Inelastic Market Hypothesis

A new way to explain market fluctuations

Paper realizzato da:
Besliu Remus
Cardoni Chiara
Favazzo Luca Maria
Ferrero Giacomo
Furlan Edoardo
Mariotti Andrea
Urbano Pierluigi
Trubiano Maia
Vinci Dario

Indice

1 Introduction

This paper aims to explain the high volatility of the stock market as a consequence of its surprising inelasticity—what is referred to as the “Inelastic Markets Hypothesis”—providing a quantification of the market’s aggregate elasticity through a new instrumental approach: the “granular instrumental variables” (GIV) approach.

1.1 Inelastic Market Hypothesis

According to the simplest "Efficient Markets" model, if an investor sells one dollar worth of bonds and buys one dollar worth of stocks, the overall valuation of the aggregate stock market should remain unchanged, since prices are assumed to reflect the present value of future dividends. However, both theoretical analysis and empirical evidence show that such a trade leads to a 5% increase in the stock market’s value. This observation implies that the stock market is surprisingly inelastic: the aggregate price-elasticity of demand for stocks is low (a 5% increase in price results in only a 1% decrease in demand, suggesting a price elasticity of just 0.2). To explain this market inelasticity, Gabaix and Koijen propose a simple model: a representative investor can allocate wealth between a pure bond fund and a mixed fund that follows a fixed allocation rule—for example, 80% equities and 20% bonds. When the investor sells $1 of the pure bond fund and reallocates it to the mixed fund, the latter is required to invest 80 cents in stocks. This pushes up stock prices, which increases the value of the mixed fund’s equity holdings, prompting it to buy even more stocks to maintain its 80% target. This recursive mechanism continues until a new equilibrium is reached, in which a $1 flow into equities results in a $5 increase in total market value. This amplification effect highlights how capital flows and demand shocks can have a large and persistent impact on asset prices and expected returns in an inelastic market environment. Through the paper richer setups are explored, showing that the ramifications of this simple model are solid. For instance, in the case the fund is more actively contrarian (it buys more when the expected excess on equities is high). Moreover the model aggregates well (if different investors have different elasticities, the total market elasticity is the size-weighted elasticity of the market participants, where the size is share of equity they hold), it clarifies how to measure net flows into the aggregate stock market and extends to an infinite horizon (price today is determined by the cumulative inflows to date and the present value of future expected flows divided by the market elasticity).

1.2 Granular Instrumental Variables Approach

To quantify market-wide elasticity, the paper introduces a new method—the Granular Instrumental Variable (GIV) approach—which uses idiosyncratic demand shocks from large institutions or sectors as exogenous variation. These shocks, identified via factor models on holdings data, are aggregated by size to form the GIV. This allows the authors to estimate the market’s sensitivity to demand shocks (multiplier ≈ 5) and investor demand elasticity (≈ 0.2). The findings suggest a lasting price impact of flows, with robust checks confirming the multiplier and showing strong links to returns, but weak ties to macroeconomic growth.

1.3 A Priori Reasons Why the Markets Would Be Inelastic

Gabaix and Koijen model can be further accompanied by 3 “a priori” reasons why the markets would be inelastic:

1. Institutional constraints and lack of macro arbitrageurs: Many investment funds operate under strict mandates and cannot freely adjust their portfolios. Moreover, there are few investors capable of acting as macro-level arbitrageurs (Hedge funds, for instance, hold less than 5% of the equity market).
2. Limited risk transfer across investor sectors: The transfer of equity risk between investor groups is minimal (about 0.6% of total market value per quarter on average), suggesting that most investors have low demand elasticity or experience similar demand shocks.
3. Empirical evidence of low aggregate elasticity: While the demand elasticity for individual stocks (“micro” elasticity) is estimated to be around 1, the aggregate or “macro” elasticity should logically be lower, since broad stock indices and bonds are less substitutable than, say, two companies in the same industry. Supporting this, recent studies on factor demand (e.g., size, value) find elasticities around 0.2—consistent with the inelastic markets hypothesis.

