

Halo Effects in Asset Prices: Evidence from Televised Business Pitches

Matthew Flynn*, Janet Payne*, Marilyn Wiley†

January 16, 2026

Abstract

This paper investigates categorical belief formation in the context of investment decision-making. We develop a simple model of belief formation in which sentiment shocks to a focal firm generate short-run price spillovers to categorically related but economically distinct firms, followed by correction as beliefs revert. We extend earlier studies showing transient price increases in response to attention-getting events using a custom-trained large language model to extract structured data from *Shark Tank* product pitches displayed to a large viewing audience. We show that a signal from these perceived reliable sources propagates to the larger industry of related firms in the stock market, even though there is no logical connection. Industries matched to high-sentiment pitches earn economically meaningful next-day abnormal returns of 10–40 basis points that dissipate within days and generate no longer-horizon alpha. Robustness tests confirm that spillovers are stronger at the industry level than at the firm level, suggesting that investors generalize enthusiasm across broad categories rather than targeting specific analogs. We propose that these results are rooted in behavioral biases. Primarily, they are explained by the cognitive bias known as the *halo effect*, whereby a positive impression of an announced product investment decision influences investor perceptions of loosely related public equities.

*McCoy College of Business, Texas State University.

†G. Brint Ryan College of Business, University of North Texas.

1 Introduction

This paper studies a particular form of belief distortion arising from categorical inference. We propose that investors map salient signals about a focal, non-investable object onto economically distinct but categorically related assets. When an object receives a strong positive or negative evaluation, investors may extrapolate that impression to a broader category, even when the connection is imprecise and the signal contains no value-relevant information. In psychology, this phenomenon is known as the *halo effect*: an overall impression of one object contaminates judgments about other, related objects along dimensions that are not directly informed by new evidence. While the halo effect is well studied in other disciplines, its role in asset pricing has received little direct empirical attention.

Our contribution builds on, but is distinct from, existing evidence that retail investors respond to attention-grabbing signals and expert cues. Prior work shows that individual investors exhibit attention driven trading behavior, particularly in response to salient information. For example, Barber et al. document that less-experienced Robinhood investors engage in speculative trading linked to platform design and salient stimuli, consistent with earlier evidence of purchase herding around attention-grabbing events (Barber and Odean (2008)). Similarly, Engelberg et al. (2012) study stock recommendations on the television program *Mad Money* and find that perceived expert endorsements generate large but transient price increases that dissipate quickly, with stronger effects in settings where arbitrage is limited. Together, these studies demonstrate that retail investors respond strongly to salient signals tied directly to specific publicly traded stocks. The lack of persistence of the price increase implies that the response is not based on fundamental information but is behaviorally induced.

We extend this literature by studying a distinct and previously unexplored mechanism: belief spillovers from salient evaluations of private, non-investable firms to publicly traded firms. In contrast to settings where investors react to explicit stock recommendations or firm-specific news, our setting features sentiment shocks that plausibly convey no information about public firms' fundamentals. We hypothesize that retail investors who cannot directly invest in a security may invest in the closest available proxy in order to act on an attention-

getting event. This assertion is supported by the work of Bernstein et al. (2017), who study private equity markets. They find that private equity investments have a positive effect on the valuation of the industry as a whole. Aldatmaz and Brown (2020) find similar results when comparing private equity investments to the value of publicly traded firms in the same industry.

We develop a simple theoretical model of halo-driven belief formation to structure our empirical tests and distinguish from alternative explanations. In the model, sentiment shocks to a focal firm affect investors' beliefs about categorically related firms, generating short-run price comovement that is unrelated to fundamentals. Because these belief spillovers are not grounded in cash-flow information, the model predicts that price effects should be immediate but transitory, should not generate long-horizon abnormal performance, and should be amplified in settings where limits to arbitrage slow correction. These predictions guide our empirical design.

This model is rooted in well-established theories of heuristic reasoning. Kahneman and Tversky (1972) show that individuals often rely on representativeness and stereotypes rather than statistical inference when forming beliefs under uncertainty, giving rise to systematic biases. The halo effect, first formalized by Thorndike (1920), is a specific manifestation of this tendency: an investor's overall positive impression of a company, person, or product in one area positively influences their judgments about other unrelated qualities or other loosely related companies, persons, or products. This bias can lead to irrational investment choices, such as overlooking critical financial data in favor of superficial attributes.

Our empirical tests use a unique data set derived from transcripts of the US television show *Shark Tank*. The show provides a unique laboratory for testing halo-driven belief formation. Each episode delivers highly salient, widely viewed evaluations of private startups by prominent investors, generating sharp sentiment shocks while plausibly conveying no direct information about the cash flows, risk, or prospects of publicly traded firms. Using a custom-trained large language model, we extract structured data on individual firm pitches from *Shark Tank* transcripts and map each pitch to thematically related public industries. This mapping allows us to examine how sentiment expressed during the show spills over into public equity markets at the category level.

Our empirical analysis yields three main findings. First, industries matched to high-sentiment pitches earn economically meaningful abnormal returns on the next trading day. A one-unit increase in pitch sentiment is associated with approximately 10 basis points of next-day abnormal return, and a one standard deviation increase in sentiment corresponds to roughly 20–25 basis points. Returns increase monotonically across sentiment quintiles, with a top–bottom spread of over 40 basis points. Second, these effects are short-lived. Abnormal returns attenuate within days and generate no persistent alpha over longer horizons in Fama–French factor models, consistent with temporary mispricing rather than information transmission. Third, sentiment-driven price responses are amplified in industries with greater illiquidity and lower trading volume, consistent with limits-to-arbitrage theories. These results suggest that retail investors can form broadly positive opinions about products or industries based on attention that then influences their trading behavior.

Our findings contribute to a growing literature on the influence of media in financial markets. Previous work has established that pessimistic financial news predicts short-term stock underperformance (Tetlock (2007)), and that both tone and volume of media coverage affect investor behavior (Fang and Peress (2009)). Moreover, individual investors are known to exhibit attention-driven trading behavior, particularly in response to salient but nonfundamental stimuli (Barber et al. (2022)). We extend these insights by showing that entertainment content can activate similar behavioral mechanisms, producing a sector halo effect.

This paper is the first, to our knowledge, to link non-financial entertainment to systematic, cross-firm price effects in public equity markets. Unlike prior studies that examine media-induced responses to explicit stock recommendations (Engelberg et al. (2012)) or market-wide sentiment shifts (Edmans et al. (2007); Lepori (2015)), our setting illustrates entire sector spillover in investor sentiment stemming from affective and categorical inference. This form of extrapolative belief formation from private firm pitches to public firm return formation is novel in the literature.

In sum, our findings reveal how narratives embedded in entertainment media can drive temporary mispricing through behavioral channels, particularly in segments of the market with heightened retail trading and subject to limits to arbitrage. By isolating a setting

with strong sentiment, limited information content, and clear attenuation, we provide clean evidence that categorical belief formation plays a meaningful role in asset price formation.

2 Halo Effects and Category-Level Belief Formation

This section provides the conceptual foundations for the halo-driven belief distortions modeled in Section 3. We draw on insights from psychology, behavioral economics, and related fields to describe how evaluative judgments formed in response to salient information about a focal object can spill over to affect beliefs about related but economically distinct objects. While the halo effect has been extensively documented in individual decision-making contexts, its implications for belief formation at the category level in asset pricing are not yet studied.

We first review the origins of the halo effect in psychology and decision-making, emphasizing how affective impressions and salient signals contaminate evaluations along dimensions unrelated to fundamentals. We then discuss the mechanisms through which halo bias arises and propagates and imprecise mental mapping across related entities. Finally, we emphasize the difficulty in distinguishing halo-driven belief spillovers from rational information transmissions. This framework motivates the model and empirical analysis that follow.

2.1 The Halo Effect in Psychology and Decision-Making

The halo effect is a well-documented cognitive bias first studied in psychology by Thorndike (1920). Thorndike described it as a “constant error” in which an observer’s general impression of a person (e.g. seeing them as overall good or bad) colors the evaluation of that person’s specific traits. A single positive attribute can create a “halo” that leads us to judge unrelated attributes more favorably, and likewise a single negative trait can cast a “horn” effect causing disproportionately harsh judgments. The halo effect thus reflects a failure to weigh all attributes objectively, letting an overall positive/negative feeling dominate one’s judgment. Psychologists have found this bias to be unconscious and pervasive. Early experiments showed that people’s liking for an instructor’s physical appearance could unconsciously alter their ratings of the instructor’s unrelated qualities such as performance (Nisbett and

Wilson, 1977). The key insight is that once a global evaluation is formed, it is cognitively difficult to prevent that affective impression from influencing specific assessments.

In marketing and consumer behavior, the halo effect is a central mechanism underlying brand equity and reputation spillovers. Positive experiences with a flagship product, brand, or endorsement routinely elevate perceptions of other products within the same brand family, independent of their intrinsic quality. Empirical evidence from marketing research shows that consumers generalize favorable attributes (such as brand prestige) across products and categories, leading to sustained loyalty and price premia, but also systematic misjudgment when quality does not, in fact, generalize. These findings highlight how coarse brand-level affect can dominate product-specific information in consumer decision-making.

Although a formal study of the halo effect on asset prices is novel in finance literature, other studies have indirectly captured related effects. Jang et al. (2016) document reputation-driven biases in equity markets, showing that investor buying behavior around earnings announcements is influenced by prior corporate reputation. More generally, finance research has also explored category-based biases. Barberis et al. (2005) suggest that many investors simplify decisions by grouping assets into broad categories (styles, industries, themes), leading to comovement within categories. Implicitly, this behavior can create halo-like spillovers: when one asset in a category experiences positive news, other assets in the same category benefit. A related concept in economics literature is categorical thinking or coarse belief partitioning (Mullainathan et al., 2008), under which individuals use mental “buckets” to partition information.

2.2 Mechanisms of Halo Bias Formation

While the existence of halo effects is well established across psychology and consumer research, the precise mechanisms through which halo-biased beliefs form and propagate in financial markets remain underexplored. In the context of asset pricing, halo effects arise not from a single cognitive failure, but from a combination of well-documented heuristics that lead investors to overgeneralize salient information across related firms or categories. We highlight two mechanisms that are particularly relevant for belief formation in markets.

A first mechanism is the *affect heuristic*, whereby individuals rely on coarse emotional

evaluations when forming judgments. Research in psychology demonstrates that people often substitute difficult probabilistic assessments with an overall sense of “goodness” or “badness,” allowing affect to guide beliefs about unrelated attributes. In an investment setting, this implies that a favorable impression of a firm can inflate expectations about its broader prospects, even in the absence of corroborating information. When applied at the category level, affect-driven judgments may spill over to other firms that investors loosely associate with the focal firm, generating halo-driven demand shocks.

