Finfluencers and the Consumification of Stock Investment

Lars Hornuf
Dresden University of Technology
lars.hornuf@tu-dresden.de

Matthias Mattusch
Dresden University of Technology
matthias.mattusch@tu-dresden.de

Abstract

Financial influencers (finfluencers) increasingly disseminate investment advice through short-form videos embedded in consumer and lifestyle content. We study whether their stock recommendations adopt to consumption context or merely replicate analyst advice. We develop a consumption-hedging view of finfluencer recommendations, arguing that retail investors may value shares of firms whose products they consume because equity ownership can hedge exposure to future increases in the prices of those products. Using a hand-collected sample of 544 explicit and reasoned TikTok stock recommendations from 98 finfluencers between 2023 and 2025, we test the predictions of the consumption-hedging view using a triple-difference design that exploits exogenous variation in consumption. We find that periods of higher consumption increase recommendations of consumer-goods firms by a factor of 2.02, generate greater engagement with finfluencers, and do not lead to inferior long-horizon performance relative to analyst recommendations. Our findings show that finfluencer investment advice provides value beyond information about expected returns by incorporating consumption objectives of retail investors.

Keywords: Financial Advice, Finfluencer, Consumification, Investment Decisions, Social Media

JEL Classification: G11, G41, D12, L82

1. Introduction

People increasingly spend their spare moments scrolling through short videos on their smartphones, whether while walking, commuting, or before going to sleep. Platforms like TikTok, driven by personalized algorithms, deliver rapidly changing visual content optimized to capture and sustain user attention (Lindström et al., 2021). TikTok alone reports over 200 million users in both Europe and the United States and over one billion globally as of 2026 (TikTok, 2025, 2026a, 2026b). A notable portion of TikTok content centers on both routine and high-end consumption, such as food and cosmetics, as well as aspirational luxury-fashion content (Hogsnes & Grønli, 2025; Yang et al., 2025; Zhang et al., 2025). These videos often create benchmarks for desirable lifestyles and foster a "fear of missing out" (FOMO), that is an anxiety that one's consumption may fall short of societal standards (FINRA & CFA Institute 2023; Przybylski et al. 2013).

Amid this context, short-form videos featuring stock recommendations present an intriguing anomaly for four reasons: (i) stock recommendations in feeds dominated by lifestyle, beauty, and entertainment content appear disconnected from other short videos,[1] (ii) most users likely do not expect stock recommendations in the format of short-videos,[2] (iii) most short videos are asynchronous, i.e. users may view them weeks after publication when financial information is outdated (Liu et al. 2022), and (iv) unlike analysts or semiprofessional analysts on platforms like StockTwits or Seeking Alpha, who cater to specialized audiences[3], finfluencers reach casual users who may have little exposure to financial topics (Guan, 2023; FINRA & CFA Institute, 2023). Finfluencers have also adapted strategies from consumer influencers such as personal storytelling, accessible language, and interactive formats to engage audiences with limited financial experience (Hayes & Ben-Shmuel, 2024). We describe this process as the *consumification of stock investment*.

The consumification of stock investment is the mirror image of financialization. Financialization turns consumption goods into objects of investment, usually through

---

[1] Popular topics range from #WhatToWear to #CookingHacks in 2025 (TikTok Newsroom, 2026b, 2026a).

[2] Two out of three TikTok user say they discover useful things on the platform beyond what they originally searched for (TikTok, 2026a, 2026b).

[3] StockTwits reaches approximately ten million semiprofessional analysts and therefore a substantially smaller audience (Stocktwits, 2025).

financial professionals. Consumification works in the opposite direction. Non-professional financial communicators like finfluencers present investment objects like stocks as relatable forms of future consumption. A consumer who envisions the future purchase of the latest smartphone, luxury car, and her favorite coffee brand may encounter financial risk if the prices of these goods rise. Without forward contracts to hedge against future consumption costs, finfluencers suggest an alternative hedge, by investing in the stocks of firms producing these goods. If the price of the good increases, the corresponding rise in the firm's stock price could offset the higher costs for the consumer.

To formalize this idea, we develop a multi-period partial equilibrium model incorporating recursive Epstein-Zin preferences and a gain-loss reference framework (Epstein & Zin, 1991, 2013; Guo & Dong He, 2017). The model introduces a "consumption-hedge channel," according to which consumers partially hedge future consumption costs by holding stocks linked to their desired goods. In line with the phenomenon of consumification, stock recommendations following the motto “invest in what you will buy” can serve a rational economic purpose. The consumption-hedge channel thus complements the classical financial information channel and requires a performance evaluation beyond traditional alphas.