1.4 Data and Facts on Equity Shares and Flows

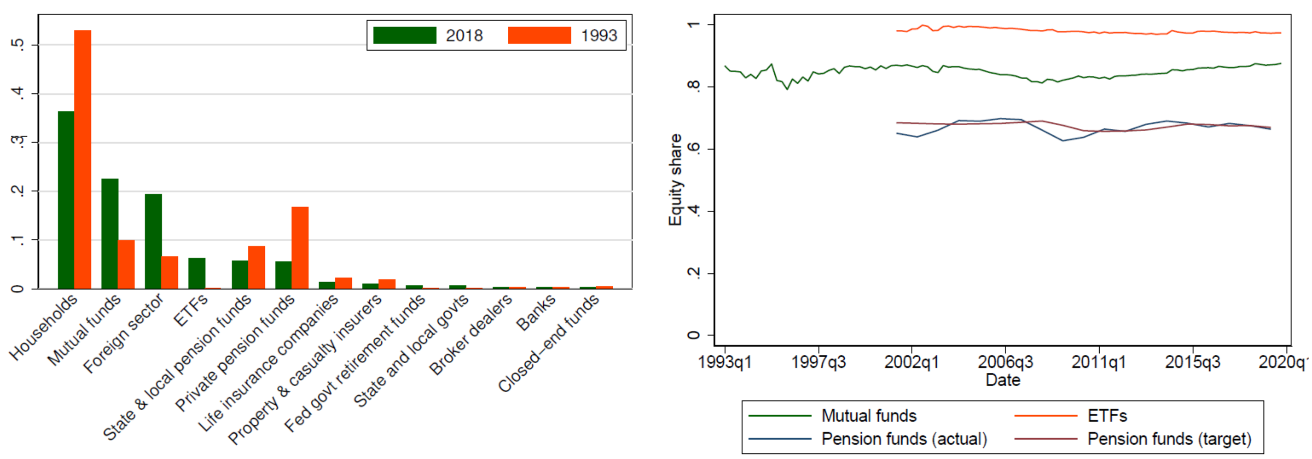

Figure 1: Equity shares of institutional investors in 1993 and 2018

The graph on the left shows the equity share of institutional sector in 1993 and 2018, with orange indicating the value in 1993 and green in 2018. On the right you can see the value-weighted average equity share of several institutional investors (the ones considered are mutual funds, ETFs, state and local pension funds). This analysis uses investor return expectations from Gallup surveys, following Greenwood and Sheleifer(2014). Data go from Q4 of 1996 to Q4 2018. As we can see, equity ownership in the last 25 years went from households to institutional investors. For mutual funds, ETFs, and pension funds, the authors use investor-level data on equity and fixed income holdings. The authors show the equity share, putting together investors based on the size of their equity portfolios rather than total assets under management. This distinction matters: if two funds hold equal total assets, but one is 100% in stocks and the other fully in bonds, the equity-weighted share would be 100%, while the asset-weighted share would be just 50%. If the percentage of equity or bond changes , the asset-weighted share changes, while the equity-weighted share remains at the same level . It is interesting to show that equity shares remain stable for major investor groups. This is mainly due to the fact that many institutional investors have constraints to keep a fixed equity allocation. The study shows that equity allocations for institutional investors remain quite stable over time. Since foreign investors are made up of similar types of institutions, this raises the question of who actually adjusts their portfolios to exploit pricing differences across asset classes. Broker-dealers hold less than 0.5% of the equity market, basically too small to manage large market shifts over time. Hedge funds are also relatively small, with less than 4% of the market in long positions before the financial crisis. Ben-David et al. (2012) shows that hedge funds sold a significant portion of their equity during the 2008 crisis, but due to their size, this only amounted to about 0.1% of the market per quarter. The drop was largely driven by investor redemptions and leverage constraints(roughly 80% of the change). The authors extend this analysis to other major downturns, like the tech crash and the 2008 financial crisis, and still find that average inter-sector flows rarely exceed 0.5% of the market. If demand were highly elastic, we'd expect large flows and big shifts in holdings when investors disagree. But since flows are small, elastic-demand models would require that investors hold nearly identical beliefs and risk preferences. In contrast, inelastic-demand models allow for constant disagreement and more variation in sentiment something that aligns better with the data. A central point in the paper is the distinction between micro and macro elasticity. At the micro level (e.g., between two similar stocks), demand is relatively elastic: if you buy Apple and sell Microsoft, relative prices adjust moderately. But at the macro level (e.g., between stocks and bonds), demand is much more rigid.

Table 1: Multipliers from the Literature

Panel A: Micro multiplier		
Study	**Methodology**	**Multiplier**
Chang, Hong and Liskovich (2014)	Index inclusion	0.7 to 2.5
Pavlova and Sikorskaya (2020)	Index inclusion	1.5
Schmickler (2020)	Dividend payouts	0.8
Frazzini et al. (2018), Bouchaud et al. (2018)	Trade-level permanent price impact	15
Panel B: Factor-level multiplier		
Ben-David, Li, Rossi and Song (2020a)	Morningstar ratings change	5.3
Peng and Wang (2021)	Fund flows	4.8
Li (2021)	Fund flows + SVAR	5.7
Panel C: Macro multiplier		
Da, Larrain, Sialm and Tessada (2018)	Pension fund rebalancing Chile	2.2
Li, Pearson and Zhang (2020b)	IPO restrictions in China	2.6–6.5

