24	-0.26	-0.26
	(0.48)	(0.44)	(0.39)	(0.53)	(0.53)	(0.49)
SIZE	—	-1.00	-1.50	—	—	-2.35
		(0.02)	(0.13)			(0.01)
VALUE	—	—	—	0.30	—	0.57
				(0.50)		(0.17)
MOM	—	—	—	—	0.86	1.27
					(0.02)	(0.00)

(c) Four Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	0.04	—	7.60	-0.13	-0.05	7.25
	(0.98)		(0.09)	(0.94)	(0.98)	(0.10)
PROF	3.56	3.86	3.99	5.61	3.63	5.76
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-3.80	-3.68	-3.17	-3.06	-3.80	-2.48
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-1.17	-0.43	-0.39	-1.74	-1.19	-0.95
	(0.41)	(0.75)	(0.77)	(0.21)	(0.40)	(0.44)
BETA	-1.24	-1.12	-1.74	-0.42	-1.26	-0.86
	(0.38)	(0.41)	(0.23)	(0.74)	(0.37)	(0.52)
SIZE	—	-3.83	-11.78	—	—	-11.13
		(0.02)	(0.02)			(0.02)
VALUE	—	—	—	4.78	—	4.64
				(0.00)		(0.00)
MOM	—	—	—	—	-0.69	1.78
					(0.62)	(0.15)

(d) Six Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\widehat{me}$	-0.22	—	9.47	-0.52	-0.55	7.30
	(0.93)		(0.06)	(0.84)	(0.82)	(0.12)
PROF	4.81	5.17	5.33	8.66	5.05	8.79
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-5.80	-5.62	-5.00	-4.48	-5.83	-3.90
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	0.78	1.69	1.76	-0.19	0.73	0.65
	(0.74)	(0.46)	(0.44)	(0.93)	(0.75)	(0.76)
BETA	-0.40	-0.28	-1.08	1.06	-0.45	0.47
	(0.84)	(0.89)	(0.58)	(0.60)	(0.82)	(0.81)
SIZE	—	-4.96	-14.89	—	—	-11.79
		(0.05)	(0.01)			(0.01)
VALUE	—	—	—	8.72	—	8.31
				(0.00)		(0.00)
MOM	—	—	—	—	-2.59	1.04
					(0.13)	(0.49)

4.3 Resolution Tests

In the previous section, we examined panel regressions of cumulative individual stock returns at various horizons H on standard asset pricing factors, augmented with instrumented market equity. There, we normalized each variable either by ranking it and mapping it onto the interval $[-1,1]$, or by standardizing it as a z-score.

In this section, we consider a specification that even more directly relates to Identification Condition 1 in Lemma 1—namely, that a price shift used for elasticity identification must resolve fully by the next trading date. To examine this quantitative restriction, we focus on the variable $\hat{b} \times \widehat{me}$, which, under the assumption that the instrument shifts prices for non-fundamental reasons (controlling for the observables X), captures the corresponding log price shift. Moreover, we no longer normalize variables in this analysis. Thus, if institutions can trade monthly, full resolution within that period would correspond to a coefficient of -1 on $(\hat{b} \times \widehat{me})$ at the one-month horizon.

The results are reported in Table 6. Qualitatively, the findings are similar to those reported earlier. Once again, the coefficient on instrumented market equity varies in sign depending on the control variables included and the return horizon considered. At the one-month horizon, the point estimates for the relationship between instrumented market equity and subsequent returns remain consistently positive across specifications. Moreover, at longer horizons, the relationship is typically positive whenever size (uninstrumented market equity) is included as a control.

Importantly, 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 0.00.[11] This provides direct evidence that the instrument fails to satisfy the full-resolution requirement for elasticity identification stated in Lemma 1.

[11] The p-values reported in the table refer to the null hypothesis of 0.

Table 6: Price shift resolution regressions. The table reports panel regressions of excess returns on standard asset pricing characteristics (profitability, investment, dividends, beta, size, value, and momentum) as well as $\hat{b} \times \widehat{me}$. Variables are not standardized. The table reports the coefficient estimates as well as the p-values corresponding to the null hypothesis of zero in parentheses. 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 0.00. The panels represent the following horizons: 1 month (return over the first month of quarter $t+1$), one quarter (return over the last two months of quarter $t+1$, i.e., months 2 and 3), 4 quarters (i.e., months 2 through 12) and 6 quarters (i.e, months 2 through 18).

(a) One Month

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\hat{b} \times \widehat{me}$	0.04	—	0.09	0.06	0.03	0.08
	(0.71)		(0.56)	(0.61)	(0.78)	(0.60)
PROF	1.24	1.30	1.37	1.56	1.18	1.55
	(0.04)	(0.03)	(0.02)	(0.01)	(0.01)	(0.01)
INV	-1.57	-1.65	-1.51	-1.14	-1.58	-1.17
	(0.07)	(0.07)	(0.09)	(0.20)	(0.07)	(0.18)
DIV	-0.14	0.12	1.96	1.89	-1.26	2.12
	(0.97)	(0.98)	(0.64)	(0.63)	(0.74)	(0.63)
BETA	0.30	0.32	0.29	0.39	0.35	0.42
	(0.34)	(0.33)	(0.35)	(0.23)	(0.27)	(0.20)
SIZE	—	-0.05	-0.18	—	—	-0.12
		(0.72)	(0.34)			(0.48)
VALUE	—	—	—	0.73	—	0.63
				(0.02)		(0.03)
MOM	—	—	—	—	-0.94	-0.76
					(0.04)	(0.07)

(b) One Quarter

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\hat{b} \times \widehat{me}$	-0.05	—	0.02	-0.05	-0.05	0.03
	(0.69)		(0.89)	(0.72)	(0.71)	(0.85)
PROF	0.96	1.16	1.17	1.03	0.97	1.27
	(0.20)	(0.13)	(0.11)	(0.18)	(0.18)	(0.07)
INV	-3.92	-3.86	-3.83	-3.82	-3.91	-3.72
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-9.23	-6.15	-5.73	-8.76	-9.02	-4.85
	(0.13)	(0.37)	(0.35)	(0.16)	(0.13)	(0.44)
BETA	-0.40	-0.40	-0.41	-0.37	-0.41	-0.40
	(0.27)	(0.27)	(0.25)	(0.30)	(0.25)	(0.25)
SIZE	—	-0.27	-0.30	—	—	-0.30
		(0.13)	(0.04)			(0.03)
VALUE	—	—	—	0.17	—	0.18
				(0.41)		(0.31)
MOM	—	—	—	—	0.18	0.28
					(0.74)	(0.61)