A second mechanism is the *representativeness heuristic* (Tversky and Kahneman, 1974), through which individuals infer group-level characteristics from a small number of salient examples. A prominent success story may come to represent an entire industry or theme in investors’ minds, leading them to extrapolate its perceived quality to peer firms. This process is closely related to coarse categorization and style-based thinking: rather than distinguishing among firms based on fine-grained fundamentals, investors form beliefs at a much less granular level, lightening cognitive load. As emphasized by Barberis et al. (2005), such categorical thinking can generate excess comovement in asset prices, even when firms’ underlying cash flows are largely independent.

Importantly, these mechanisms do not require investors to consciously ignore information or act irrationally in a narrow sense. Instead, halo effects emerge as a byproduct of cognitive shortcuts that simplify decision-making in complex environments. The same heuristics that are adaptive in everyday judgment can, in financial markets, produce belief distortions that propagate across firms and categories.

These mechanisms give rise to a central tension for asset pricing: observed spillovers across firms may reflect either rational information transmission—where news about one firm signals information about others—or halo-driven belief formation, where investors overgeneralize sentiment beyond what fundamentals justify. Distinguishing between these channels is essential for understanding when price comovement reflects shared economic risks and when it represents temporary mispricing, and is a challenging econometric problem. In the next section, we formalize this distinction in a stylized model that isolates halo-driven belief spillovers in the absence of shared fundamentals.

3 A Model of Halo-Driven Mispricing and Correction

We now construct a stylized model of asset prices incorporating halo-biased investors. The model captures how halo-driven mispricing and temporary price divergence arise and subsequently correct in financial markets. We maintain simplicity for interpretability, focusing on belief formation for a focal firm and a representative peer within an industry.

To begin, consider two firms, indexed by $i \in \{A, B\}$, belonging to the same investor-defined industry or thematic category. Firm A denotes the focal firm that receives a salient news shock, and Firm B denotes a representative peer. Although the model is written in terms of a single peer firm B , it should be interpreted as a generic member of a broader category of firms that different investors may associate with the focal firm through heterogeneous, potentially imprecise, cognitive mappings.

Time unfolds over three discrete periods. At $t = 0$, no new information has arrived and asset prices reflect prior beliefs about fundamentals. At $t = 1$, Firm A is subject to an exogenous news shock. At $t = 2$, fundamentals are realized. Each firm i has a fundamental value V_i that is realized at $t = 2$. The variable V_i may be interpreted as the present value of future cash flows or, more generally, the firm's true economic worth. For simplicity, and abstracting from any true common factor, we assume that the fundamental values of Firms A and B are independent, so that there is no informational linkage between the two firms' projects or cash flows. Formally, let

$$V_A = \mu + \varepsilon_A, \quad V_B = \mu + \varepsilon_B, \quad (1)$$

where μ denotes the common prior expected value of firms in the industry or group, and ε_A and ε_B are firm-specific idiosyncratic shocks. We assume that these shocks have zero mean and are mutually independent,

$$\mathbb{E}[\varepsilon_i] = 0, \quad \varepsilon_A \perp \varepsilon_B. \quad (2)$$

At $t = 0$, no firm-specific information has been revealed. Under risk-neutral pricing, rational expectations, and the absence of arbitrage, the price of each firm equals the expected

value of its fundamentals. Consequently, initial prices satisfy

$$P_A(0) = P_B(0) = \mathbb{E}[V_i] = \mu. \quad (3)$$

3.1 Investor Beliefs and Halo Bias

We now introduce heterogeneity in investor beliefs. There are two types of investors. The first type consists of rational investors, who correctly understand that the fundamental values of Firms A and B are independent and update beliefs using Bayes' rule. The second type consists of halo-biased investors, who overreact to category-level signals and behave as if the valuation-relevant fundamentals of Firms A and B are positively related.

We formalize halo bias using a parameter $\rho \in [0, 1]$, which captures the strength of biased investors' perceived association between Firm A and a representative peer firm B . A fully rational investor corresponds to $\rho = 0$, whereas a more strongly biased investor behaves as if $\rho > 0$. Importantly, ρ represents an average perceived association strength across investors, rather than as an objective measure of firm-to-firm similarity. Different investors may associate the focal firm with different subsets of peer firms, but in the aggregate these heterogeneous cognitive mappings generate a positive perceived linkage at the category level. Intuitively, halo-biased investors believe that positive news about Firm A contains information relevant for firms they associate with the same category, even when no such informational linkage exists.

3.2 Information Structure

At $t = 1$, Firm A receives a public and salient signal about its fundamental value. Specifically, investors observe:

$$s_A = V_A + \eta, \quad (4)$$

where η is a noise term with mean zero. For tractability, we assume that s_A is an unbiased signal of V_A and contains no direct information about the fundamentals of other firms.

Firm B receives no new firm-specific information at $t = 1$; its fundamental value will be revealed at $t = 2$. As a result, any price response of Firm B at $t = 1$ reflects belief updating

based on the signal to Firm A rather than the arrival of new information about Firm B itself.

3.3 Belief Updating

Rational investors. Rational investors update their beliefs about Firm A based on the signal s_A , forming the posterior expectation $\mathbb{E}[V_A | s_A]$, which is approximately equal to s_A under the assumption of zero-mean noise. Because V_B is independent of V_A by construction, rational investors do not revise their beliefs about Firm B :

$$\mathbb{E}[V_B | s_A] = \mu. \quad (5)$$

Accordingly, under rational expectations, the price of Firm A responds to the signal at $t = 1$, while the true value of Firm B remains at its prior price μ .

Halo-biased investors. Halo-biased investors, in contrast, employ a mis-specified belief-updating process. Rather than treating the signal s_A as purely firm-specific, they behave as though it is informative about the valuation-relevant fundamentals of firms they associate with the same category as Firm A . A simple reduced-form representation of this belief distortion is given by:

$$\mathbb{E}^b[V_B | s_A] = \mu + \rho(s_A - \mu), \quad (6)$$

where $\mathbb{E}^b[\cdot]$ denotes expectations under biased beliefs. When $\rho = 0$, biased and unbiased investors agree on value. When $\rho > 0$, a positive (negative) surprise for Firm A leads halo-biased investors to revise upward (downward) their valuation of Firm B . As before, ρ captures the average strength of halo-driven belief spillovers across heterogeneous investors, rather than any one direct firm-to-firm correlation. This belief distortion is consistent with one or more of the halo mechanisms discussed in Section 2.2.

3.4 Prices at $t = 1$

Let a fraction $\theta \in [0, 1]$ of investors be halo-biased, with the remaining investors rational. If arbitrage is limited, prices reflect the aggregation of heterogeneous beliefs in the market. In this case, the price of Firm B at $t = 1$ is given by a weighted average of rational and biased expectations:

$$P_B(1) = (1 - \theta) \mathbb{E}[V_B \mid s_A] + \theta \mathbb{E}^b[V_B \mid s_A]. \quad (7)$$

Substituting from equations (5) and (6) yields

$$P_B(1) = \mu + \theta \rho(s_A - \mu). \quad (8)$$

Equation (8) captures halo-driven mispricing at $t = 1$. The term $\theta \rho(s_A - \mu)$ represents a halo premium (or discount) embedded in Firm B 's price due to category-level sentiment induced by news about Firm A . When $s_A > \mu$, halo-biased investors push the price of Firm B above its fundamental value; when $s_A < \mu$, they depress the price below fundamentals.

3.5 Resolution of Uncertainty and Price Correction at $t = 2$

At $t = 2$, the fundamental value of Firm B is revealed, for example through the realization of cash flows or the arrival of firm-specific information. At this point, uncertainty regarding Firm B 's fundamentals is resolved. The realized value of Firm B is given by

$$V_B = \mu + \varepsilon_B. \quad (9)$$

Once Firm B 's fundamentals are observed, prices reflect realized values. Accordingly, the price of Firm B at $t = 2$ equals its fundamental value:

$$P_B(2) = V_B. \quad (10)$$

3.6 Price Dynamics and Halo-Induced Mispricing

Combining the results from the previous sections, the price path of Firm B over time is given by

$$P_B(0) = \mu, \quad (11)$$

$$P_B(1) = \mu + \theta\rho(s_A - \mu), \quad (12)$$

$$P_B(2) = \mu + \varepsilon_B. \quad (13)$$

Equation (12) highlights the presence of halo-induced mispricing at $t = 1$. For example, suppose Firm A experiences a positive surprise such that $s_A - \mu = \Delta > 0$. In this case, the price of Firm B at $t = 1$ is

$$P_B(1) = \mu + \theta\rho\Delta, \quad (14)$$

so that the deviation of price from the prior fundamental mean is

$$P_B(1) - \mu = \theta\rho\Delta. \quad (15)$$

This term represents the halo premium embedded in Firm B 's price following the sentiment shock to Firm A .

At $t = 2$, this mispricing corrects as Firm B 's fundamental value is revealed. Taking expectations with respect to ε_B , we have $\mathbb{E}[V_B] = \mu$, implying that prices revert to the prior mean on average. When Firm A experiences a positive surprise but Firm B does not, the price of Firm B declines from $t = 1$ to $t = 2$ as the halo-induced optimism unwinds. The expected magnitude of this correction is

$$\mathbb{E}[P_B(2) - P_B(1)] = -\theta\rho(s_A - \mu). \quad (16)$$

Together, these dynamics generate a pattern of short-term price divergence followed by convergence to fundamentals: prices temporarily deviate from intrinsic values in response to category-level sentiment shocks, and subsequently adjust as firm-specific information arrives.

3.7 Limits to Arbitrage

Halo-induced price divergence at $t = 1$ arises from biased belief formation and need not rely on limits to arbitrage. However, limits to arbitrage play a central role in determining the magnitude and persistence of the resulting mispricing. Although rational investors recognize that Firm B may be mispriced by an amount $\theta\rho(s_A - \mu)$, they face uncertainty regarding the realization of Firm B 's idiosyncratic shock ε_B .

If ε_B is sufficiently positive, Firm B 's realized fundamental value may ultimately justify its elevated price, creating risk for arbitrageurs who attempt to short the stock at $t = 1$. As a result, even rational traders who identify the mispricing may optimally limit their positions.

When rational traders are risk-averse or capital constrained, as in the limits-to-arbitrage framework of Shleifer and Vishny, 1997, they will not fully offset biased demand. In this case, halo-biased beliefs exert a larger influence on equilibrium prices, allowing mispricing to persist until firm-specific fundamentals are revealed. Once uncertainty is resolved at $t = 2$, prices converge to realized values and the halo premium dissipates. Thus, while halo-driven price divergence arises from biased belief formation, the extent to which such divergence is reflected in equilibrium prices depends on the severity of limits to arbitrage faced by rational traders.

3.8 Continuous-Time Interpretation

Although the model is presented in discrete time for expositional clarity, its intuition extends naturally to an event-time, continuous-time environment in which information arrives and prices adjust asynchronously. The key ingredients remain unchanged: a salient signal about a focal firm generates a category-level belief shock among halo-biased investors, and prices subsequently converge as firm-specific fundamentals are revealed. Consistent with the categorical nature of halo bias, these spillovers need not concentrate on a uniquely identified set of “closest” peers; rather, they operate through investor-defined categories.