Table 1 summarizes elasticity estimates from the literature. At the micro level, the multiplier (i.e., the price impact per percentage point bought) is around 1. At the macro level, observed multipliers in Chile, China, and the U.S. range between 2.2 and 6.5 — consistent with the idea that $1 of buying can increase market value by $5. If the **inelastic markets hypothesis** is true, then **investor flows and demand shocks** have a direct and measurable impact on prices, replacing the "dark forces" of traditional asset pricing models with observable data. This might allow the studying of the concrete behavior of investors to better understand how prices are formed. Firms also influence the market more than previously thought **stock buybacks**, which are neutral in classic models, can raise market value—by about $2.2 for every $1 repurchased. Given their growing role, corporate actions are a major source of market volatility and growth. The authors propose a **general equilibrium model** focused on flows and inelasticity, able to replicate real-world market features such as volatility, the equity premium, and the slow mean reversion of the price-dividend ratio. The model also links these flows to the **stochastic discount factor (SDF)**, offering a more realistic foundation for macro-financial modeling.

2 Models

In examining the origins of financial fluctuations within the framework of the Inelastic Markets Hypothesis, we develop a series of models that underscore the disproportionate impact of cash flows on asset prices. Rather than attributing price changes solely to shifts in fundamentals such as dividends or earnings, the model emphasizes that the structure of market demand—in particular, its inelasticity—amplifies even small net equity flows. In simple terms, the market's limited capacity to absorb trades leads to a *multiplier effect*, whereby modest cash inflows result in substantial price adjustments.

2.1 The Two-Period Model and the Multiplier Effect

The core of the analysis is represented by the equation

$$p = \frac{f}{\zeta},$$

where:

- p represents the percentage change in the asset's price;
- f denotes the net equity flow into the market, measured as a percentage of the total market capitalization;
- ζ is the macro elasticity of demand for equities, indicating the sensitivity of the aggregate demand to price changes.

This relationship implies that if the market demand is highly inelastic (i.e., ζ is low), even a small positive cash flow f will induce a large percentage change in price. For instance, when $\zeta = 0.2$, the multiplier $1/\zeta = 5$ suggests that a 1% inflow will correspond to a 5% price increase.

2.2 Microeconomic Derivation

The model further disaggregates the analysis by considering the demand of individual funds. For a particular fund i, the deviation in the demand for equities is modeled as

$$q_i = -\zeta_i\, p + \kappa_i\, \delta\, d + f_i$$

where:

- q_i is the deviation in demand for equities for fund i;
- ζ_i measures the individual elasticity of fund i's demand to price deviations;
- p denotes the deviation in the asset price from its equilibrium;
- d represents the deviation in expected dividends, capturing a fundamental input;
- δ is the price, so $\Delta d/P$ is the dividend–price ratio;
- κ_i captures the sensitivity of the fund to excess returns;
- f_i denotes the fund-specific cash flow.

When aggregated across all funds, the deviation in demand becomes:

$$q = -\zeta p + \kappa \delta d + f$$

where ζ, κ, and f now represent the aggregate elasticities and cash flows, respectively. Assuming the aggregate supply of shares is fixed (i.e., $q = 0$), equilibrium forces the price to adjust so that

$$p = \frac{f + \kappa\, \delta\, d}{\zeta}$$

- **Flows-driven Pricing**. If dividend deviations are absent (i.e., $d = 0$), the equation simplifies to:
$$p = \frac{f}{\zeta}$$
- **Fundamentals-driven Pricing**. In the absence of cash flows (i.e., $f = 0$), the price deviation is driven by dividend deviations:
$$p = \frac{\kappa \delta d}{\zeta}$$

Since κ/ζ is typically small—owing to fund inertia or mandate rigidity—the impact of changes in expected dividends is often muted compared to that of cash flows.

2.3 The Infinite-Horizon Model: Incorporating Dynamics

The analysis is extended to an infinite-horizon framework by modeling the equilibrium asset price at time t as the discounted sum of future net equity inflows and dividend deviations:

$$p_t = \sum_{\tau} \left(\frac{\rho}{\zeta} f_\tau + \frac{\delta}{\zeta} d_\tau \right)$$

where:

- p_t represents the price at time t;
- f_τ and d_τ denote the net equity flows and dividend deviations at a future time period τ, respectively;
- ρ is a discount factor that accounts for the present value of future cash flows.