(c) Four Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\hat{b} \times \widehat{me}$	-0.10	—	0.23	-0.02	-0.18	0.14
	(0.88)		(0.74)	(0.98)	(0.76)	(0.83)
PROF	7.95	8.69	8.81	9.58	7.43	9.44
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-15.69	-15.67	-15.33	-13.73	-15.71	-13.92
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-34.65	-24.36	-19.35	-25.44	-41.11	-21.19
	(0.15)	(0.45)	(0.43)	(0.31)	(0.09)	(0.39)
BETA	-0.79	-0.79	-0.86	-0.27	-0.41	-0.12
	(0.60)	(0.59)	(0.56)	(0.86)	(0.78)	(0.93)
SIZE	—	-0.96	-1.27	—	—	-1.00
		(0.23)	(0.03)			(0.06)
VALUE	—	—	—	3.25	—	2.49
				(0.00)		(0.00)
MOM	—	—	—	—	-5.37	-4.64
					(0.00)	(0.00)

(d) Six Quarters

Regressor	(a)	(b)	(c)	(d)	(e)	(f)
$\hat{b} \times \widehat{me}$	-0.67	—	-0.24	-0.55	-0.79	-0.37
	(0.37)		(0.77)	(0.44)	(0.25)	(0.62)
PROF	12.73	14.07	13.95	15.31	11.70	14.75
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
INV	-24.21	-23.29	-23.65	-21.48	-24.31	-21.72
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
DIV	-29.71	-3.09	-8.47	-16.37	-39.48	-12.40
	(0.36)	(0.94)	(0.79)	(0.62)	(0.20)	(0.69)
BETA	-0.70	-0.87	-0.79	0.04	-0.04	0.40
	(0.68)	(0.61)	(0.64)	(0.98)	(0.98)	(0.81)
SIZE	—	-2.04	-1.71	—	—	-1.23
		(0.05)	(0.04)			(0.10)
VALUE	—	—	—	4.40	—	3.50
				(0.00)		(0.00)
MOM	—	—	—	—	-9.22	-8.11
					(0.00)	(0.00)

4.4 Unrestricted Elasticity Estimates

The results in the previous section show that, in several specifications, the relationship between instrumented market equity and returns is positive, not negative. The idea that investors might allocate more to such "expensive" stocks is consistent with the price-shift build-up patterns documented in the literature (see Binsbergen et al., 2023). In a static framework, higher allocations to more expensive stocks, as measured by instrumented market equity, could therefore lead to erroneously estimated negative price elasticities.

KY19 impose the assumption that demand curves are downward sloping, that is, they restrict elasticities to be positive, by requiring $\beta_{0,i,t} < 1$. In this section, we reestimate their demand equations (37) without imposing this restriction to assess how binding the constraint is in practice. We consider the sample of institutions with at least 1,000 stock holdings at any given point in time. Over the sample period 1980Q1-2017Q4, this yields 9,289 institution-date pairs for which we estimate the demand equation.

Results. Panel A of Table 7 reports the unconstrained beta estimates of the demand equations for a random subsample of 1,000 date-manager pairs (columns 1 to 4) and compares them with the constrained estimates (imposing $\beta_{0,i,t} < 1$) in column 5. Columns 1 and 3 use the institution-specific instruments ($\widehat{me}_{i,t}(n)$), whereas columns 2 and 4 report estimates based on the pooled instrumented market equity ($\widehat{me}_{HH,t}(n)$). For robustness, we further consider specifications where zero-weight observations are included (columns (1) and (2)) and where they are not included (columns (3) and (4)). Panel A shows that 41-53% of the unconstrained beta estimates violate KY19's assumption that $\beta_{0,i,t} < 1$. Conditional on violating the constraint, the average estimate ranges between 1.24 and 1.38.

Next, in Panel B of Table 7, we examine the implications of these beta estimates for asset-level elasticities. According to KY19, for each date-manager pair, the point estimate of $\beta_{0,i,t}$ can be used to compute the matrix of cross-elasticities as follows:

$$\Psi_{i,t} = \mathbf{I} - \beta_{0,i,t}\, \text{diag}(\mathbf{w})^{-1} \times (\text{diag}(\mathbf{w}) - \mathbf{w}\mathbf{w}'), \tag{40}$$

where $\mathbf{w}$ is the vector of positive portfolio weights for institution i at time t, and $\mathbf{I}$ and

Table 7: Summary Statistics for Estimated βs and Elasticities.
The table reports summary statistics for estimated betas and demand elasticities considering five different specifications:
(1) Instrument is $IVme$, zero-weight observations are included.
(2) Instrument is $IVme_{HH}$, zero-weight observations are included.
(3) Instrument is $IVme$, zero-weight observations are not included.
(4) Instrument is $IVme_{HH}$, zero-weight observations are not included.
(5) Instrument is $IVme$, zero-weight observations are included, and an upper bound of 0.99 is imposed for the point estimate of $\beta_{0,i,t}$.

	Panel A: Estimated βs					**Panel B: Implied Elasticities**				
Statistic	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Mean	1.02	1.03	0.92	0.90	0.85	−0.04	−0.07	0.06	0.05	0.15
Std	0.41	0.50	0.37	0.40	0.24	0.42	0.44	0.38	0.40	0.26
p25	0.79	0.72	0.71	0.63	0.78	−0.31	−0.38	−0.17	−0.22	0.01
p50	1.03	1.01	0.94	0.90	0.99	−0.08	−0.06	0.01	0.06	0.01
p75	1.26	1.29	1.16	1.14	0.99	0.19	0.22	0.27	0.31	0.20
$\Pr[\beta > 1]$	0.53	0.51	0.43	0.41	0.00			—		
$\mathbb{E}[\beta \mid \beta > 1]$	1.30	1.38	1.24	1.26	—			—		
$\Pr[\hat{\varepsilon} < 0]$			—			0.58	0.55	0.49	0.46	0.00

w have corresponding dimensions.

Focusing on the diagonal elements of $\Psi_{i,t}$ across a random sample of 300 date-manager pairs, we obtain the distribution of demand elasticities reported in Panel B of Table 7. For each specification considered in columns (1) through (5) in Panel A we report the corresponding demand elasticity estimates in Panel B. The table shows that when we remove the constraint $\beta_{0,i,t} < 1$, the fraction of estimated elasticities that are negative varies from 0.46 to 0.58 across specifications.