To illustrate, let $t \geq 0$ denote event time, with the signal about Firm A arriving at $t = 0$. Let m_t denote a category-level sentiment state that jumps on impact and then mean-reverts

as uncertainty resolves. A parsimonious representation is

$$dm_t = -\kappa m_t dt, \quad m_{0+} = \theta \rho(s_A - \mu), \quad (17)$$

where $\kappa > 0$ governs the speed at which halo-induced sentiment dissipates.¹

Let the price of a representative peer firm B be decomposed into fundamentals and a transitory halo component,

$$P_B(t) = \mathbb{E}[V_B \mid \mathcal{I}_t] + m_t, \quad (18)$$

where $\mathcal{I}_t$ denotes the information available at time t . Prior to the revelation of V_B , the conditional expectation term remains anchored at μ in this stylized setting, while the halo component decays at rate κ . When firm-specific information about B arrives and V_B is resolved, the halo component vanishes and prices converge to realized values.

This continuous-time interpretation delivers the same qualitative predictions as the discrete-time model: an immediate category-level price response at event time, followed by gradual correction as firm-level fundamentals become salient. The speed and shape of the reversal depend on the rate of information revelation and the persistence of halo-driven sentiment.

3.9 Testable Implications

The model of halo-driven belief formation developed in Section 3 yields a set of empirically testable implications that distinguish sentiment-based category spillovers from rational information transmission. In particular, the model predicts short-run price comovement across broadly defined categories following salient firm-specific news, followed by subsequent correction as firm-level fundamentals are revealed.

Hypothesis 1 (Direction and Magnitude of Short-Run Halo Spillovers). Following a salient sentiment shock to a focal firm, peer firms within the same industry or thematic category experience abnormal returns in the short run, despite no new information about their own fundamentals. Positive (negative) sentiment shocks generate positive (negative) abnormal returns for peer firms relative to firms outside the treated category. As implied

¹More generally, one can allow m_t to follow a stochastic mean-reverting process; equation (17) captures the deterministic component that is sufficient for the model's comparative statics.

by equation (8), the magnitude of the short-run response is increasing in the strength of the sentiment shock ($s_A - \mu$), the prevalence of halo-biased investors θ , and the average perceived association strength ρ .

Hypothesis 2 (Reversion and the Role of Arbitrage Frictions). Halo-induced price responses among peer firms are temporary. As firm-specific information arrives and uncertainty regarding peer fundamentals is resolved, prices revert toward intrinsic values, as characterized by equation (16). While rational investors may recognize that peer firms are mispriced following sentiment shocks, limits to arbitrage can delay correction, allowing halo-driven price deviations to persist in the short run and unwind gradually over time.

Mechanism Consistency: Fundamentals. A halo-based interpretation implies that short-run peer-price responses need not be accompanied by commensurate changes in peer firms' fundamental prospects. We therefore examine whether post-event fundamental outcomes and fundamentals-based expectations for peer firms change following sentiment shocks. A lack of systematic improvement in such measures is consistent with halo-driven mispricing rather than rational updating about shared cash flows.

Additionally, the model predicts that halo effects should manifest most clearly at the category or industry level, reflecting categorical thinking and coarse belief partitioning. While investors may readily generalize salient narratives to broad themes or industries, heterogeneous investors are unlikely to independently map a focal firm to the same precise firm analog. As a result, sentiment-driven belief spillovers are expected to aggregate at the category level rather than at the individual firm level.

4 Institutional Setting and Data: *Shark Tank* as a Laboratory

This section introduces the institutional setting and data used to study halo-driven belief spillovers in asset prices and to address the associated identification challenges. Empirically identifying halo effects in financial markets is challenging, as shocks to a focal firm often contain relevant information about firms in its network, making it difficult to distinguish

belief-driven price responses from rational information transmission. For example, consider a major information shock to a large technology firm, such as the announcement of a breakthrough in proprietary chip design or artificial intelligence infrastructure. Such an event is likely to affect not only the focal firm's valuation but also the prices of related firms throughout the technology sector. In this case, price responses among peer firms may reflect a mixture of rational information updating and sentiment-driven extrapolation by investors who generalize optimism about the focal firm to the broader category. Disentangling these channels is empirically difficult because both mechanisms operate simultaneously.

We therefore focus on a setting that generates salient information shocks while remaining largely orthogonal to the cash flows and risk exposures of publicly traded firms. The television program *Shark Tank* provides such an environment: it delivers widely observed evaluations of private startups whose prospects are unlikely to contain material information about the fundamentals of public companies. This section describes the information environment of *Shark Tank*, the construction of the dataset from episode transcripts, and the procedure used to map televised pitches to public firms and industries.

4.1 *Shark Tank* as an Experimental Setting

Shark Tank is a nationally televised program in which entrepreneurs pitch early-stage private businesses to a panel of prominent angel investors (the “Sharks”²) in an effort to secure funding. Each episode features multiple pitches that receive highly visible, evaluative reactions from the Sharks, ranging from enthusiastic endorsement to outright rejection. These reactions are conveyed through both explicit investment decisions and qualitative commentary, making them salient to a broad viewing audience.

Several features make *Shark Tank* a particularly well-suited setting for studying halo-driven belief spillovers as formalized in Section 3. First, each pitch generates a discrete, salient sentiment shock whose tone and intensity vary substantially across episodes, products, and narratives. These shocks are inherently qualitative and evaluative, providing a

²The panel of investors (the “Sharks”) consists of prominent entrepreneurs and venture capitalists with well-established public reputations across technology, consumer products, retail, real estate, and finance. Appendix Table A1 provides brief background profiles for the primary Sharks featured in the sample.

natural source of heterogeneous belief updating rather than precise fundamental signals. Second, pitches are broadcast to a broad audience and occur at well-defined points in time, allowing for clean event-time identification of price responses. Third, and most importantly, the focal firms are private startups whose business prospects are unlikely to be mechanically linked to the cash flows or risk exposures of publicly traded firms. This institutional feature aligns closely with the model's assumption that any perceived correlation between the focal firm and related public firms reflects investors' subjective categorization rather than true informational overlap. Consequently, observed price comovement among public firms categorized as related to a pitch is difficult to rationalize through rational information transmission and instead maps naturally to halo-driven belief formation operating at the category level.

Episodes typically last one hour and feature approximately four pitches, each following a common structure: entrepreneurs present their product and valuation, discuss market opportunity and business metrics, respond to questioning by the Sharks, and engage in negotiations when interest is expressed. While individual pitches are edited for broadcast, the underlying information and evaluative commentary are preserved in full transcripts. Our sample spans over 15 seasons and more than 350 episodes, providing a large and diverse set of sentiment shocks across industries and time.

4.2 Pitch Characteristics and Sentiment Variation

Although *Shark Tank* pitches follow a common format, the evaluative content conveyed to viewers varies substantially across episodes. Pitches differ in the tone and intensity of the Sharks' reactions and the degree of expressed confidence or skepticism in the pitch/firm. In some cases, Sharks offer enthusiastic endorsements and compete to invest; in others, they issue blunt critiques or collectively opt out. This variation generates a rich distribution of sentiment shocks that are plausibly interpreted by viewers as informative judgments by expert investors.

This heterogeneity in pitch-level sentiment is central to our empirical design. In the model developed in Section 3, the magnitude of halo-driven spillovers depends on both the strength of the sentiment shock and the prevalence of biased belief formation. The variation in pitch

characteristics and sentiment across pitches and episodes therefore provides a natural source of identifying variation to test the model’s predictions.

4.3 Transcript Data and LLM-Based Extraction

Because *Shark Tank* transcripts are conversational, multi-speaker, and lack uniform structure, we employ a large language model (LLM), OpenAI GPT-4o, to extract structured information at scale. Rule-based or keyword-driven natural language processing approaches are poorly suited to this setting, as they struggle to capture contextual meaning, such as speaker roles and negotiation dynamics. In contrast, GPT-based models leverage contextual embeddings to identify entities, deal terms, valuations, and evaluative language with high accuracy. Recent evidence shows that LLMs can perform comparably to human annotators while substantially reducing time and cost (Calderon et al., 2025).

We develop and iteratively refine a custom GPT pipeline to parse episode transcripts into structured variables, including company and product descriptions, entrepreneur characteristics, pitch-level sentiment, deal outcomes, and episode metadata. Model performance is evaluated against a manually annotated validation set using standard metrics, including precision, recall, and the F1 score. These tests indicate highly reliable extraction. Full technical details on model construction, prompts, and replication materials are provided in Appendix A.

4.4 Mapping Pitches to Public Firms and Industries

To study how pitch narratives propagate to public equity markets, each *Shark Tank* pitch must be linked to thematically related publicly traded firms. We implement this mapping using text-based similarity measures between pitch transcripts and firm business descriptions from Compustat. After standard preprocessing (lowercasing, stopword removal, stemming), we construct term-frequency inverse-document-frequency (TF-IDF) representations and compute cosine similarity scores between each pitch and all public firm descriptions.

We employ two complementary mapping approaches. In our baseline specification, each pitch is assigned to a broad industry grouping by identifying the most thematically similar

SIC category based on text similarity. This approach captures categorical associations that reflect economically meaningful industry classifications and aligns with the model's emphasis on coarse belief formation and category-level spillovers. The resulting industry portfolios serve as the primary measure of exposure to pitch-induced sentiment shocks. Figure 1 provides a conceptual illustration of the mapping procedure. In addition to the baseline industry-level mapping, we construct an alternative firm-level matching that links each pitch to the single public firm with the highest textual similarity score. This mapping is used in supplementary analyses to assess the sensitivity of results to narrower definitions of exposure.

4.5 Descriptive Statistics

The transcript extraction process yields structured data on 1,286 individual *Shark Tank* pitches. These pitch-level variables are merged with financial market data from CRSP and Compustat, which provide daily stock returns, trading volume, and firm characteristics for publicly traded firms. The resulting dataset links *Shark Tank* pitches to public market outcomes.

Table 1 reports descriptive statistics for key pitch-level variables, separately for pitches that receive an investment deal and those that do not. Entrepreneurs who receive deals request an average of \$263,149 in funding in exchange for approximately 12.7 percent equity, corresponding to an implied ask valuation of roughly \$3.5 million. Final deal terms typically involve larger equity stakes (mean of 22.9 percent) and lower implied valuations (mean of \$2.3 million), indicating substantial renegotiation during the pitch process. The average change in implied valuation from ask to deal is negative, with a mean reduction of approximately \$1.2 million.

Non-equity deal structures are relatively uncommon. Royalty-based agreements appear in roughly 7.5 percent of completed deals, while debt components are present in fewer than 5 percent. Among pitches that do not receive a deal, requested funding amounts and implied valuations are somewhat higher on average, while sentiment scores are markedly lower.