This formulation emphasizes that asset prices are determined not only by present market conditions but also by the anticipation of future flows. Consequently, even persistent but modest cash inflows can have lasting effects on asset prices.

2.4 Implications and Interpretation

The model challenges traditional asset pricing theories by highlighting the significance of liquidity and inelastic demand. The multiplier effect, encapsulated by:

$$p = \frac{f}{\zeta}$$

illustrates that markets with low demand elasticity amplify the impact of cash flows—leading modest inflows to produce disproportionately large price movements. Furthermore, the model suggests that:

- **Dominance of Flows:** Uninformed or mechanical trading flows (such as those due to institutional rebalancing or pension fund contributions) can drive asset prices more significantly than fundamental factors like dividends.
- **Feedback Mechanisms:** An initial cash inflow not only directly raises prices but also induces portfolio rebalancing, resulting in a self-amplifying loop that reinforces the initial price increase.

Thus, the Inelastic Markets Hypothesis offers a robust explanation for the persistent volatility observed in equity markets by demonstrating that liquidity constraints and cash flows can be as influential—or even more so—than traditional economic fundamentals.

3 Empirical Estimation of Aggregate Market Elasticity

The authors posits that aggregate stock prices are highly sensitive to capital flows due to a low elasticity of demand. The authors introduce and implement a novel econometric strategy to estimate the price elasticity of demand in the aggregate stock market. The main finding is that this elasticity is low—approximately $\zeta = 0.2$—implying that an additional dollar invested in the stock market increases total market capitalization by around five dollars. This magnitude of price impact, which significantly exceeds predictions from standard asset pricing theory, fundamentally reshapes our understanding of market fluctuations.

The primary challenge in estimating the causal effect of flows on prices lies in the endogeneity between investor demand and price movements. To address this, the authors propose the use of a granular instrumental variable (GIV) approach, which exploits the idiosyncratic demand shocks of large institutional investors. These shocks, extracted as residuals from a factor model, represent variations in equity demand that are unrelated to market-wide fundamentals or common risk factors. Crucially, these idiosyncratic components are aggregated into an instrument using a size-weighted average, where the weights reflect each institution's share of total equity holdings.

Formally, if Δq_{it} denotes the change in equity demand by institution i at time t, it is modeled as the sum of systematic components driven by observable factors and an institution-specific residual ν_{it}. The granular instrument is then defined as:

$$\text{GIV}_t = \sum_i s_i \nu_{it},$$

where s_i is the share of institution i's equity holdings relative to the aggregate market. This GIV is used as an instrument in a two-stage least squares (2SLS) estimation framework to correct for the endogeneity of flows.

In the first stage, aggregate capital flows F_t are regressed on the instrument:

$$F_t = \alpha + \gamma \cdot \text{GIV}_t + \varepsilon_t,$$

yielding predicted flows $\hat{F}_t$ that are orthogonal to return expectations and macro fundamentals. In the second stage, the change in the logarithm of the aggregate price index is regressed on these predicted flows:

$$\Delta \log P_t = \beta \cdot \hat{F}_t + \eta_t,$$

where β represents the inverse of the aggregate elasticity, $\beta = 1/\zeta$.

The estimated coefficient implies a low elasticity of demand, with $\zeta \approx 0.2$. This estimate is interpreted as a macro-level response: when investors allocate capital into equities equivalent to 1% of the total market value, the price level increases by approximately 5%. Conversely, withdrawals of a similar magnitude cause a symmetric decline. These results imply that flows have first-order effects on prices, contrary to the near-zero multipliers predicted by traditional models with infinitely elastic demand.

An important feature of this price impact is its persistence. The authors find that flow-induced changes in price levels are not quickly reversed, suggesting that these effects are not merely temporary liquidity shocks but rather reflect new equilibrium valuations in an inelastic setting. This finding is consistent with the theoretical framework presented in earlier chapters, where institutions face rigid allocation mandates and limited flexibility in adjusting portfolio weights, leading to persistent deviations in aggregate demand.

The robustness of the elasticity estimate is confirmed using a variety of datasets, including Flow of Funds data, mutual fund and ETF flows, and 13F filings from institutional investors. Moreover, the results are stable across different time periods, including crisis episodes and boom years, which adds credibility to the claim that inelasticity is a structural feature of the market rather than a cyclical anomaly. Even when disaggregated across investor types, most institutions exhibit similarly low elasticities. This supports the view that institutional frictions—such as fixed policy mandates, rebalancing rules, and benchmark tracking—are pervasive.