4.5 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 scenario where identification could succeed is if investors mistakenly believed, throughout the entire sample period, that the cross-sectional price dispersion induced by $\widehat{me}$ would lead to fully resolving returns within a single trading period, even though this consistently failed to occur.

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

More broadly, our results call for a reassessment of how and whether existing elasticity and multiplier estimates should be used to inform asset-pricing models and to evaluate the efficacy of central bank purchase programs. We highlight that empirical research must explicitly account for the temporal structure of supply shocks and price formation. Failure to do so risks conflating frictionless dynamic-optimization behavior with investor inelasticity, leading to incorrect inferences about investor behavior and the functioning of financial markets.

References

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.

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.

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.

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

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," October 2025. Preliminary. Most recent version.

Gabaix, Xavier and Ralph S. J. Koijen, "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," NBER Working Paper 28967, National Bureau of Economic Research 2021.

_ 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.

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, Alan Moreira, and Tyler Muir, "Whatever It Takes? The Impact of Conditional Policy Promises," *American Economic Review*, January 2025, *115* (1), 295–329.

_ , 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.

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

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.

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.

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.

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

$$V(W,s) = \upsilon \log W + A(s), \tag{41}$$

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

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

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 $A(s)$:

$$\begin{aligned}
0 \;=\; &-\rho A(s) \;+\; \upsilon\Big[r_f + (\boldsymbol{\mu} - r_f \mathbf{1})^\top \boldsymbol{\theta} - \tfrac{1}{2}\boldsymbol{\theta}^\top \boldsymbol{\Sigma}\, \boldsymbol{\theta}\Big] \\
&+ \lambda \sum_{s'} q(s'|s)\Big[(1-\pi_C)A(s') - A(s) + \widetilde{C} \cdot \log\big(1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}\big)\Big],
\end{aligned} \tag{42}$$

where we define:

$$\widetilde{C} \equiv \pi_C + (1-\pi_C)\, \upsilon. \tag{43}$$

Differentiating (42) with respect to each θ_j yields the first-order conditions

$$0 \;=\; \upsilon\Big[(\mu_j - r_f) - (\boldsymbol{\Sigma}\boldsymbol{\theta})_j\Big] + \lambda\,\widetilde{C}\sum_{s'} q(s'|s)\,\frac{\kappa^{(j)}_{s,s'}}{1+\boldsymbol{\theta}^\top \boldsymbol{\kappa}_{s,s'}}, \qquad j = 1,\ldots,J. \tag{44}$$

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

$$\left.\frac{d\theta_C^*(s)}{d\beta}\right|_{\beta=0} \tag{45}$$

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

$$0 = A(s)\left[\mu - r_f - \gamma\sigma^2\theta_C(s)\right] + \lambda\,\pi_R\,C(0)\,\kappa(s)\left[1 + \theta_C(s)\,\kappa(s)\right]^{-\gamma}, \tag{46}$$

where we define:

$$C(s) \equiv \pi_C + (1 - \pi_C)A(s). \tag{47}$$

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

$$\mu - r_f - \gamma\sigma^2\,\theta_C^*(s) = 0 \tag{48}$$

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

$$f(\theta, \beta; s) \equiv A(s)\left[\mu - r_f - \gamma\sigma^2\theta\right] + \lambda\,\pi_R C(0)\,\kappa(s)\,[1 + \theta\,\kappa(s)]^{-\gamma}, \tag{49}$$

The FOC (46) is $f(\theta_C(s),\beta;s)=0$. At $\beta=0$, we have $\kappa(s)=0$, so

$$f(\theta,0;s)=A(s)\big[\mu-r_f-\gamma\sigma^2\,\theta\big], \tag{50}$$

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 $(\theta_0,0)$:

$$\partial_\theta f(\theta_0,0;s)=-A(s)\gamma\sigma^2. \tag{51}$$

Moreover, as shown above, at $\beta=0$ we have $A(s)=A(0)$. For $\partial_\beta f$, the first term in (49) contributes:

$$\frac{\partial A(s)}{\partial\beta}\bigg|_{\beta=0}\underbrace{\big[\mu-r_f-\gamma\sigma^2\theta_0\big]}_{=0}=0, \tag{52}$$

because the FOC implies $\mu-r_f-\gamma\sigma^2\theta_0=0$. The derivative of the second term in (49) is:

$$\lambda\pi_R C(0)\left[\frac{\partial\kappa}{\partial\beta}\cdot[1+\theta_0\kappa]^{-\gamma}+\kappa\cdot(-\gamma)[1+\theta_0\kappa]^{-\gamma-1}\cdot\theta_0\cdot\frac{\partial\kappa}{\partial\beta}\right]. \tag{53}$$

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:

$$\partial_\beta f(\theta_0,0;s)=\lambda\pi_R C(0)\cdot(-s)\cdot 1=-\lambda\pi_R C(0)\,s. \tag{54}$$

By the implicit function theorem we then obtain:

$$\frac{d\theta_C^*(s)}{d\beta}\bigg|_{\beta=0}=-\frac{\partial_\beta f}{\partial_\theta f}\bigg|_{(\theta_0,0)}=-\frac{\lambda\,\pi_R\,C(0)}{A(0)\,\gamma\sigma^2}\,s. \tag{55}$$

It follows that the shifter-process elasticity is

$$\eta_C(s) = -\frac{1}{s}\left.\frac{\frac{d\theta_C^*(s)}{d\beta}}{\theta_C^*(0)}\right|_{\beta=0} = \pi_R\frac{\lambda\, C(0)}{(\mu - r_f)A(0)}. \tag{56}$$

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:

$$\eta_C \to \infty \quad \text{as } \lambda \to \infty. \tag{57}$$

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:

$$\begin{aligned} 0 = &-\rho A(0) + (1-\gamma)A(0)\left[r_f + (\mu - r_f)\theta(0) - \tfrac{1}{2}\gamma\theta(0)^2\sigma^2\right] \\ &+ \lambda\left[\left[\pi_C + (1-\pi_C)A(0)\right] - A(0)\right], \end{aligned} \tag{58}$$

and solving for $A(0)$ yields:

$$A(0) = \frac{\lambda\pi_C}{\lambda\pi_C + \rho - (1-\gamma)\left[r_f + (\mu - r_f)\theta(0) - \frac{1}{2}\gamma\theta(0)^2\sigma^2\right]}. \tag{59}$$