Pitch-level sentiment exhibits substantial dispersion across episodes and differs systematically by deal outcome. Pitches that receive a deal have a higher average sentiment score (1.37) than those that do not (0.73), though the distribution of sentiment is wide in both

groups. This variation in evaluative tone—orthogonal to fundamentals of public firms—plays a central role in our empirical analysis. Overall, the summary statistics illustrate both the scale of the dataset and the presence of meaningful heterogeneity in pitch characteristics and sentiment intensity.

5 Empirical Strategy

5.1 Halo Effect - Short-Run Effects of Pitches on Industry Returns

We begin by testing Hypothesis 1 from Section 3, which concerns the direction and magnitude of short-run halo spillovers following salient sentiment shocks. As shown in equation (8), when a focal firm receives a positive (negative) sentiment shock, halo-biased investors partially extrapolate this signal to economically distinct but categorically related firms, generating abnormal price movements at the category level. Importantly, these spillovers arise despite the absence of new information about peer firms' fundamentals.

To operationalize this prediction, we examine whether the sentiment conveyed during *Shark Tank* pitches induces short-run abnormal returns in public equity portfolios matched to the pitch's industry category. The empirical design mirrors the structure of the model: pitch sentiment serves as the observable analogue of the sentiment shock ($s_A - \mu$), while industry-level abnormal returns capture the aggregate price response of peer firms that investors loosely associate with the focal firm through categorical belief formation.

Specifically, we estimate the following baseline specification:

$$\begin{aligned} \text{AbRet}_{it} = & \beta_1 \text{Sentiment}_t + \beta_2 \text{DealMade}_t + \beta_3 \text{ValuationMean}_t + \beta_4 \text{ValuationChange}_t \\ & + \beta_5 \text{AbRet}_{i,t-1} + \beta_6 \text{MktRet}_t + \beta_7 \text{Size}_i + \beta_8 \text{Illiquidity}_i + \beta_9 \text{IdioVol}_i + \alpha_i + \lambda_t + \varepsilon_{it}, \end{aligned} \quad (19)$$

where AbRet_{it} denotes the abnormal return for industry i on trading day t , defined as the industry portfolio return net of the contemporaneous S&P 500 return. We examine short event windows surrounding the episode air date, focusing on both overnight returns (close-to-open) and next-day returns (open-to-close), which correspond closely to the model's

prediction of immediate belief-driven price responses.

The coefficient of primary interest is β_1 , which captures the sensitivity of industry-level returns to pitch sentiment. In the model, pitch sentiment serves as the empirical analogue of the sentiment shock ($s_A - \mu$), and a positive (negative) estimate of β_1 reflects halo-driven extrapolation of favorable (unfavorable) narratives to categorically related public firms. Consistent with equation (8), the magnitude of β_1 reflects the combined influence of the prevalence of halo-biased investors and the average perceived association strength between the focal pitch and the matched industry.

The remaining covariates play a disciplining role rather than representing alternative channels emphasized by the theory. *DealMade* indicates whether a pitch results in an on-air investment agreement, while *ValuationMean* and *ValuationChange* capture the level and revision of implied valuations during negotiation. Including these variables allows us to assess whether observed return responses are driven by economically interpretable deal information, as opposed to the narrative tone of the pitch itself. Lagged industry returns control for short-term return persistence, while contemporaneous market returns absorb broad market movements. Industry size, Amihud illiquidity, and idiosyncratic volatility capture cross-sectional differences in trading frictions and risk that may influence short-run price dynamics.

All specifications are estimated with and without industry fixed effects (α_i) and air-date fixed effects (λ_t). Industry fixed effects absorb time-invariant differences across industries, while air-date fixed effects control for episode-specific conditions, common shocks, and market-wide sentiment on the broadcast date. This fixed-effects structure ensures that identification of β_1 is driven by variation in sentiment across industries within the same episode and by variation across episodes for a given industry, rather than by aggregate market movements or persistent industry characteristics. In extended specifications, we interact pitch sentiment with individual Shark indicators to test whether sentiment spillovers vary with the prominence of the endorsing investor, though such heterogeneity is not required for the baseline halo mechanism to operate.

Finally, to provide transparent evidence that does not rely solely on functional form assumptions, we complement the regression analysis with nonparametric sorts of next-day industry returns by pitch sentiment. These sorts directly illustrate whether industries asso-

ciated with more positive (negative) televised narratives systematically earn higher (lower) short-run returns, consistent with sentiment-driven belief spillovers.

5.2 Return Dynamics and Long-Horizon Abnormal Performance

Hypothesis 2 implies that halo-driven price effects are temporary and unwind as beliefs revert in the absence of fundamental information. In the model, category-level price responses at $t = 1$ arise from belief extrapolation rather than from information about peer firms' cash flows, and equation (16) predicts that these deviations should dissipate over time. This prediction has two empirically observable implications: first, sentiment-induced return effects should attenuate over short horizons following the initial shock; second, such effects should not generate persistent abnormal performance over longer holding periods.

We evaluate these implications using complementary approaches. We first examine the dynamic evolution of matched-industry returns over the days following *Shark Tank* airings, tracing whether the initial sentiment response decays as predicted. We then assess whether these short-run effects translate into long-horizon alpha by estimating Fama–French three-factor regressions over extended holding periods.

For each holding horizon, we estimate the following specification:

$$r_{pt} - r_{ft} = \alpha_p + \beta_1(r_{mt} - r_{ft}) + \beta_2\text{SMB}_t + \beta_3\text{HML}_t + \varepsilon_t, \quad (20)$$

where r_{pt} denotes the return on a portfolio of matched industries at time t , r_{ft} is the risk-free rate, and the regressors include the excess market return, the size factor (*SMB*), and the value factor (*HML*). Factor data are obtained from Ken French's research library. The dependent variable is thus the portfolio excess return.

We estimate this model across five cumulative holding horizons, ranging from 50 to 250 trading days, and for three complementary portfolio constructions:

1. *Equal-weighted portfolios*, which weight all matched industries equally;
2. *Deal portfolios*, which hold industries whose matched pitches result in on-air investment agreements;

3. *Long-short sentiment portfolios*, which go long industries associated with positive-sentiment pitches and short industries associated with negative-sentiment pitches.

Portfolios are formed using overlapping event cohorts. Each matched industry enters the portfolio on the first trading day on or after the episode air date and remains included for the specified holding horizon. This approach maximizes statistical power while preserving the event-time interpretation of the model. Because overlapping holding periods induce serial correlation in portfolio returns, standard errors are adjusted using the Newey–West procedure.

The coefficient of primary interest is α_p , which captures abnormal returns not explained by factor exposures. Under the halo-driven mispricing mechanism, α_p should be economically small and statistically insignificant across horizons, indicating that the price impact of televised pitches reflects temporary belief distortions rather than the revelation of value-relevant information about public firms. In contrast, persistent positive or negative values of α_p would suggest that *Shark Tank* airings convey information that is gradually incorporated into prices, contrary to the model’s central identifying assumption.

Accordingly, this analysis serves a dual purpose. First, it tests the prediction that halo-induced price responses unwind over time. Second, it provides a validation check on the empirical setting itself: the absence of long-horizon alpha supports the interpretation that *Shark Tank* functions as a sentiment shock rather than a channel of fundamental information transmission to public markets.

5.3 Limits to Arbitrage

The model developed in Section 3 predicts that halo-driven price pressures need not be fully arbitrated away in the short run. Even when rational investors recognize that category-level price responses are not justified by fundamentals, trading frictions and inventory risk can prevent immediate correction. As a result, equation (8) implies that the magnitude of short-run mispricing should be increasing in the severity of limits to arbitrage, while equation (16) implies that correction may be delayed rather than instantaneous.

We test this implication by examining whether the short-run sensitivity of industry re-

turns to pitch sentiment varies systematically with cross-sectional limits to arbitrage that directly affect the ability of rational traders to respond to perceived mispricing. Guided by the model, we focus on industry-level trading frictions—specifically illiquidity and trading volume—which govern how quickly and at what scale arbitrage capital can be deployed.

Formally, we estimate the following interaction specification:

$$\text{AbRet}_{it+1} = \alpha + \beta_1 \text{Sentiment}_t + \beta_2 \text{Limit}_i + \beta_3 (\text{Sentiment}_t \times \text{Limit}_i) + \gamma' \mathbf{X}_t + \varepsilon_{it}, \quad (21)$$

where the dependent variable is the next-day abnormal return for industry i , defined as the industry return net of the market return on the first trading day following the episode air date. Sentiment_t measures the average tone of the pitch narratives matched to the industry, and Limit_i denotes an industry-level proxy for limits to arbitrage.

We consider two measures that are most closely aligned with the model's mechanism:

- *Illiquidity quartile*, based on the Amihud illiquidity ratio, capturing price impact per unit of trading volume; and
- *Industry trading volume quartile*, defined as average dollar trading volume of constituent firms, proxying for market depth available to arbitrageurs.

Both measures are constructed as prior-month cross-sectional quartiles to ensure they are predetermined with respect to episode airings. The coefficient of interest is β_3 , which captures whether sentiment-driven return responses are amplified in industries where arbitrage is more constrained. Under the halo-driven mispricing mechanism, β_3 should be economically meaningful, reflecting stronger price pressure when rational investors face higher trading frictions.

All specifications include the same baseline control set used throughout the empirical analysis, including deal outcomes, valuation measures, lagged industry returns, contemporaneous market returns, and episode-level characteristics. The main specifications are estimated without fixed effects to preserve the cross-sectional variation in arbitrage constraints that is central to identification in this test. Fixed-effects variants are reported in the Appendix and yield qualitatively similar conclusions.³ We estimate models that include each

³Appendix Table A5 reports specifications including industry and air-date fixed effects. We do not report

interaction individually as well as a joint specification that allows illiquidity and trading volume constraints to operate simultaneously.

By focusing on trading-based limits to arbitrage, this design provides a sharp test of the model's central implication: if halo effects reflect temporary belief distortions rather than rational information transmission, their impact should be strongest precisely where arbitrage capital is least able to respond.

6 Results

6.1 Halo Effect - Short-Run Effects of Pitches on Industry Returns

Table 2 reports results from our baseline tests of Hypothesis 1, which examines whether sentiment conveyed during *Shark Tank* pitches generates short-run halo spillovers in public equity markets. The dependent variable is the abnormal return of the industry portfolio matched to each pitch, defined as the industry return net of the contemporaneous market return and measured either overnight (Columns 1 and 3) or over the following trading day (Columns 2 and 4).