Interestingly, the estimated flows are only weakly correlated with macroeconomic fundamentals like GDP or earnings growth. In contrast, they exhibit a strong relationship with investor expectations, as proxied by survey data and forward-looking sentiment indicators. This disconnect further reinforces the inelastic markets hypothesis: flows and demand shocks, rather than changes in fundamental value, appear to be the proximate drivers of asset price fluctuations. Prices, in this framework, react not because new information about cash flows arrives, but because supply and demand imbalances arise in a market structure where few participants are willing or able to absorb shocks.

The implications of this finding are profound. In standard theory, prices should incorporate new information about fundamentals, and the influence of uninformed flows should be arbitraged away by rational investors. However, if markets are inelastic, even uninformed or behaviorally-driven flows can permanently shift prices. This challenges the Modigliani-Miller irrelevance theorem in the aggregate and implies that policies such as stock buybacks or public equity purchases can have substantial real effects.

The empirical results also motivate a deeper investigation into the determinants of flows themselves. Possible drivers include regulatory constraints, institutional investor mandates, behavioral biases, or liquidity constraints. Understanding these origins could enable a fully microfounded theory of price formation in inelastic markets.

4 Policy and Corporate Implications

4.1 How Inelastic Market Reshape Governament Interventions

Typical QE stresses mostly bonds that are long-term, but the IMH thinks buying equity by central banks could affect prices of assets more. Market inelasticity suggests a potential $5 rise in market capitalization from $1 of equities, a greater multiplier than bond purchases show. Important examples of this are: Hong Kong (1998) the government equity purchases ($\sim$ 6% of market cap) led to a 24% surge in stock prices; and Japan (2010s) the BOJ overall ETF program (about 5% of the market) resulted in broad valuation increases without broad market reversals. This model as a clear trade-off as it lowers the cost of capital for firms stimulating investments, but at the same time it creates risks of asset bubbles and inflating asset prices, benefiting wealthier households. Under the traditional view, fiscal stimulus impacts aggregate demand via consumption channels (e.g., tax cuts, unemployment benefits). IMH suggests that policies increasing equity demand (e.g., tax incentives for stock ownership) may have direct, amplified price effects, generating substantial wealth effects beyond mere consumption increases. These wealth effects may extend through flow-induced asset price booms, further deepening disparities between asset-rich and asset-poor households. Having now a look at the implications of the IMH on financial stability and macroprudential policy, we can see that this model warns that flow reversals (e.g., mass selloffs) can trigger severe market disruptions, as witnessed during the 2008 global financial crisis. Situations like this require new macroprudential policies capable of addressing the inelastic nature of financial markets. The main policy tools we dispose of are:

- **Circuit Breakers**: Intended to prevent panic selling, these may inadvertently worsen illiquidity in the face of persistent flow shocks.
- **Capital Buffers**: Regulations requiring institutions to hold more liquid assets may help absorb shocks stemming from flow imbalances.

4.2 The Influence of Firms on Market Prices

Regarding share buybacks and equity issuance, traditional views (such as the Modigliani-Miller theorem) assert that corporate payout decisions do not affect firm value. However, the IMH posits that in an inelastic market, buybacks reduce equity supply while demand remains rigid, thus raising stock prices. Through this new perspective companies can leverage buybacks to increase their stock prices during times of low flow, potentially timing the market to their advantage. This dynamic helps explain the surge in buybacks seen in recent decades, which plays a role in both market volatility and the overall valuation of equities. Given the inelasticity of equity markets, firms may find it more costly to issue equity during periods of high prices, incentivizing them to take on more debt instead. This creates a feedback loop: higher debt levels increase financial fragility, but inelasticity discourages firms from deleveraging. This paradox leads to a system where equity issuance remains costly, especially when market prices are high. In an inelastic market, firms can exploit overpriced equity to fuel stock-financed acquisitions. This is particularly relevant in stock market bubbles, where inflated stock prices facilitate acquisitions even when the fundamental value may not justify them, as seen in the 1999–2000 tech bubble.

5 Microstructural Interpretation

Now we will move from GK's macro-finance perspective towards the microstructural approach analyzed by Bouchaud. In particular, we will consider the application of the Latent Liquidity Theory to GK's findings. Traditional economists believe that the price of an asset tends to represent the fundamental value of the asset itself; the assumption here is that prices are martingales:

$$P_t = \mathbb{E}[P_{t+1}|\mathcal{F}_t] \tag{1}$$

Where P is the price at a certain time t and $\mathcal{F}$ is the information set. This means that the price today is equal to the price we expect to see in the future given our current knowledge of the market. A discount rate can be added to this model, such that the price today is equal to the expected price discounted by a certain amount per period. The noise in the price of assets, which we observe empirically, is the result of uninformed trades that should affect the price only for a brief period of time. Real price movements only occur when new information gets added to the current information set, which is when the perceived fundamental value of the asset changes. The process that leads from the old price to the new price is called price-discovery and should happen instantaneously. This model is the Efficient Market Hypothesis (EMH), which is the most affirmed and academically recognized theory. According to EMH, the microstructure of the markets should be irrelevant in the long term.