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:

$$A(0) = \frac{\lambda\pi_C}{\lambda\pi_C + \rho - (1-\gamma)\left[r_f + \frac{(\mu - r_f)^2}{2\gamma\sigma^2}\right]}. \tag{60}$$

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

$$n_1 = \frac{\theta_D W_\tau}{P_\tau}, \qquad n_0 = \frac{(1-\theta_D) W_\tau}{P_{0,\tau}}. \tag{61}$$

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

$$\begin{aligned} W_t &= n_1 P_t + n_0 P_{0,t} = \theta_D W_\tau \frac{P_t}{P_\tau} + (1-\theta_D) W_\tau \frac{P_{0,t}}{P_{0,\tau}} \\ &= \theta_D W_\tau \frac{P_t}{P_\tau} + (1-\theta_D) W_\tau e^{r_f(t-\tau)} > 0 \quad \text{a.s.} \end{aligned} \tag{62}$$

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

$$\theta_C(0) = \frac{\mu - r_f}{\gamma\sigma^2} \in (0,1) \tag{63}$$

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:

$$\theta_D^*(s;\lambda,\beta) \in (0,1) \qquad \text{for all } s \in \{-1,0,+1\}. \tag{64}$$

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

$$\theta_D^*(s;\lambda,\beta) = \theta_C(0) + O\left(\frac{1}{\lambda}\right) + O(\beta), \tag{65}$$

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

$$d\theta_t = \theta_t(1-\theta_t)\big(\mu - r_f - \sigma^2\theta_t\big)\,dt + \sigma\,\theta_t(1-\theta_t)\,dB_t, \qquad \theta_0 = \theta_D. \tag{66}$$

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

$$\mathbb{E}\left[\int_0^\Delta \theta_t^2\,dt\right] \;\leq\; \mathbb{E}[\Delta] \;=\; \frac{1}{\lambda} \;<\infty, \tag{67}$$

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

$$\frac{du(W_t)}{u(W_t)} = \alpha_t dt + \zeta_t dB_t. \tag{68}$$

where we define:

$$\alpha_t \equiv (1-\gamma)\Big[r_f + \theta_t(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_t^2\Big], \tag{69}$$

$$\zeta_t \equiv (1-\gamma)\sigma\,\theta_t. \tag{70}$$

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

$$\frac{u(W_\Delta)-u(W_0)}{u(W_0)}=\int_0^\Delta \frac{u(W_t)}{u(W_0)}\alpha_t dt+\int_0^\Delta \frac{u(W_t)}{u(W_0)}\zeta_t dB_t. \tag{71}$$

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:

$$\mathbb{E}\left[\int_0^\Delta \frac{u(W_t)}{u(W_0)}\zeta_t dB_t\right]=\mathbb{E}\left[\int_0^\Delta Z_t\,\zeta_t\, dB_t\right]=\mathbb{E}\Big[\,\mathbb{E}\Big[\int_0^\Delta Z_t\,\zeta_t\, dB_t\,\Big|\,\Delta\Big]\Big]=0. \tag{72}$$

Then

$$\lambda\,\mathbb{E}\left[\frac{u(W_\Delta)-u(W_0)}{u(W_0)}\right]=\lambda\,\mathbb{E}\int_0^\Delta Z_t\,\alpha_t\, dt=\lambda\int_0^\infty e^{-\lambda t}\,\mathbb{E}[Z_t\alpha_t]\,dt. \tag{73}$$

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

$$R(t)\equiv \mathbb{E}[(Z_t-1)\alpha_t], \tag{74}$$

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

$$dZ_t = Z_t\alpha_t\, dt + Z_t\zeta_t\, dB_t \tag{75}$$

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

$$m(t)=1+\int_0^t \mathbb{E}[Z_i\alpha_i]\,di \;\le\; 1+F\int_0^t m(i)\,di, \tag{76}$$

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

$$|R(t)|=\big|\,\mathbb{E}[(Z_t-1)\alpha_t]\big|\le F\,\mathbb{E}|Z_t-1|\le F\,(e^{F\cdot t}-1). \tag{77}$$

As a result, we obtain:

$$\lambda\int_0^\infty e^{-\lambda t}\,|R(t)|\,dt\;\le\; F\,\lambda\int_0^\infty e^{-\lambda t}\big(e^{F\cdot t}-1\big)\,dt\;=\;\frac{F^2}{\lambda-F}\;=\;O\Big(\frac{1}{\lambda}\Big), \tag{78}$$

where $\lambda > F$ is required to ensure that the integral is finite. Combining (73) 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

$$\begin{aligned} &\lambda\, \mathbb{E}\left[\frac{u(W_\Delta) - u(W_0)}{u(W_0)}\right] \\ =&\lambda(1-\gamma)\, \mathbb{E}\left[r_f \Delta \,+\, (\mu - r_f) \int_0^\Delta \theta_t\, dt \,-\, \tfrac{1}{2}\gamma\sigma^2 \int_0^\Delta \theta_t^2\, dt \;\middle|\; \theta_0 = \theta_D \right] + O(\lambda^{-1}). \end{aligned} \tag{79}$$

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 $\varphi \in C^2([0,1])$

$$\mathbb{E}_\theta[\varphi(\theta_t)] = \varphi(\theta) + \mathbb{E}_\theta \int_0^t \mathscr{L}\varphi(\theta_r)\, dr, \tag{80}$$

where the dynamics of θ given in (66) imply the generator

$$\mathscr{L}\varphi(x) = x(1-x)\big(\mu - r_f - \sigma^2 x\big)\, \varphi'(x) + \tfrac{1}{2}\sigma^2 x^2 (1-x)^2\, \varphi''(x). \tag{81}$$

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

$$\mathbb{E}_\theta[\varphi(\theta_t)] \;= \varphi(\theta) + t\, \mathscr{L}\varphi(\theta) + O(t^2) \tag{82}$$

$$\mathbb{E}_\theta \int_0^t \varphi(\theta_r) dr \;= t\, \varphi(\theta) + \tfrac{1}{2} t^2 \mathscr{L}\varphi(\theta) + O(t^3) \tag{83}$$

Integrating the first expansion (82) over $s \in [0,t]$ yields the second (83). Conditioning on $\theta_0 = \theta_D$ and Δ and taking expectations yields:

$$\mathbb{E}\left[\int_0^\Delta \varphi(\theta_r)\,dr\right] = \mathbb{E}\left[\Delta\,\varphi(\theta_D) + \tfrac{1}{2}\Delta^2 \mathscr{L}\varphi(\theta_D) + O(\Delta^3)\right]$$
$$= \frac{\varphi(\theta_D)}{\lambda} + \frac{\mathscr{L}\varphi(\theta_D)}{\lambda^2} + O(\lambda^{-3}), \tag{84}$$