The most consistent and economically meaningful predictor of next-day abnormal returns is pitch sentiment. In Column (2), a one-unit increase in average pitch sentiment about the non-tradeable asset is associated with an approximately 10 basis point increase in next-day matched-industry excess returns ($p < 0.01$). This magnitude is economically meaningful: given the observed dispersion in sentiment across pitches, a one standard deviation increase in sentiment score is associated with a 21.7 basis point increase in industry returns on the next trading day. Because the AFINN sentiment measure is constructed by assigning integer scores to individual words and averaging across the full pitch transcript, a one standard deviation change corresponds to a modest shift in narrative tone rather than an extreme change in language. For example, moving from a neutral pitch to one that contains several additional positively valenced descriptors (e.g., “strong,” “exciting,” “innovative”) is sufficient

these specifications in the main text as the fixed effects structure in this setting absorbs substantial identifying variation in arbitrage constraints.

to generate sentiment changes of this magnitude.

The sentiment effect remains present after absorbing episode-specific conditions. Column (4) includes air-date and industry fixed effects, which control for common shocks and episode-level conditions shared across industries. Even in this specification, sentiment retains explanatory power, with a coefficient of 5.5 basis points per unit of sentiment ($p < 0.05$). This persistence suggests that pitch tone contains information beyond general episode-level variation and is consistent with the model's prediction that halo-biased belief formation operates through heterogeneous perceived associations rather than through common information arrival.

By contrast, economically interpretable deal characteristics—such as whether a deal was made on air, the level of agreed valuations, or valuation revisions during negotiation—do not systematically predict short-run industry returns. In the context of *Shark Tank*, where the focal firms are private startups and the broadcast content is unlikely to convey material information about publicly traded firms' cash flows, this pattern is consistent with price responses being driven primarily by narrative tone rather than by economically informative signals embedded in deal terms. We also observe that sentiment effects are concentrated in next-day returns rather than overnight returns. While we do not take a strong stance on the precise timing of belief updating, this pattern is consistent with price responses emerging shortly after episodes air rather than being instantaneously incorporated at the market close.

Taken together, Table 2 provides evidence of short-run, sentiment-driven category spillovers. Investors appear to extrapolate positive impressions formed during compelling *Shark Tank* pitches to publicly traded firms in the same thematic category, despite the absence of new information about those firms' cash flows. This behavior mirrors the halo mechanism described in the model: salient narratives generate demand shocks at the category level, producing economically meaningful price distortions.

Finally, Table 3 provides a nonparametric assessment of the sentiment–return relationship by sorting next-day industry returns into quintiles based on pitch sentiment. This approach avoids functional-form assumptions and offers a transparent view of how returns vary across the sentiment distribution. Results display a broadly monotonic pattern. Industries associated with the lowest-sentiment pitches experience negative next-day abnormal returns of

approximately 25 basis points, while industries associated with the highest-sentiment pitches earn positive abnormal returns of roughly 17 basis points. Intermediate quintiles lie between these extremes, with mean returns generally increasing in sentiment. The spread between the highest- and lowest-sentiment quintiles is approximately 42 basis points, indicating economically meaningful variation across the sentiment distribution.

Although individual quintile estimates are noisy, the overall ordering of returns suggests that the halo effect scales with the strength of the sentiment shock. This pattern aligns with the linear structure implied by equation (8), under which mispricing is proportional to the magnitude of the sentiment innovation ($s_A - \mu$). Rather than being driven by a small number of extreme observations, the results indicate that more positive narratives are systematically associated with higher short-run returns across the distribution.

Taken together, the regression evidence and nonparametric sentiment sorts indicate that pitch sentiment generates economically meaningful and graded category-level price responses. The monotonic relation between sentiment intensity and next-day returns supports a halo-based mechanism in which investor extrapolation scales smoothly with the strength of the narrative signal. Having established the existence and structure of these sentiment spillovers, we next examine whether such effects persist or unwind over time, as predicted by the model. In the Appendix, we report several complementary analyses that further probe the scope of the halo effect.⁴

6.2 Return Dynamics and Long-Horizon Abnormal Performance

Hypothesis 2 of the model predicts that halo-induced price responses are transitory. Because category-level price pressure arises from belief extrapolation rather than from information about peer firms' fundamentals, equation (16) implies that sentiment-driven return effects should attenuate as beliefs correct and uncertainty is resolved. A necessary implication of this mechanism is that *Shark Tank*-related price responses should decay over short horizons and should not generate persistent abnormal performance.

⁴Appendix Table A3 allows sentiment effects to vary by individual Shark; results are qualitatively similar. Appendix Table A4 uses a narrow firm-level match and finds little evidence of price responses, consistent with halo effects operating through coarse categorical beliefs rather than precise firm mapping.

We begin by examining the short-run dynamics of sentiment-induced price responses. Figure 2 traces the evolution of the estimated effect of pitch sentiment on matched-industry returns over subsequent trading days, with the event day defined as the first trading day on or after the episode air date. The figure shows that the initial return response attenuates rapidly. Estimated coefficients decline in magnitude over the days following the broadcast and lose statistical precision, with no evidence of persistent effects beyond the immediate short run. This pattern is consistent with the resolution of halo-induced mispricing rather than with gradual incorporation of fundamental information, and closely mirrors the price dynamics implied by equations (8) and (16).

Having established that sentiment-driven return effects dissipate quickly, we next evaluate whether any residual abnormal performance remains over longer horizons. Specifically, we estimate Fama–French three-factor regressions on portfolios formed around episode air dates. Table 4 reports estimated portfolio alphas over 50-, 100-, 150-, 200-, and 250-day holding horizons for three portfolio constructions: (i) an equally weighted portfolio of all matched industries (Panel A), (ii) a deal-weighted portfolio that emphasizes industries associated with pitches receiving an on-air investment (Panel B), and (iii) a long–short portfolio that is long industries matched to high-sentiment pitches and short industries matched to low-sentiment pitches (Panel C).

Across all horizons and portfolio constructions, estimated alphas are statistically indistinguishable from zero. This absence of persistent abnormal performance stands in sharp contrast to the economically meaningful short-run return responses documented in the previous section, even for portfolios designed to explicitly exploit sentiment variation. The absence of long-horizon alpha plays an important diagnostic role. If *Shark Tank* pitches conveyed economically relevant information about the future cash flows or risk exposures of publicly traded firms, such information would be expected to manifest in persistent abnormal performance. Instead, the results indicate that the return effects associated with pitch sentiment dissipate fully over time, consistent with the model’s characterization of halo-driven mispricing rather than rational information transmission.

Accordingly, this analysis serves a dual purpose. First, it directly tests the model’s prediction that halo-induced price responses unwind as firm-specific information arrives and

beliefs adjust. Second, it provides a validation check on the empirical setting itself: the lack of long-horizon alpha supports the interpretation that *Shark Tank* functions as a source of salient sentiment shocks rather than as a channel through which durable fundamental information enters public markets.

Taken together with the rapid attenuation of returns documented earlier, the Fama–French results reinforce the view that *Shark Tank* airings primarily shift investor attention and affect in the short run. This pattern aligns closely with the price dynamics implied by equations (8) and (16), and underscores the suitability of the setting for isolating belief-based mechanisms from changes in intrinsic value.

6.3 Limits to Arbitrage

We next examine whether the short-run return responses documented above are amplified in settings where arbitrage is more costly or constrained, as implied by the model. In Section 3, halo-driven mispricing persists at $t = 1$ because rational traders face uncertainty about peer firms’ fundamentals and are unable or unwilling to fully offset sentiment-driven demand pressure. A central implication is that halo effects should be strongest precisely where arbitrage frictions are most severe.

To test this mechanism, we interact pitch sentiment with two industry-level proxies for limits to arbitrage: illiquidity and trading volume. These measures capture frictions that directly impede arbitrage activity. Table 5 reports regressions of next-day industry returns on pitch sentiment, the arbitrage-friction measures, and their interactions.

Across specifications, we find consistent evidence that sentiment-driven return effects are amplified in industries that are more difficult to arbitrage. The interaction between pitch sentiment and illiquidity is negative and statistically significant, indicating that the marginal effect of sentiment on returns is larger in more illiquid industries. Similarly, the interaction between sentiment and industry trading volume is negative and significant, implying that sentiment-induced price pressure is stronger when industry-level trading activity is lower. In both cases, the sign and magnitude of the interaction terms indicate that limits to arbitrage magnify the impact of sentiment shocks on prices.

These results align closely with the structure of the model. In equation (8), the magnitude

of mispricing depends not only on the strength of the sentiment shock ($s_A - \mu$) and the prevalence of halo-biased investors, but also on the ability of rational traders to counteract biased demand. When liquidity is low or trading capacity is limited, arbitrageurs face higher costs and greater risk in correcting mispricing, allowing sentiment-driven price deviations to persist more strongly in the short run.

Importantly, the amplification effects we document operate through trading frictions rather than through economically interpretable information channels. The main effects of deal characteristics and valuation terms remain weak, while the interaction results point squarely toward a behavioral mechanism constrained by market frictions. This pattern is difficult to reconcile with rational information transmission, but is a natural implication of the limits-to-arbitrage framework of Shleifer and Vishny (1997) applied to category-level belief spillovers. These findings reinforce the interpretation that *Shark Tank*-related return responses reflect sentiment-induced demand shocks that are only partially arbitrated away. Where arbitrage is easiest, price effects are muted; where it is most difficult, halo-driven mispricing is most pronounced. This cross-sectional variation strengthens the case that the observed return patterns at least partially arise from behavioral belief formation interacting with trading constraints.

7 Conclusion

This paper studies how salient narratives shape investor beliefs and asset prices through the halo effect: a cognitive bias in which a favorable (or unfavorable) impression of a focal object contaminates evaluations of related but economically distinct objects. While the halo effect has been extensively documented in psychology, marketing, and consumer behavior, its implications for belief formation and price dynamics in financial markets have remained largely unexplored. We develop a simple theoretical framework in which sentiment shocks to a focal firm induce category-level demand spillovers among loosely associated peer firms, generating short-run mispricing that unwinds as beliefs correct.

The model delivers three central predictions. First, sentiment shocks generate immediate price comovement among categorically related assets even in the absence of shared funda-

mentals. Second, these effects are temporary and reverse as information about peer firms' fundamentals is revealed. Third, mispricing is amplified when arbitrage is constrained, allowing sentiment-driven demand pressure to persist in the short run. We further test these predictions empirically.

To test our predictions, we exploit *Shark Tank* as a natural laboratory for studying halo-driven belief formation. The setting is uniquely well suited for this purpose. Televised pitches generate discrete, emotionally salient sentiment shocks that are widely observed, occur at well-defined points in time, and concern private startups whose prospects are unlikely to convey material information about the cash flows or risk exposures of publicly traded firms. This structure allows us to study belief spillovers in an environment where rational information transmission is implausible, sharpening interpretation of observed price responses.

Using a custom large language model to extract structured data from 1,286 *Shark Tank* pitches, we map each pitch to thematically related public industries and examine industry-level stock returns around episode air dates. We find that pitch sentiment is a strong and robust predictor of next-day abnormal industry returns. A one-unit increase in sentiment is associated with approximately 5–10 basis points of abnormal return, and moving from the lowest to highest sentiment quintile corresponds to return differences of roughly 40 basis points over a single trading day. These effects are economically meaningful, particularly given the short horizon and the absence of any new fundamental information about the affected industry.