5.1 Orden-driven markets

However, empirical evidence shows that there are at least three problems with this the EMH:

- Excess volatility problem (Shiller, 1980): the volatility of assets' prices exceeds the volatility that can be justified by just new information and changes in fundamental value.
- Excess trading problem: investors tend to trade too frequently without any apparent reason, leading to sub-optimal strategies.
- Trend-following problem: we can observe long-lasting trends in the prices of assets that cannot be predicted through the EMH.

Models, like those proposed by Kyle in 1985 and Glosten & Milgrom also in 1985, try to assess these problems of the EMH by adding noise traders, who can actually create real price movements and who are considered to be absolutely necessary for the markets to behave efficiently. However, these models do not explain the underlying reason real prices differ from fundamental value. To try to give a better understanding of the way prices behave, another possible approach is the theory of order-driven markets. This new paradigm views price changes as a result of the sequence of orders that investors send to the exchange, rather than fundamental value shifts. This view is supported by the fact that, empirically, we see that endogenous factors account for more than 80% of total market volatility. The consequence of this new theory, which seems to be supported by empirical data, appears obvious: every kind of trade can affect prices in the long run, even uninformed ones (Gabaix & Koijen, 2021). In contrast with the EMH, here changes in prices are not the result of a price discovery phenomenon, but they are merely the result of a process of price formation, influenced primarily by supply and demand.

5.2 Square-root law

We should now pay attention to metaorders, which are big market orders divided into many smaller operations. Empirically, we observe that the impact on prices of metaorders follows the square-root law, which means that the impact does not grow linearly with volume of the metaorder, but with the square root of its volume. The formula to describe this behavior is as follows:

$$I(Q,T) \approx Y\sigma_T\sqrt{\frac{Q}{V_T}} \tag{2}$$

In the formula, $I(Q,T)$ represents the average price impact of a metaorder of size Q executed over a time period T. The coefficient Y is a dimensionless constant, typically around 0.5 in U.S. equity markets, and reflects how sensitive the market is to trading activity. The term σ denotes the volatility of the asset over the execution time T, capturing how much the price typically fluctuates during that period. For this law to be applicable, Q must be consistently smaller than the volume V.

It is relevant to observe that Q is divided only by the volume and not by the market capitalization of the asset, meaning that the impact that the order has on the price is far greater than what a traditional model would predict. The variation of the price that results from a 1$ investment is indicated with the letter M. We empirically observe that a M is equal to 1 if we invest 1$ in a single stock and to 5 if we invest in an index.

5.3 Latent Liquidity thory

The first theories which attempted to explain the square-root law appeared during the mid-nineties and, finally, in 2011 the Latent Liquidity Theory (LLT) was proposed: it is based on a dynamical description of supply and demand, providing a natural statistical interpretation for the square-root law and further predictions, in particular concerning the decay of the impact of the metaorder, which is necessary to recover GK's multiplier at long times. Essentially, the theory assumes that each long-term investor in the market has a reservation price (to buy or to sell) that he or she updates as a function of time, due to incoming news, price changes, noise, etc. The collection of all these trading intentions constitutes the available liquidity at any instant of time—although most of it remains 'latent', i.e. is not immediately posted in the public order book. When the market price hits the reservation price of a given buy (sell) investor, his/her order is executed and becomes a sell (buy) intention, but at a price significantly higher (lower) than the execution price. Reservation prices remain sticky during a typical memory time Tm and, when revised, tend to distribute themselves around the new market price: this means that investors tend to mix their own price estimate with the market price, which represents what other investors believe. This seems reasonable, following Black's intuition that no one can really claim to know better than others, so a realignment of beliefs seems natural. Intuitively, on short time scales T<Tm, liquidity is essentially static, preventing price motion. But as belief realignment happens, memory of previous intentions is erased and the impact of past trades on the price becomes permanent. In general:

1. The impact of a metaorder of volume Q executed over some time $T < T_m$ is given by the square-root law defined by equation 1.
2. Once the metaorder is fully executed, its impact decreases with time from its peak value, determined by equation 1, down to a value that depends on the information content that triggered the trade.
3. In the absence of any information, the long-term and permanent part of the impact is given by the equation:

$$I_\infty(Q,T) = \frac{1}{2}\sigma_1 T_m \frac{Q}{V_1 T_m} \tag{3}$$

Where the subscript '1' corresponds to 1 day and Tm is measured in days. Therefore, this result is independent of the execution time T, is linear in Q, is directly proportional to the volatility on the scale of the memory time Tm and inversely proportional to the volume traded by the market on the same time scale V1Tm. In other words, it means that, while the square-root law describes the short-term, in the long-term remains a smaller component, which is linear in Q. The longer market participants stick to their beliefs, the stronger the anchor to a reference price and the smaller the long-term impact. If liquidity was infinite (i.e. $V1 = \infty$), or if beliefs were permanent (i.e. $T_m = \infty$), then markets would be perfectly elastic and uninformed transactions would not impact prices.