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:

$$\mathscr{L}(x)(\theta_D) = \theta_D(1-\theta_D)(\mu - r_f - \sigma^2\theta_D) \tag{85}$$
$$\mathscr{L}(x^2)(\theta_D) = 2\theta_D^2(1-\theta_D)\big(\mu - r_f - \sigma^2\theta_D\big) + \sigma^2\theta_D^2(1-\theta_D)^2 \tag{86}$$

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

$$\mathbb{E}\left[\int_0^\Delta \theta_t\,dt\right] = \frac{\theta_D}{\lambda} + \frac{b(\theta_D)}{\lambda^2} + O(\lambda^{-3}), \tag{87}$$
$$\mathbb{E}\left[\int_0^\Delta \theta_t^2\,dt\right] = \frac{\theta_D^2}{\lambda} + \frac{2\theta_D^2(1-\theta_D)\big(\mu - r_f - \sigma^2\theta_D\big) + \sigma^2\theta_D^2(1-\theta_D)^2}{\lambda^2} + O(\lambda^{-3}). \tag{88}$$

Substituting (87), (88), and $\mathbb{E}[\Delta] = 1/\lambda$ into (79) yields:

$$\lambda\,\mathbb{E}\left[\frac{u(W_\Delta) - u(W_0)}{u(W_0)}\right] = (1-\gamma)\left[r_f + \theta_D\cdot(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\right]$$
$$+ \frac{1-\gamma}{\lambda} g(\theta_D) + O(\lambda^{-1}), \tag{89}$$

where we define the correction polynomial relative to continuous rebalancing:

$$g(\theta_D) \equiv (\mu - r_f)\,\theta_D(1-\theta_D)\big((\mu - r_f) - \sigma^2\theta_D\big)$$
$$-\ \gamma\sigma^2\theta_D^2(1-\theta_D)\big((\mu - r_f) - \sigma^2\theta_D\big)\ -\ \tfrac{1}{2}\gamma\sigma^4\theta_D^2(1-\theta_D)^2. \tag{90}$$

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:

$$\kappa(s) = e^{-\beta s} - 1 = -s\, d\beta + O(d\beta^2). \tag{91}$$

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

$$\mathbb{E}\big[(1+\theta_\Delta \kappa(s))^{1-\gamma} \mid \theta_0 = \theta_D\big] = 1 - (1-\gamma)\, \mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D]\, s\, d\beta + O(d\beta^2), \tag{92}$$

and by the Dynkin expansion,

$$\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D] = \theta_D + \frac{b(\theta_D)}{\lambda} + O(\lambda^{-2}). \tag{93}$$

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

$$\begin{aligned} 0 = & -\rho A(s) + (1-\gamma)A(s)\left[r_f + \theta_D \cdot (\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\right] + \frac{1-\gamma}{\lambda} A(s)\, g(\theta_D) \\ & + \lambda\left\{\pi_R C(0)\, \mathbb{E}\big[(1+\theta_\Delta \kappa(s))^{1-\gamma} \mid \theta_0 = \theta_D\big] + (1-\pi_R)\, C(s) - A(s)\right\} \\ & + O(\lambda^{-1}) + O(\lambda\, d\beta^2), \end{aligned} \tag{94}$$

Substituting (92) and (93) 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:

$$\begin{aligned} 0 = & (\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} - \lambda\, \pi_R \frac{C(0)}{A(s)}\left(1 + \frac{b'(\theta_D)}{\lambda}\right) s\, d\beta \\ & + O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1} d\beta). \end{aligned} \tag{95}$$

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:

$$f_\theta \left.\frac{d\theta_D^*}{d\beta}\right|_0 + f_\beta = 0, \tag{96}$$

with

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

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

$$\left.\frac{d\theta_D^*}{d\beta}\right|_{\beta=0} = -\frac{f_\beta}{f_\theta} = -\frac{\lambda\,\pi_R C(0)}{A(0)\,\gamma\sigma^2}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right)s + O(\lambda^{-1}). \tag{98}$$

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

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

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

$$\varepsilon_D = \frac{C(0)}{A(0)\,\gamma\sigma^2\theta_{D,0}^*}\left(\lambda + b'(\theta_{D,0}^*) + \frac{g''(\theta_{D,0}^*)}{\gamma\sigma^2}\right) + O(\lambda^{-1}). \tag{100}$$

Comparing (99) and (100), we see that $\eta_D = \pi_R\varepsilon_D + O(\lambda^{-1})$. Since $\varepsilon_D = O(\lambda)$, dividing yields:

$$\frac{\eta_D}{\varepsilon_D} = \pi_R + \frac{O(\lambda^{-1})}{\varepsilon_D} = \pi_R + O(\lambda^{-2}). \tag{101}$$

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

$$\frac{\hat{\varepsilon}_D}{\varepsilon_D} = \frac{\eta_D}{\varepsilon_D} = \pi_R + O(\lambda^{-2}) \to \pi_R \quad \text{as } \lambda \to \infty. \tag{102}$$

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:

$$q(1|0) = 1, \qquad\qquad q(0|0) = 0, \tag{103}$$

$$q(2|1) = \pi_R, \qquad\qquad q(1|1) = 1 - \pi_R, \tag{104}$$

$$q(2|2) = 1. \tag{105}$$

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 (89), the Dynkin expansion (93), 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

$$\kappa_{s,s'} = e^{\beta(v(s') - v(s))} - 1. \tag{106}$$

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

$$\kappa_{0,1} = e^{\beta(v(1)-v(0))} - 1 = e^0 - 1 = 0 \qquad \text{for all } \beta. \tag{107}$$

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:

$$\begin{aligned} 0 = -\rho A(0) + (1-\gamma)A(0)\Big[r_f + \theta_D(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\Big] \\ + \frac{1-\gamma}{\lambda}A(0)\,g(\theta_D) + \lambda\big[C(1) - A(0)\big] + O(\lambda^{-1}). \end{aligned} \tag{108}$$

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

$$0 = (\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} + O(\lambda^{-1}). \tag{109}$$