Consistent with the model, these return effects are short lived. Dynamic event-time estimates show that sentiment-induced price pressure attenuates rapidly over subsequent trading days and fully dissipates within a week. Moreover, Fama–French factor regressions reveal no evidence of persistent alpha across a range of holding horizons and portfolio constructions, including portfolios that explicitly exploit sentiment variation. The absence of long-horizon abnormal performance further serves a diagnostic role: confirming that *Shark Tank* operates as a sentiment shock rather than a channel of fundamental information transmission, validating the setting as a clean environment for studying belief-based mispricing.

We further show that sentiment spillovers are amplified in industries where arbitrage is more difficult. Interactions between pitch sentiment and industry illiquidity or trading

volume indicate that price responses are stronger when trading frictions limit the ability of rational investors to quickly counteract sentiment-driven flows. This pattern aligns closely with limits-to-arbitrage theories and reinforces the interpretation that the observed return effects arise from temporary demand pressure rather than rational updating.

Together, this evidence supports a halo-based view of category-level belief formation in financial markets. Investors appear to extrapolate positive impressions formed during compelling narratives to broadly related public firms, even when such generalization lacks a fundamental basis. These belief spillovers produce economically meaningful but fragile price distortions that unwind as sentiment fades and informed capital responds. More broadly, the findings highlight how entertainment and media-driven narratives can shape asset prices through psychological channels. While our chosen context is *Shark Tank* for econometric cleanliness, the mechanisms we document are likely relevant in many settings where salient stories capture investor attention. Any number of highly visible events may generate similar halo-driven spillovers within asset categories.

Understanding these dynamics is important for both researchers and practitioners. For academics, the results underscore the value of integrating cognitive biases into models of belief formation and price dynamics. For market participants, the findings suggest that sentiment shocks can create predictable, short-lived price pressure without conveying durable information about value. More generally, this paper contributes to a growing literature illustrating that how information is framed and perceived can matter as much as the information itself for asset prices.


References

- Aldatmaz, S. and Brown, G. W. (2020). Private equity in the global economy: Evidence on industry spillovers. *Journal of corporate finance*, 60:101524.
- Barber, B. M., Huang, X., Odean, T., and Schwarz, C. (2022). Attention induced trading and returns: Evidence from robinhood users. *Journal of Finance*, 77(6):3141–3190.
- Barber, B. M. and Odean, T. (2008). All that glitters: The effect of attention and news

- on the buying behavior of individual and institutional investors. *The review of financial studies*, 21(2):785–818.
- Barberis, N., Shleifer, A., and Wurgler, J. (2005). Comovement. *Journal of Financial Economics*, 75(2):283–317.
- Bernstein, S., Lerner, J., Sorensen, M., and Strömberg, P. (2017). Private equity and industry performance. *Management Science*, 63(4):1198–1213.
- Calderon, N., Reichart, R., and Dror, R. (2025). The alternative annotator test for llm-as-a-judge: How to statistically justify replacing human annotators with llms. *arXiv preprint arXiv:2501.10970*.
- Edmans, A., Garcia, D., and Norli, Ø. (2007). Sports sentiment and stock returns. *The Journal of finance*, 62(4):1967–1998.
- Engelberg, J., Sasseville, C., and Williams, J. (2012). Market madness? the case of mad money. *Management Science*, 58(2):351–364.
- Fang, L. and Peress, J. (2009). Media coverage and the cross-section of stock returns. *The journal of finance*, 64(5):2023–2052.
- Jang, W.-Y., Lee, J.-H., and Hu, H.-C. (2016). Halo, horn, or dark horse biases: Corporate reputation and the earnings announcement puzzle. *Journal of Empirical Finance*, 38:209–226.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3):430–454.
- Lepori, G. M. (2015). Investor mood and demand for stocks: Evidence from popular tv series finales. *Journal of Economic Psychology*, 48:33–47.
- Mullainathan, S., Schwartzstein, J., and Shleifer, A. (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics*, 123(2):577–619.
- Nisbett, R. E. and Wilson, T. D. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology*, 35(4):250–256.

- Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of finance*, 62(3):1139–1168.
- Thorndike, E. L. (1920). A constant error in psychological ratings. *Journal of Applied Psychology*, 4(1):25–29.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

Figure 1: Mapping Shark Tank Beverage Pitches to Industry Portfolios. This figure illustrates the matching process from individual Shark Tank pitches to industry-level portfolios. Examples include Poppi Beverages (Season 10, Episode 8), Copper Cow Coffee (Season 12, Episode 25), and Little Saints Mocktails (Season 16, Episode 1). Pitch transcripts are compared with public firm descriptions using cosine similarity, and matched to an economically meaningful industry grouping, here SIC 2080 (Beverages).



This figure illustrates the matching process from individual Shark Tank pitches to industry-level portfolios. Examples include Poppi Beverages (Season 10, Episode 8), Copper Cow Coffee (Season 12, Episode 25), and Little Saints Mocktails (Season 16, Episode 1). Pitch transcripts are compared with public firm descriptions using cosine similarity, and matched to an economically meaningful industry grouping, here SIC 2080 (Beverages).

Figure 2: Coefficient of Pitch Sentiment on Matched Industry Returns after Air Date.

This figure plots estimated coefficients on average pitch sentiment from a sequence of regressions examining forward matched-industry stock returns following the airing of a Shark Tank episode. The horizontal axis denotes the number of trading days after the episode's event trading day (defined as the first trading day on or after the air date). Each point represents the estimated coefficient on average pitch sentiment for the corresponding forward return horizon, measured in percentage points. Vertical bars denote 95% confidence intervals. All regressions include the full set of baseline controls used in the main analysis, including lagged industry returns and market returns. The dashed horizontal line indicates zero. The figure illustrates a short-run pricing response to pitch sentiment that attenuates over subsequent trading days.




Table 1: Descriptive Statistics

Variable	N	Mean	SD	P25	Median	P75
<i>Pitches That Received a Deal</i>						
Ask Amount (\$)	772	263,149	309,757	100,000	200,000	300,000
Ask Equity (%)	768	12.7	7.76	8	10	15
Deal Amount (\$)	772	292,797	359,468	100,000	200,000	321,250
Deal Equity (%)	763	22.9	14.9	12	20	30
Royalty Deal (Indicator)	773	0.075	0.264	0	0	0
Debt Deal (Indicator)	773	0.049	0.216	0	0	0
Ask Valuation (\$)	768	3,494,992	5,362,018	750,000	1,708,333	4,000,000
Deal Valuation (\$)	763	2,303,185	3,624,387	452,273	1,000,000	2,500,000
Deal – Ask Valuation (\$)	759	-1,195,464	3,133,129	-1,500,000	-500,000	-50,000
Sentiment Score	773	1.37	2.17	0	1	3
<i>Pitches That Did Not Receive a Deal</i>						
Ask Amount (\$)	510	317,875	414,592	100,000	200,000	400,000
Ask Equity (%)	505	14.8	9.52	10	10	20
Ask Valuation (\$)	505	3,643,504	5,534,161	625,000	1,500,000	4,000,000
Sentiment Score	510	0.73	2.18	0	0	2

This table reports pitch-level summary statistics from *Shark Tank* episode transcripts, reported separately for pitches that receive a deal and those that do not. *Ask Amount* is the dollar amount requested by entrepreneurs, and *Ask Equity* is the equity share offered (percent). *Deal Amount* and *Deal Equity* are the agreed investment and equity terms for pitches that receive a deal. *Royalty Deal* and *Debt Deal* are indicators for whether the final agreement includes a royalty or debt component. *Ask Valuation* and *Deal Valuation* are implied valuations at the ask and deal terms; *Deal – Ask Valuation* is the change in implied valuation (deal minus ask), reported only for deals. *Sentiment Score* is computed using the AFINN lexicon (term scores from -5 to $+5$) and is measured as the average sentiment across words in the pitch.

Table 2: Effect of Pitch Characteristics on Subsequent Industry Returns (in Percentage Points)

Dependent Variable:	Overnight Ret	Next-Day Ret	Overnight (FE)	Next-Day (FE)
Avg. Sentiment	-0.046 (0.029)	0.100** (0.034)	-0.061+ (0.031)	0.055* (0.024)
Deal Made	3.950*** (0.994)	0.064 (1.153)	4.246 (5.525)	-1.264 (0.786)
Mean Deal Valuation (\$M)	-0.007 (0.011)	-0.000 (0.013)	0.003 (0.010)	-0.013 (0.009)
Avg. Valuation Change (\$M)	0.007 (0.011)	0.000 (0.013)	-0.003 (0.010)	0.013 (0.009)
# Pitches	0.252 (0.197)	-0.269 (0.229)	0.162 (0.204)	-0.108 (0.149)
# Deals	-0.197 (0.297)	0.174 (0.345)	0.223 (0.354)	0.343 (0.303)
Lagged Industry Return	0.026 (0.042)	-0.075 (0.049)	0.033 (0.042)	0.015 (0.053)
Market Return (S&P 500)	0.528*** (0.058)	-0.433*** (0.067)		
Industry Size (\$B)	0.009 (0.008)	0.004 (0.009)	0.059 (0.046)	-0.008 (0.013)
Illiquidity (log Amihud)	-196.427 (794.582)	3832.628*** (922.055)	900.056 (567.161)	248.085 (497.720)
Idiosyncratic Volatility	2.187 (6.522)	-9.259 (7.569)	-6.796 (10.595)	-17.006 (19.762)
Num.Obs.	686	686	686	686
R2	0.138	0.095	0.698	0.863
R2 Within			0.064	0.027
Std.Errors	IID	IID	Air Date	Air Date

This table reports regressions relating *Shark Tank* pitch characteristics to subsequent matched-industry stock returns around episode airings. Industries are matched to pitches using large language model (LLM)-based classification and aggregated to the episode $\times$ industry level. The unit of observation is an industry (SIC-based) $\times$ event trading day, where the event trading day is defined as the first trading day on or after the episode air date. The dependent variables are industry overnight returns (close-to-open) and next-day returns (close-to-close), measured in percentage points. *Deal Made* is an indicator equal to one if at least one pitch in the episode-industry cell results in a deal. *Avg. Sentiment* is the mean AFINN sentiment score across pitches in the episode-industry cell. Lagged industry returns are prior-day close-to-close returns, and market returns are measured using the S&P 500 return. Columns (3) and (4) include air-date and industry fixed effects; market-wide returns and time-invariant industry characteristics are absorbed by these fixed effects and therefore omitted. Standard errors are clustered by air date in fixed-effects specifications and assumed IID otherwise. $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 3: Next-Day Industry Returns Sorted by Pitch Sentiment

Sentiment Quintile	<i>N</i>	Mean Sentiment	Mean Next-Day Return	Std. Dev.	Std. Error	<i>t</i> -stat
Q1 (Lowest)	134	-1.98	-0.249	2.93	0.253	-0.98
Q2	134	0.52	-0.122	1.68	0.145	-0.84
Q3	133	1.68	-0.257	1.47	0.128	-2.01
Q4	133	2.59	-0.116	1.88	0.163	-0.71
Q5 (Highest)	133	4.83	0.167	1.64	0.142	1.17
Top – Bottom			0.416		0.290	1.43

This table reports average next-day industry returns sorted into quintiles based on the average sentiment of *Shark Tank* pitches associated with each industry–episode observation. Sentiment is measured using the AFINN lexicon and averaged across pitches within an episode–industry cell. Next-day returns are measured in percentage points. Quintiles are formed using the full sample of episode–industry observations. The Top–Bottom row reports the difference in mean next-day returns between the highest- and lowest-sentiment quintiles. *t*-statistics are computed using standard errors based on the cross-sectional dispersion of returns within each quintile and are intended for descriptive comparison.