5.4 GK's multiplier

Now we can compare this new equation with GK's multiplier M (we could not compare equation 1, because it is not linear in Q and it predicts only a transient effect). Assuming $Q = 1\%M$ (i.e. a total order size of 1% of the market capitalization, but any other number would do), the long-term impact on the price expressed in percent is GK's multiplier, hence:

$$\frac{1}{2}\sigma_1 \frac{M}{V_1}\sqrt{\frac{1}{T_m}} \tag{4}$$

The precise value of the memory time is difficult to pin down, especially since we expect market participants to be characterized by a broad distribution of frequencies. But it is plausible that an effective value of Tm is between a few days and a few weeks.

5.5 Discussion

One could expect that large capitalization stocks, for which volatility is smaller and the fraction of market caps exchanged daily is larger, have a relatively smaller multiplier, which means that larger capitalization stocks are less 'inelastic'. But another plausible specification is that the memory time Tm is such that the volatility over that time scale reaches a certain universal threshold value, meaning that investors tend to realign their beliefs when the price level has changed 'appreciably', say when = 10% (corresponding to Tm =16 days for a 2.5% daily volatility). This is tantamount to setting $Tm = \frac{\Delta^2}{\sigma_1^2}$, which finally leads to:

$$M \propto \frac{\sigma_1^2 M}{V_1} \tag{5}$$

The overall range of variation predicted by equation (4) turns out to be rather restricted. The multiplier M is seen to decrease by a factor roughly equal to 2 between small cap stocks ($M \sim 100$M \$) and large cap stocks ($M \sim 1000$B \$).

High-frequency traders and market makers only seek to exploit short-term statistical arbitrage opportunities, without any long-term view about fundamental value. Therefore, their activity makes prices unpredictable and propagates the high-frequency value of volatility to long time scales. The resulting volatility has no reason to match the fundamental volatility. Hence, one plausible explanation for the excess-volatility puzzle is that the trading-induced volatility is much larger than the fundamental volatility.

It is only over very long time scales (> 5 years) that some mean reversion around the fundamental value can be observed. This means that GK's multiplier M for long periods is significantly smaller than the one measured on monthly time scales. Unfortunately, on a long term it is very hard to measure.

Consequently, no-arbitrage at high frequencies is secured by HFT/market-making activities. But since these activities are highly competitive, one expects that the average profitability of liquidity provision is in fact close to zero. This condition is enough to enforce that spread and volatility are related. The conclusion is that order flow, whether informed or non-informed, is probably the major source of volatility in financial markets.

As said before, the IMH is a very controversial topic, since it tries to offer an alternative to the most widely spread theory of market efficiency. For this reason, we now present a summary of the main criticisms that could be made to this model:

1. Limited scope of application, since the hypotesis requires strong assumptions and works only under specific circumstances (e.g. it focuses on institutional flows, underestimating the effect of retail and algorithmic trade).
2. It cannot fully explain why the price distorsion remains despite the presence of arbitrageurs, because they should be able to predict predictable patterns and restore elasticity.
3. The presence of a large amount of historical data supporting market elasticity.
4. It focuses on demand shocks but does not consider how supply may respond to price changes, reducing the inelasticity.

6 Regulatory Evidence from Trusts

Trusts are legal structures where assets are managed by a trustee on behalf of beneficiaries. Trusts are governed by fiduciary law, with a central duty being the duty of prudence, that is, investing cautiously and responsibly, considering the limited financial knowledge the beneficiaries may have. Under the traditional Prudent Man Rule, trustees were required to ensure that every asset held in the trust was individually safe and prudent. This meant avoiding highly volatile and speculative investments, as legal liability could result from holding even a single asset that might later be deemed inappropriate. As a result, trust portfolios were heavily tilted toward stable, court-defensible assets such as mature, dividend-paying firms. However, this regulatory framework was fundamentally redirected with the introduction of the Uniform Prudent Investor Act (UPIA), adopted in a staggered manner across U.S. states between 1986 and 2006. Under UPIA, prudence is evaluated based on the entire portfolio. This reform incorporated principles of modern portfolio theory, emphasizing diversification and allowing trustees to hold riskier and less traditional assets, as long as their whole portfolio remained prudent. We therefore can generate three central predictions:

1. Before UPIA, trusts will exhibit a tilt toward prudent stocks, significantly more than other institutional investors.
2. After UPIA, trusts will reallocate away from safer stocks in the direction of riskier ones, particularly those that improve portfolio diversification.
3. If markets are inelastic, these reallocation patterns will impact prices, leading more

cautious stock prices to fall and less safe ones to rise.