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 (109), so that $f(\theta^*_D(0;\beta), \beta) = 0$. Evaluating the partial derivatives at $(\theta^*_{D,0}, 0)$:

$$f_\theta(\theta^*_{D,0}, 0) = -\gamma\sigma^2 + \frac{g''(\theta^*_{D,0})}{\lambda} + O(\lambda^{-1}), \tag{110}$$

$$f_\beta(\theta^*_{D,0}, 0) = O(\lambda^{-1}). \tag{111}$$

Applying the implicit function theorem yields:

$$\left.\frac{d\theta_D^*(0)}{d\beta}\right|_{\beta=0} = -\left.\frac{f_\beta}{f_\theta}\right|_{(\theta_{D,0}^*,0)} = -\frac{O(\lambda^{-1})}{-\gamma\sigma^2 + O(\lambda^{-1})} = O(\lambda^{-1}). \tag{112}$$

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:

$$\eta(0) = -\frac{1}{v(0)} \cdot \frac{1}{\theta_{D,0}^*} \cdot \left.\frac{d\theta_D^*(0)}{d\beta}\right|_{\beta=0} = O(\lambda^{-1}). \tag{113}$$

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 (100) applies here unchanged, giving $\varepsilon_D = O(\lambda)$ as $\lambda \to \infty$. Since $\eta(0) = O(\lambda^{-1})$, we obtain:

$$\frac{\eta(0)}{\varepsilon_D} = \frac{O(\lambda^{-1})}{O(\lambda)} = O(\lambda^{-2}). \tag{114}$$

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

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = \frac{\eta(0)}{\varepsilon_D} = O(\lambda^{-2}) \to 0 \quad \text{as } \lambda \to \infty. \tag{115}$$

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 (44) (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:

$$0 = \upsilon\left[(\mu - r_f) - \sigma^2\theta\right] + \lambda\,\widetilde{C}\,\frac{\kappa_{0,1}}{1 + \theta\,\kappa_{0,1}}. \tag{116}$$

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:

$$q(2|1) = \pi_B, \qquad q(1|1) = 1 - \pi_B, \tag{117}$$

$$q(0|2) = \pi_R, \qquad q(2|2) = 1 - \pi_R, \tag{118}$$

$$q(0|0) = 1. \tag{119}$$

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 (89), the Dynkin expansion (93), 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

$$\kappa_{1,2} = d\beta + O(d\beta^2), \tag{120}$$

$$\kappa_{1,1} = 0. \tag{121}$$

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

$$\lambda \Big[\pi_B C(2)\, \mathbb{E}\big[(1 + \theta_\Delta \kappa_{1,2})^{1-\gamma} \mid \theta_0 = \theta_D\big] + (1 - \pi_B) C(1) - A(1) \Big]. \tag{122}$$

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 (93) for $\mathbb{E}[\theta_\Delta \mid \theta_0 = \theta_D]$, we obtain:

$$\begin{aligned}&\mathbb{E}\left[(1+\theta_\Delta\kappa_{1,2})^{1-\gamma} \mid \theta_0 = \theta_D\right] \\ =&1+(1-\gamma)\Big(\theta_D + \frac{b(\theta_D)}{\lambda}\Big)\kappa_{1,2} + O(d\beta^2) + O(\lambda^{-2}d\beta). \end{aligned} \tag{123}$$

Combining with the diffusive terms (89), the full HJB at state $s=1$ takes the form

$$\begin{aligned} 0 = &-\rho A(1) + (1-\gamma)A(1)\Big[r_f + \theta_D(\mu - r_f) - \tfrac{1}{2}\gamma\sigma^2\theta_D^2\Big] + \frac{1-\gamma}{\lambda}A(1)\,g(\theta_D) \\ &+\lambda\Big[\pi_B C(2)\Big(1+(1-\gamma)\Big(\theta_D + \frac{b(\theta_D)}{\lambda}\Big)\kappa_{1,2}\Big) + (1-\pi_B)C(1) - A(1)\Big] \\ &+O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1}d\beta). \end{aligned} \tag{124}$$

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

$$\begin{aligned} 0 = &(\mu - r_f) - \gamma\sigma^2\theta_D + \frac{g'(\theta_D)}{\lambda} + \frac{\lambda}{A(1)}\,\pi_B C(2)\,\kappa_{1,2}\Big(1+\frac{b'(\theta_D)}{\lambda}\Big) \\ &+O(\lambda^{-1}) + O(\lambda\, d\beta^2) + O(\lambda^{-1}d\beta), \end{aligned} \tag{125}$$

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 (125), the FOC at state $s=1$ becomes

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

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 (126), so that $f(\theta_D^*(1;\beta),\beta) = 0$. Evaluating the partial derivatives at $(\theta_{D,0}^*, 0)$:

$$f_\theta(\theta_{D,0}^*, 0) = -\gamma\sigma^2 + \frac{g''(\theta_{D,0}^*)}{\lambda} + O(\lambda^{-1}), \tag{127}$$

$$f_\beta(\theta_{D,0}^*, 0) = \frac{\lambda}{A(0)} \pi_B C(0) \left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda}\right) + O(\lambda^{-1}). \tag{128}$$

Applying the implicit function theorem yields:

$$\left.\frac{d\theta_D^*(1)}{d\beta}\right|_{\beta=0} = -\left.\frac{f_\beta}{f_\theta}\right|_{(\theta_{D,0}^*,0)} = \frac{\lambda\,\pi_B C(0)}{A(0)\,\gamma\sigma^2}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right) + O(\lambda^{-1}). \tag{129}$$

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

$$\begin{aligned}
\eta(1) &= -\frac{1}{1}\frac{1}{\theta_{D,0}^*}\left.\frac{d\theta_D^*(1)}{d\beta}\right|_{\beta=0} \\
&= -\frac{\lambda\,\pi_B C(0)}{A(0)\,\gamma\sigma^2\,\theta_{D,0}^*}\left(1 + \frac{b'(\theta_{D,0}^*)}{\lambda} + \frac{g''(\theta_{D,0}^*)}{\lambda\,\gamma\sigma^2}\right) + O(\lambda^{-1}).
\end{aligned} \tag{130}$$

Comparison with the price elasticity. As noted previously, in the proof of Proposition 3, when $\beta = 0$, the true price elasticity derived in (100) applies independent of the properties of the Markov process for the state s. Thus, the expression derived in (100) for the price elasticity ε_D applies here unchanged. Dividing $\eta(1)$ by (100), yields

$$\frac{\eta(1)}{\varepsilon_D} = -\pi_B + O(\lambda^{-2}). \tag{131}$$

Recognizing that the estimated elasticity $\hat{\varepsilon}$ recovers the shifter-process elasticity η (by conditions 2 and 3 of Lemma 1), we obtain

$$\frac{\hat{\varepsilon}}{\varepsilon_D} = \frac{\eta(1)}{\varepsilon_D} = -\pi_B + O(\lambda^{-2}) \to -\pi_B \quad \text{as } \lambda \to \infty. \tag{132}$$