Table 4: Fama–French Regression Results for Portfolio Alpha Across Specifications and Holding Horizons

Holding Period:	50-day	100-day	150-day	200-day	250-day
<i>Panel A: All Products</i>					
Alpha	-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Market (Mkt–RF)	0.944*** (0.022)	0.912*** (0.021)	0.907*** (0.020)	0.904*** (0.019)	0.902*** (0.019)
Size (SMB)	0.744*** (0.055)	0.737*** (0.043)	0.727*** (0.040)	0.732*** (0.038)	0.727*** (0.037)
Value (HML)	0.277*** (0.041)	0.271*** (0.035)	0.257*** (0.036)	0.267*** (0.033)	0.272*** (0.033)
Observations	2300	2877	3086	3250	3350
R^2	0.824	0.843	0.853	0.862	0.867
Std. Errors	Newey–West	Newey–West	Newey–West	Newey–West	Newey–West
<i>Panel B: Deal Portfolios</i>					
Alpha	-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Market (Mkt–RF)	0.944*** (0.012)	0.912*** (0.009)	0.907*** (0.009)	0.904*** (0.008)	0.902*** (0.008)
Size (SMB)	0.744*** (0.020)	0.737*** (0.017)	0.727*** (0.016)	0.732*** (0.015)	0.727*** (0.014)
Value (HML)	0.277*** (0.015)	0.271*** (0.012)	0.257*** (0.012)	0.267*** (0.011)	0.272*** (0.011)
Observations	2300	2877	3086	3250	3350
R^2	0.824	0.843	0.853	0.862	0.867
Std. Errors	Newey–West	Newey–West	Newey–West	Newey–West	Newey–West

Continued on next page

Table 4 continued

	50-day	100-day	150-day	200-day	250-day
<i>Panel C: Long–Short Sentiment Portfolios</i>					
Alpha	0.000 (0.001)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
Market (Mkt–RF)	-0.200*** (0.045)	-0.164*** (0.030)	-0.123*** (0.025)	-0.128*** (0.022)	-0.119*** (0.021)
Size (SMB)	-0.555*** (0.079)	-0.294*** (0.053)	-0.255*** (0.044)	-0.215*** (0.038)	-0.191*** (0.035)
Value (HML)	-0.225*** (0.056)	-0.160*** (0.039)	-0.162*** (0.033)	-0.167*** (0.030)	-0.188*** (0.027)
Observations	1021	1651	2118	2453	2677
R^2	0.108	0.065	0.055	0.056	0.055
Std. Errors	Newey–West	Newey–West	Newey–West	Newey–West	Newey–West

This table reports Fama–French three-factor regressions estimating portfolio alpha across five holding horizons (50 to 250 trading days) under three portfolio constructions. The dependent variable in all models is the daily excess return (portfolio return minus the risk-free rate). Portfolios are constructed following the airing of *Shark Tank* episodes and are formed using overlapping event cohorts, with constituent industries entering on the first trading day on or after the episode air date and remaining in the portfolio for the specified holding horizon. *Panel A* reports results for equally weighted portfolios of all industries appearing on the show. *Panel B* reports results for portfolios composed of industries associated with pitches that receive a deal. *Panel C* reports results for long–short sentiment portfolios that are long industries associated with high-sentiment pitches and short industries associated with low-sentiment pitches. Regressors include the excess market return (Mkt–RF), the Fama–French size factor (SMB), and the value factor (HML). Standard errors (reported in parentheses) are Newey–West adjusted to account for serial correlation induced by overlapping portfolio holding periods. Statistical significance is denoted as: $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, and $^{***}p < 0.001$.

Table 5: Interaction of Pitch Sentiment and Limits to Arbitrage on Next-Day Industry Returns

Dependent variable:	Next-Day Industry Return (%)		
Avg. Sentiment	0.327** (0.102)	0.396*** (0.117)	0.493*** (0.128)
Illiquidity Quartile	0.275** (0.099)		0.212* (0.101)
Industry Volume Quartile		0.418** (0.140)	0.351* (0.144)
Size Quartile	-0.042 (0.084)	-0.156 (0.108)	-0.180+ (0.108)
Idiosyncratic Volatility Quartile	0.021 (0.081)	-0.008 (0.083)	-0.016 (0.083)
Sentiment $\times$ Illiquidity	-0.085* (0.036)		-0.061 (0.040)
Sentiment $\times$ Industry Volume		-0.111** (0.042)	-0.087* (0.047)
Deal Made	-0.305 (1.243)	-0.329 (1.242)	-0.362 (1.240)
Lagged Industry Return	0.142** (0.046)	0.129** (0.047)	0.129** (0.046)
Market Return (S&P 500)	-0.351*** (0.077)	-0.363*** (0.076)	-0.348*** (0.077)
Mean Deal Valuation (\$M)	-0.006 (0.015)	-0.005 (0.015)	-0.008 (0.015)
Avg. Valuation Change (\$M)	0.006 (0.015)	0.005 (0.015)	0.008 (0.015)
# Pitches	-0.266 (0.274)	-0.234 (0.272)	-0.260 (0.273)
# Deals	0.256 (0.420)	0.274 (0.419)	0.263 (0.418)
Observations	544	544	544
R^2	0.102	0.106	0.113
Std. Errors	IID	IID	IID

This table reports pooled OLS regressions estimating whether the effect of *Shark Tank* pitch sentiment on next-day matched-industry returns varies with limits to arbitrage. The unit of observation is an industry (SIC) $\times$ event trading day, where the event trading day is the first trading day on or after the episode air date. The dependent variable is the industry next-day return (close-to-close), measured in percentage points. *Avg. Sentiment* is the mean AFINN sentiment score across pitches within an episode-industry cell. Limits-to-arbitrage measures are constructed as prior-month cross-sectional quartiles within each month: *Illiquidity Quartile* (Amihud illiquidity), *Industry Volume Quartile* (industry trading volume), *Size Quartile*, and *Idiosyncratic Volatility Quartile*. Columns (1)–(2) include one interaction at a time; Column (3) includes both interactions jointly. Standard errors are assumed IID. $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table A1: Profiles of Primary Sharks

Shark	Domain of Expertise / Notable Background
Mark Cuban	Technology entrepreneur; owner of Dallas Mavericks; venture investor
Lori Greiner	Retail and product design; inventor with over 100 patents
Daymond John	Founder of FUBU (fashion); branding and apparel investor
Kevin O’Leary	Venture capitalist; background in software and consumer products
Barbara Corcoran	Real estate entrepreneur; small business investor
Robert Herjavec	Cybersecurity entrepreneur; founder of Herjavec Group

A Custom GPT for Transcript Parsing

This appendix provides technical details on the construction and use of a custom large language model (LLM) for extracting structured data from *Shark Tank* transcripts.

A.1 What is a Custom GPT?

Custom GPTs are specialized large language models built on top of OpenAI’s GPT architecture and fine-tuned for domain-specific tasks. Unlike general-purpose GPT models, which are trained broadly on internet-scale data, custom GPTs incorporate tailored instructions, examples, and feedback to optimize performance for a well-defined problem. This makes them particularly well-suited for extracting structured information from unstructured text, such as television transcripts, where consistent annotation is required.

A.2 Suitability for Transcript Parsing

Financial research often relies on structured datasets (e.g., CRSP, Compustat, IBES), but unstructured text such as transcripts poses unique challenges. Traditional natural language processing (NLP) methods—based on rules, bag-of-words, or shallow machine learning—tend to struggle with conversational, multi-speaker formats. By contrast, GPT-based models leverage contextual embeddings and few-shot learning, enabling them to capture nuances in speaker attribution, deal terms, and valuation details.

In this setting, a custom GPT is especially valuable for three reasons: (1) Accuracy. GPT models can achieve human-comparable accuracy in information extraction, as shown by recent work (Calderon et al., 2025). (2) Scalability. Once tuned, the model can parse hundreds of episodes quickly, dramatically reducing the cost relative to manual annotation. (3) Flexibility. Prompts and instructions can be refined iteratively to handle edge cases (e.g., multiple Sharks making separate offers, royalty-based deals).

A.3 Construction of the Custom GPT

The specialized GPT was created using OpenAI’s Custom GPT Builder, with the following workflow:

1. **Manual Labeling:** A subset of episodes was annotated by hand to create a high-quality training set.
2. **Supervised Fine-Tuning:** The model was trained to extract structured variables, including entrepreneur names, product descriptions, deal terms, valuations, and Shark participation.
3. **Reinforcement Learning with Human Feedback (RLHF):** Outputs were iteratively evaluated, with corrections fed back into the model to refine its weights and instructions.
4. **Prompt Optimization:** The extraction prompt was refined in stages—testing on small samples, identifying common errors, and adjusting instructions—until consistent accuracy was achieved.
5. **Scaling:** The final optimized GPT was applied across all episodes, yielding structured data on 1,286 individual pitches.

A.4 Access and Final Prompt

The custom GPT used in this paper is available at the following link: <https://chatgpt.com/g/g-678fcea6c9d08191ae67eb099555bb92-transcript-analyzer>

The full extraction prompt used in the final trained GPT is given below:

I need you to extract data from shark tank transcripts. I will upload the transcripts, then you will extract the following data: Pitch Number, Company Name, Entrepreneur(s), Equity Ask, Deal Status, Deal Details, Sharks Involved in Negotiations, Shark Who Secured the Deal, Additional Offers, Product Name, Product Category, Pitch Summary Season/Episode. Be sure that you extract data for all the pitches contained within each transcript. There should be roughly 4 pitches per episode, and we need to capture all of them. Output the extracted data one episode at a time, do not add data together from episode to episode. This data will be used in a research project, so accuracy of the extracted data is paramount. Are you ready to start receiving transcripts?

B Variable Definitions

This appendix provides definitions of all variables used in the analysis, grouped by type. Variables are drawn from *Shark Tank* transcripts, CRSP/Compustat, or the Fama–French data library.