This regulatory change applies only to trusts, and not mutual funds, hedge funds, or pension funds (which are regulated under the federal ERISA law). Nevertheless, the significant market presence of trusts, holding between 14% and 26% of total institutional equity, ensures that their collective rebalancing has the potential to meaningfully influence asset prices.

6.1 Data

The authors collected data on the different dates of UPIA implementation between 1986 and 2006, which is necessary to assess the law's influence on the stock market. They then gathered equity positions of all institutional investors managing over $100 million in the US and used a proprietary classification scheme to specifically identify trusts. Trusts are particularly interesting to study because they are subject to stringent legal fiduciary duties that require them to invest in a cautious and legally defendable way, making it easier to see how changes in fiduciary law lead to shifts in market demand and asset prices—thus supporting the inelastic market hypothesis. To measure the "prudence" of equity investments, a prudence index (ψ) was created based on CAPM beta, firm age, dividend yield, profitability, stock return volatility, and S&P 1,500 membership; the index ranges from -6 (most prudent) to 6 (least prudent) and is built using only the information available to investors at decision time.

6.2 Results on Holdings

Using the prudence index (ψ), the authors show that one year before UPIA, trusts held a significantly larger portion of "prudent" equity compared to other institutional investors. They compare trusts to non-trust institutions within the same State over time using a difference-in-difference approach. Splitting stocks into four categories based on ψ, they show that after UPIA, trusts decreased their allocation to the most prudent stocks ($\psi \leq -3$) by about 4.0% relative to non-trusts, while increasing their holdings in stocks with moderate prudence levels ($-2 \leq \psi \leq 0$) by approximately 3.2%—implying that roughly 80% of the funds withdrawn from the most prudent stocks were reallocated into this less prudent category.

Moreover, when comparing portfolio weight differences before and after UPIA, the gap between trusts and non-trusts essentially flattens out, suggesting that the binding constraints of the old "prudent man" rule were the primary driver of these differences and that UPIA effectively eliminated them.

6.3 Stock Return around the Introduction of the UPIA

The authors observe that UPIA causes a significant shift of money from prudent stocks into less prudent ones. If equity markets were perfectly elastic, such demand changes wouldn't affect asset prices; however, consistent with the inelastic market hypothesis, prices of prudent stocks decline while those of less prudent stocks increase after UPIA. By maintaining the four-category classification based on ψ and constructing portfolios, they use the Fama

French five-factor model to compute risk-adjusted alphas.

They find that, in the year following UPIA, portfolios of prudent stocks yield a negative monthly alpha (-0.191%) and those of less prudent stocks a positive alpha (0.256%). A long–short strategy that goes long on less prudent stocks and short on prudent stocks generates an abnormal return of about 0.447% per month (roughly 5.3% per year), while higher ψ categories show minimal effects. A permutation-based falsification test further confirms that the timing of the regulatory change is crucial.

The authors then estimate price elasticity by regressing stock returns on "residual demand" (RD)—which represents the change in trust demand not offset by other investors—using the model:

$$R_i = \delta + \beta RD_i + \gamma X_{i,-3} + \varepsilon_i$$

Here, β indicates the sensitivity of stock returns to changes in RD; elasticity is interpreted as the percentage change in stock price induced by a one-percent change in RD. With elasticities computed as $\varepsilon = \frac{1}{\beta}$ and estimated in the range of 0.11 to 0.36, even a considerable residual demand shift results in only modest price changes. This provides strong evidence that U.S. equity prices are inelastic, meaning they do not respond substantially to demand changes.

Bibliography

1. Gabaix, X., & Koijen, R. S. J. (2021). *In search of the origins of financial fluctuations: The inelastic markets hypothesis*. NBER Working Paper No. 28967.
2. Bouchaud, J.-P. (2021). *The inelastic market hypothesis: A microstructural interpretation*. arXiv:2108.00242.
3. Cassella, S., Rizzo, A. E., Spalt, O., & Zimmerer, L. (2024). *Inelastic U.S. equity markets: New evidence from a reform of fiduciary duties*. AFA.