□

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

Starting from the market-clearing condition in changes

$$dm_{j,t} = -\frac{dn_{j,t}}{d\beta_j} d\beta_j \tag{133}$$

we obtain the relation

$$x_j P_{j,t} d\beta_j = P_{j,t} dm_{j,t} \cdot \left(-\frac{dn_{j,t}/n_{j,t}}{d\beta_j} \frac{n_{j,t}}{x_j} \right)^{-1}. \tag{134}$$

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

$$\begin{aligned}
-\frac{dn_{j,t}/n_{j,t}}{d\beta_j} &= -\frac{d\left(\frac{\theta^*_{j,t} W_t}{P_{j,t}}\right)/d\beta_j}{n_{j,t}} = -\left(\frac{d\theta^*_{j,t}}{d\beta_j} \frac{W_t}{P_{j,t}} + \frac{dW_t}{d\beta_j} \frac{\theta^*_{j,t}}{P_{j,t}} - \frac{s_t W_t \theta^*_{j,t}}{P_{j,t}} \right) \frac{1}{n_{j,t}} \\
&= -\left(\frac{1}{s_t} \frac{d\theta^*_{j,t}/\theta^*_{j,t}}{d\beta_j} n_{j,t} s_t + P_{j,t} n_{j,t} s_t \frac{\theta^*_{j,t}}{P_{j,t}} - s_t n_{j,t} \right) \frac{1}{n_{j,t}} \\
&= [\eta_{j,t} + (1 - \theta^*_{j,t})] \cdot s_t,
\end{aligned} \tag{135}$$

where we use the relations:

$$\frac{dW}{d\beta_j} = \frac{d(\sum_{j=0}^{J} n_{j,t} \cdot e^{\beta_j s_t} \tilde{P}_{j,t})}{d\beta_j} = n_{j,t} e^{\beta_j s_t} \tilde{P}_{j,t} s_t = P_{j,t} n_{j,t} s_t, \tag{136}$$

$$\frac{dP^{-1}_{j,t}}{d\beta_j} = \frac{d(e^{\beta_j s_t} \tilde{P}_{j,t})^{-1}}{d\beta_j} = -(e^{\beta_j s_t} \tilde{P}_{j,t})^{-2} s_t e^{\beta_j s_t} \tilde{P}_{j,t} = -\frac{s_t}{P_{j,t}}. \tag{137}$$

Combining (135) with market clearing in changes (133) yields the stated result (25) 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:

$$\left\{ \frac{dm_{j,\tau}}{x_j} \right\}_{\tau=t}^{\infty} = \left\{ d\beta_j s_\tau \cdot [\eta_{j,\tau} + (1 - \theta^*_{j,\tau})] \cdot \frac{n_{j,\tau}}{x_j} \right\}_{\tau=t}^{\infty}. \tag{138}$$

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:

$$dm_{i,t} = -\frac{dn_{i,t}}{d\beta_j} d\beta_j. \tag{139}$$

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:

$$\begin{aligned} \frac{dn_{i,t}}{d\beta_j} &= \frac{d}{d\beta_j}\left(\frac{\theta^*_{i,t} W_t}{P_{i,t}}\right) = \frac{1}{P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t}\frac{dW_t}{d\beta_j}\right] + \theta^*_{i,t} W_t \cdot \frac{d(P^{-1}_{i,t})}{d\beta_j} \\ &= \frac{1}{P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t}\frac{dW_t}{d\beta_j}\right] + 0. \end{aligned} \tag{140}$$

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

$$\begin{aligned} \frac{dn_{i,t}/n_{i,t}}{d\beta_j} &= \frac{1}{n_{i,t} P_{i,t}}\left[\frac{d\theta^*_{i,t}}{d\beta_j} W_t + \theta^*_{i,t} P_{j,t} n_{j,t} s_t\right] \\ &= \frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{n_{i,t} P_{i,t}} + \theta^*_{i,t}\frac{P_{j,t} n_{j,t} s_t}{n_{i,t} P_{i,t}}. \end{aligned} \tag{141}$$

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

$$\frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{n_{i,t} P_{i,t}} = \frac{d\theta^*_{i,t}}{d\beta_j}\frac{W_t}{\theta^*_{i,t} W_t} = \frac{1}{\theta^*_{i,t}}\frac{d\theta^*_{i,t}}{d\beta_j} = \frac{1}{s_t}\frac{d\ln\theta^*_{i,t}}{d\beta_j} s_t = -\eta_{ij,t} s_t. \tag{142}$$

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

$$\theta^*_{i,t}\frac{P_{j,t} n_{j,t} s_t}{n_{i,t} P_{i,t}} = \theta^*_{i,t}\frac{\theta^*_{j,t} W_t \cdot s_t}{\theta^*_{i,t} W_t} = \theta^*_{j,t} s_t. \tag{143}$$

Consequently:

$$-\frac{dn_{i,t}/n_{i,t}}{d\beta_j} = [\eta_{ij,t} - \theta^*_{j,t}] s_t. \tag{144}$$

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

$$\frac{dm_{i,t}}{x_i} = -\frac{dn_{i,t}}{d\beta_j} d\beta_j \frac{1}{x_i} = [\eta_{ij,t} - \theta^*_{j,t}] s_t \cdot d\beta_j \cdot \frac{n_{i,t}}{x_i}. \tag{145}$$

This gives the final result

$$\left\{ \frac{dm_{i,\tau}}{x_i} \right\}_{\tau=t}^{\infty} = \left\{ d\beta_j s_\tau \cdot [\eta_{ij,\tau} - \theta^*_{j,\tau}] \cdot \frac{n_{i,\tau}}{x_i} \right\}_{\tau=t}^{\infty}, \tag{146}$$

as stated in (26).