Table A2: Variable Definitions

Variable	Definition	Source
<i>Pitch-Level Variables</i>		
Ask Amount (\$1000)	Dollar amount (in thousands) requested by entrepreneur in exchange for equity.	Transcript
Ask Equity (%)	Percentage of equity offered by entrepreneur for Ask Amount.	Transcript
Deal Amount (\$1000)	Dollar amount (in thousands) agreed upon by Shark(s).	Transcript
Deal Equity (%)	Percentage of equity agreed upon by Shark(s) in final deal.	Transcript
Royalty Deal	Indicator = 1 if deal involved a royalty component.	Transcript
Debt Deal	Indicator = 1 if deal involved a debt component.	Transcript

Table A2 continued

Variable	Definition	Source
Ask Valuation (\$1000)	Implied firm valuation at ask stage, computed as Ask Amount / Ask Equity.	Transcript
Deal Valuation (\$1000)	Implied firm valuation at deal stage, computed as Deal Amount / Deal Equity.	Transcript
Deal – Ask Valuation (\$1000)	Difference between Deal Valuation and Ask Valuation.	Transcript
Sentiment Score	Average sentiment of pitch transcript using AFINN lexicon (–5 to +5).	Transcript
Deal Made	Indicator = 1 if any deal was struck with a Shark.	Transcript
<i>Episode-Level Variables</i>		
Pitches	Number of pitches in the episode.	Transcript
Deals	Number of successful deals in the episode.	Transcript
Shark Indicators	Binary indicators for each Shark (Cuban, Greiner, Daymond, Kevin, Barbara, Robert) and Guest Shark.	Transcript
<i>Market Outcome Variables</i>		
Overnight Return	Close-to-open return of matched industry/firm minus S&P 500 return.	CRSP/Compustat
Next-Day Return	Open-to-close return of matched industry/firm minus S&P 500 return.	CRSP/Compustat
Firm-Level Return	Individual firm return matched to pitch via cosine similarity.	CRSP/Compustat
Portfolio Excess Return	Return on portfolio of matched industries minus risk-free rate.	CRSP/Compustat
Alpha	Intercept from Fama–French regressions of portfolio returns.	Author calc.

Arbitrage Proxies

Table A2 continued

Variable	Definition	Source
Size	Industry quartile based on average market capitalization.	Compustat
Illiquidity	Amihud ratio: avg. absolute return / dollar volume.	CRSP
Idiosyncratic Volatility	Residual volatility from rolling factor model at industry level.	CRSP/Compustat
Industry Volume	Average daily trading volume of industry constituents.	CRSP
<i>Factor Model Variables</i>		
Mkt–RF	Market return minus risk-free rate.	Ken French Library
SMB	Fama–French size factor (small minus big).	Ken French Library
HML	Fama–French value factor (high minus low book-to-market).	Ken French Library

Figure 1: Time Variation in the Relationship Between Pitch Sentiment and Next-Day Industry Returns

This figure plots year-specific estimates of the coefficient on pitch sentiment from regressions of next-day industry returns on sentiment. For each calendar year, the point estimate represents the marginal effect of average pitch sentiment on next-day returns, estimated using a pooled specification that allows the sentiment coefficient to vary by year. Vertical bars denote 95% confidence intervals constructed using robust standard errors. The dashed horizontal line indicates a zero effect. While the precision of the estimates varies across years, the point estimates are predominantly positive throughout the sample period.




Table A3: Shark Identity Controls and Sentiment Interactions

	Overnight Return	Next-Day Return	Overnight + Interactions	Next-Day + Interactions
Deal Made	3.927*** (0.998)	0.108 (1.148)	3.856*** (0.996)	0.106 (1.157)
Avg. Sentiment	-0.046 (0.030)	0.102** (0.034)	-0.132* (0.066)	0.137* (0.077)
Mean Deal Valuation	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Avg. Valuation Change	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Pitches	0.255 (0.199)	-0.214 (0.229)	0.248 (0.198)	-0.259 (0.230)
Deals	-0.275 (0.314)	0.046 (0.362)	-0.261 (0.302)	0.207 (0.351)
Lagged Industry Return	0.024 (0.042)	-0.084+ (0.048)	0.025 (0.042)	-0.081+ (0.049)
Market Return (S&P 500)	52.140*** (5.804)	-44.842*** (6.681)	52.320*** (5.817)	-44.114*** (6.758)
Industry Size	0.000	0.000	0.000	0.000

Continued on next page

Table A3 continued

	Overnight Return	Next-Day Return	Overnight + Interactions	Next-Day + Interactions
	(0.000)	(0.000)	(0.000)	(0.000)
Illiquidity (Amihud)	-165.214	3823.855***	-131.235	3882.559***
	(800.910)	(921.880)	(799.481)	(928.861)
Idiosyncratic Volatility	2.111	-8.514	0.902	-9.677
	(6.552)	(7.542)	(6.568)	(7.630)
Shark: Cuban	-0.049	0.089		
	(0.167)	(0.192)		
Shark: Greiner	0.165	0.188		
	(0.173)	(0.199)		
Shark: Daymond	0.250	0.459*		
	(0.203)	(0.234)		
Shark: Kevin	-0.118	-0.305		
	(0.201)	(0.231)		
Shark: Barbara	0.106	-0.363		
	(0.211)	(0.243)		
Shark: Robert	0.091	0.231		
	(0.197)	(0.226)		

Continued on next page

Table A3 continued

	Overnight Return	Next-Day Return	Overnight + Interactions	Next-Day + Interactions
Guest Shark	0.340	0.079	0.361	-0.066
	(0.327)	(0.377)	(0.294)	(0.341)
Sentiment $\times$ Cuban			0.048	-0.076
			(0.068)	(0.079)
Sentiment $\times$ Greiner			0.067	0.036
			(0.067)	(0.078)
Sentiment $\times$ Daymond			0.143+	-0.010
			(0.078)	(0.090)
Sentiment $\times$ Kevin			0.006	-0.105
			(0.076)	(0.088)
Sentiment $\times$ Barbara			0.115	-0.081
			(0.084)	(0.098)
Sentiment $\times$ Robert			0.071	0.011
			(0.078)	(0.090)
Observations	686	686	686	686
R^2	0.145	0.116	0.146	0.100

Std. Errors	IID	IID	IID	IID
<i>Notes:</i> This table reports robustness regressions augmenting the baseline specifications with controls for individual shark identity and interactions between pitch sentiment and shark indicators. The unit of observation is an industry (SIC) $\times$ event trading day, where the event trading day is the first trading day on or after the episode air date. All specifications include the full set of pitch characteristics and market, size, liquidity, and volatility controls used in the main analysis. Columns (3) and (4) interact average pitch sentiment with shark identity indicators to test whether the pricing response to sentiment varies across individual sharks. Standard errors are assumed IID. $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.				

Table A4: Shark Tank Pitch Sentiment and Matched-Firm Returns

	Next-Day Return		Overnight Return	
	(1) Pooled	(2) Air-Date FE	(3) Pooled	(4) Air-Date FE
Pitch sentiment	0.054 (0.097)	-0.015 (0.230)	0.002 (0.044)	-0.012 (0.109)
Deal valuation (\$M)	0.091 (0.064)	0.093 (0.095)	0.044 (0.040)	-0.025 (0.059)
Valuation change (\$M)	0.149 (0.097)	0.026 (0.092)	0.029 (0.042)	-0.064 (0.052)
Market return (S&P 500)	-0.379 (0.676)		0.474** (0.170)	
Size (\$B)	0.017 (0.018)	0.010 (0.058)	0.009 (0.008)	-0.029 (0.030)
Illiquidity ($\log(1 + \text{Amihud})$)	-110,926 (178,777)	98,245 (512,011)	-148,539* (75,256)	-170,843 (244,771)
Idiosyncratic volatility	12.890 (29.265)	11.466 (51.064)	6.350 (11.489)	14.028 (20.722)
Observations	237	237	238	238
R^2	0.020	0.848	0.089	0.794
Within R^2		0.014		0.064
SEs clustered by air date	Yes	Yes	Yes	Yes

Notes: This table reports firm-level regressions examining the relation between *Shark Tank* pitch characteristics and short-run returns of the top matched public firm. The unit of observation is a matched firm-event, defined as a public firm (CUSIP) on the first CRSP trading day on or after the episode air date. Columns (1)–(2) use next-day close-to-close returns, and columns (3)–(4) use overnight returns. Pitch-level variables are aggregated to the firm-event level when multiple pitches map to the same firm on the same air date. Pooled specifications include the market return; air-date fixed effects absorb market-wide variation and therefore omit the market return. Standard errors are clustered by air date. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A5: Limits to Arbitrage Interactions (Air-Date Fixed Effects)

Dependent variable:	Next-Day Industry Return (%)		
	Illiquidity	Industry Vol.	Full
Avg. Sentiment	0.091 (0.106)	0.229* (0.097)	0.228* (0.114)
Illiquidity Quartile	0.078 (0.089)		0.025 (0.087)
Industry Volume Quartile		0.285** (0.100)	0.277** (0.100)
Size Quartile	0.004 (0.056)	-0.079 (0.060)	-0.081 (0.061)
Idiosyncratic Volatility Quartile	-0.063 (0.071)	-0.085 (0.070)	-0.084 (0.069)
Sentiment $\times$ Illiquidity	-0.021 (0.038)		0.006 (0.044)
Sentiment $\times$ Industry Volume		-0.074* (0.035)	-0.080* (0.041)
Deal Made	-1.979* (0.933)	-1.920* (0.886)	-1.892* (0.877)
Lagged Industry Return	0.141 (0.098)	0.147 (0.099)	0.149 (0.097)
Mean Deal Valuation (\$M)	-0.012 (0.008)	-0.013 (0.008)	-0.013 (0.008)
Avg. Valuation Change (\$M)	0.012 (0.008)	0.013 (0.008)	0.013 (0.008)
# Pitches	-0.089 (0.144)	-0.071 (0.139)	-0.086 (0.143)
# Deals	0.684* (0.305)	0.662* (0.287)	0.666* (0.290)
Observations	544	544	544
R^2	0.868	0.872	0.872
R^2 (Within)	0.060	0.088	0.089
Fixed Effects	Air date	Air date	Air date
Std. Errors	by air date	by air date	by air date

This table reports fixed-effects regressions estimating whether the relationship between pitch sentiment and next-day matched-industry returns varies with limits to arbitrage. The unit of observation is an industry (SIC) $\times$ event trading day, where the event trading day is the first trading day on or after the episode air date. The dependent variable is the next-day industry return (close-to-close), measured in percentage points. *Avg. Sentiment* is the mean AFINN sentiment score across pitches within an episode-industry cell. Limits-to-arbitrage measures are constructed as prior-month cross-sectional quartiles within each month, including *Illiquidity Quartile* (Amihud illiquidity) and *Industry Volume Quartile* (industry trading volume); Size and idiosyncratic volatility quartiles are included as controls. All specifications include air-date fixed effects; accordingly, market-wide variation is absorbed. Standard errors are clustered by air date. $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.