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 (26). For the case $i = 0$ (safe asset), $j = 1$ (risky asset), we have:

$$\frac{dm_{0,t}}{x_0} = d\beta_1 \cdot s_t \cdot [\eta_{01,t} - \theta^*_{1,t}] \cdot \frac{n_{0,t}}{x_0}. \tag{147}$$

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

$$\begin{aligned}
\eta_{01,t} &= -\frac{1}{s_t} \frac{d \ln \theta^*_{0,t}}{d\beta_1} \bigg|_{\beta_1=0} = -\frac{1}{s_t} \frac{d \ln(1 - \theta^*_{1,t})}{d\beta_1} \bigg|_{\beta_1=0} \\
&= -\frac{1}{s_t} \cdot \frac{1}{1 - \theta^*_{1,t}} \cdot \left(-\frac{d\theta^*_{1,t}}{d\beta_1} \right) \bigg|_{\beta_1=0} \\
&= \frac{1}{1 - \theta^*_{1,t}} \cdot \left(-\frac{1}{s_t} \frac{d \ln \theta^*_{1,t}}{d\beta_1} \right) \cdot \theta^*_{1,t} \bigg|_{\beta_1=0} \\
&= -\eta_{1,t} \cdot \frac{\theta^*_{1,t}}{1 - \theta^*_{1,t}}.
\end{aligned} \tag{148}$$

Substituting this into (147) yields:

$$\begin{aligned}\frac{dm_{0,t}}{x_0} &= d\beta_1 \cdot s_t \cdot \left[-\eta_{1,t} \cdot \frac{\theta^*_{1,t}}{1-\theta^*_{1,t}} - \theta^*_{1,t}\right] \cdot \frac{n_{0,t}}{x_0} \\ &= d\beta_1 \cdot s_t \cdot \left[-\frac{\theta^*_{1,t}}{1-\theta^*_{1,t}}\right] \cdot \left[\eta_{1,t} + (1-\theta^*_{1,t})\right] \cdot \frac{n_{0,t}}{x_0} \\ &= -d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1-\theta^*_{1,t})] \cdot \frac{\theta^*_{1,t}}{1-\theta^*_{1,t}} \cdot \frac{n_{0,t}}{x_0}. \end{aligned} \tag{149}$$

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:

$$\begin{aligned}\frac{dm_{0,t}}{x_0} &= -d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1-\theta^*_{1,t})] \cdot \frac{n_{1,t}P_{1,t}}{n_{0,t}P_{0,t}} \cdot \frac{n_{0,t}}{x_0} \\ &= -d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1-\theta^*_{1,t})] \cdot \frac{n_{1,t}}{x_0} \cdot \frac{P_{1,t}}{P_{0,t}}. \end{aligned} \tag{150}$$

Multiplying both sides by x_0:

$$dm_{0,t} = -d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1-\theta^*_{1,t})] \cdot n_{1,t} \cdot \frac{P_{1,t}}{P_{0,t}}. \tag{151}$$

But we also know from the result (25) for asset $j = 1$ that:

$$dm_{1,t} = d\beta_1 \cdot s_t \cdot [\eta_{1,t} + (1-\theta^*_{1,t})] \cdot n_{1,t}. \tag{152}$$

Therefore:

$$dm_{0,t} = -dm_{1,t} \cdot \frac{P_{1,t}}{P_{0,t}}, \tag{153}$$

or in process form:

$$\left\{dm_{0,\tau}\right\}_{\tau=t}^{\infty} = \left\{-dm_{1,\tau} \cdot \frac{P_{1,\tau}}{P_{0,\tau}}\right\}_{\tau=t}^{\infty}. \tag{154}$$

A.6 Proof of Corollary 1

The assumptions in Corollary 1 are the same as those in Proposition 2, except that we consider continuous trade. Further, equation (25) implies the following relation for the incremental holding upon inception of a dynamic the program:

$$dm_{1,t} = s_t d\beta_1 \cdot [\eta_{1,t} + (1 - \theta^*_{1,t})] \cdot n_{1,t}, \tag{155}$$

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 $s \in \{-1, +1\}$:

$$\eta_C(s) = \pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)}. \tag{156}$$

Moreover, given the considered marginal exposure we can write:

$$\theta^*_C(s) = \theta^*_C(0) + \theta_1 s_t d\beta_1, \tag{157}$$

for some constant θ_1. Plugging these results into (155) yields

$$dm_{1,t} = s_t d\beta_1 \cdot [\pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)} + (1 - (\theta^*(0) + \theta_1 s_t d\beta_1))] \cdot n_{1,t}, \tag{158}$$

$$= s_t d\beta_1 \cdot [\pi_R \cdot \frac{\lambda \cdot C(0)}{(\mu - r_f) \cdot A(0)} + (1 - \theta^*(0))] \cdot n_{1,t}. \tag{159}$$

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, m_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, $\Delta m_j = -(n_j^{+} - n_j^{-})$, implying

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

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 sum of the past 12 months of dividend payouts. 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

The first step is to construct the portfolio weights of institutional money managers. 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. The crosswalk between CUSIP and PERMNO provided by CRSP is used. Information on stock prices (PRC) and shares outstanding (SHROUT) at quarter-end as well as share codes (SHRCD) are obtained from CRSP Monthly Stock File. Stock characteristics are from the ”IV.dta” file presented earlier. Positions whose share codes are not 10, 11, 12 or 18 are discarded. 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.

The ”household sector” ($MGRNO = 0$, ”HH”) is then defined following KY19 as holding residual shares outstanding (if any) for all securities in the universe. Likewise, following KY19, all positions of a given MGRNO at a given FDATE ina stock with share code 12, 18 or with missing return or characteristics are aggregated and labelled as the ”outside asset” ($PERMNO = 0$). If a given MGRNO has assets under management (AUM) below 10M at any given time and or if it has no holding in the outside asset, its holdings are rebated to the HH sector. Once these modifications are made, AUMs are recomputed portfolio weights are computed. For any row, RWEIGHT is the ratio of that position's weight in the corresponding MGRNO's portoflio (at the FDATE considered) and the weight of the outside asset. Based on data made publicly available by KY19 for 2007Q4, it is possible to compare the proximity between the relative portfolio weights generated by this procedure and those in their file. Over 2007Q4, the median absolute error $\frac{|rweight - rweight_{KY}|}{rweight_{KY}}$ is 0.0005.

For any given MGRNO/FDATE, rows with zero values for RWEIGHT 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. Once the investment universe of each institution is defined, it is possible to construct $IVme_{it}(n)$ for a stock n and an institution i at time t as in KY19, formula (19). It is possible to compare the generated values of IVme for the HH sector to the values retrieved from the file IV.dta. Over the sample period, the median absolute error $\frac{|IVme - IVme_{KY}|}{IVme_{KY}}$ is 0.04.

Finally, only MGRNO/FDATE pairs with NHOLDING of at least 1000 are retained for the following analysis.