Our empirical analysis examines 544 TikTok stock recommendations from 98 finfluencers between 2023 and 2025, using a triple-difference design to distinguish between financial and consumption-hedge motivations. Our results confirm that finfluencers disproportionately recommend consumer-linked firms during high-consumption periods (e.g., pre-Christmas), with recommendation rates increasing by 94.8% to 101.8%, depending on the specification. These patterns are absent among analysts, supporting the idea that finfluencers distinctively cater to consumers’ hedging needs. Finfluencer recommendations perform worse than analyst recommendations in the short term but outperform them over a one-quarter of a year horizon. This pattern suggests that finfluencers do not base their recommendations on a short-term informational advantage.

The stronger longer-term performance may partly reflect how we identify recommendations. Unlike studies that rely on sentiment classifications, such as Kakhbod et al. (2023), we consider only explicit and reasoned stock recommendations. This approach filters out noisy content that may not reflect a clear investment decision. At the same time, financial characteristics remain important in finfluencer recommendations. We identify a financial channel through which finfluencers favor stocks with high momentum, volatility, or upside potential, consistent with users' FOMO-driven preferences.

This article contributes to finance, consumer, and information systems research on finfluencers (Haase et al., 2026; Mardon et al., 2026; Mölders et al., 2025; Pandey et al., 2026; Zülch et al., 2026). First, it shifts the object of study from sentiments to explicit, actionable stock recommendations, which enables a more precise evaluation of finfluencer recommendations. Second, it shows that finfluencers are not merely analysts, but connect investment decisions with consumption-related objectives. Third, because information that becomes quickly outdated travels poorly through delayed video feeds, platforms like TikTok encourage finfluencers to produce more temporally robust, longer-term recommendations. This mechanism can to some extent align creator incentives with audience needs despite competing commercial and attention-based incentives (Qian & Jain, 2024; Yang et al., 2025).

By framing investment as a means of securing desired lifestyles, finfluencers encourage broader participation in financial markets, particularly among financially inexperienced users (FINRA & CFA Institute 2023; Mardon et al. 2026). While concerns about fraudulent practices, such as pump-and-dump schemes, remain significant (Rockwell & Conklin 2026; Zülch et al. 2026), our findings suggest that finfluencers also serve a valuable economic function, namely translating finance into a familiar and accessible framework for everyday users. Recognizing this function has implications for regulation, which should aim to foster a healthy equity culture without stifling innovation in financial communication (Securities and Exchange Commission, 2024).

The remainder of the paper is organized as follows: Section 2 contains our theoretical model and its predictions. In Section 2, we relate our research question to the existing literature. Section 3 describes the sample and variables. In Section 4, we detail our identification strategy and empirical model. Section 5 summarizes the empirical results. Section 6 concludes the paper.

2. Theoretical Model

2.1 Agent

Our theoretical model incorporates key economic features of finfluencer activity. It builds on the economic characteristics of conventional influencer activity by introducing consumption preferences and a reference level of desired consumption, relative to which consumption shortfalls are perceived as losses—an expression of FOMO. Finfluencers add a forward-looking, intertemporal dimension: risk preferences, loss aversion, and the smoothing of consumption over time enter the agent's utility considerations and decision problem. Combining these preferences with the available consumption and investment opportunities, as well as the link between them, gives rise to an endogenous portfolio choice and determines whether, and to what extent, the stock of the firm producing the desired good serves as an optimal consumption hedge. Throughout, we model finfluencers and their audiences as a single representative agent with aligned interests. Building on Guo and He (2017), we formalize this mechanism and derive empirically testable predictions.[4]

We consider an infinitely lived agent in a discrete-time economy, $t = 1,2,3, \ldots$, with two risky assets, $j \in \{S, M\}$. Asset S represents the single stock of interest, while asset M denotes the market numéraire. The agent chooses the consumption share $c_t \in [0,1]$ and the portfolio weight $\omega_t \in [0,1]$ allocated to asset S, with the remaining share invested in the market numéraire asset M. Borrowing and short-selling are not allowed. The agent takes as given the joint dynamics of asset returns and goods-price growth, which link investment returns to future consumption.

At each date t, the agent chooses c_t and ω_t to maximize her recursive lifetime utility $U_t(W_t, p_t)$, taking current nominal wealth W_t, and the current price vector $\boldsymbol{p_t}$ as given

$$\max_{c_t \in [0,1],\, \omega_t \in [0,1]} U_t(W_t, \boldsymbol{p_t}). \qquad (1)$$

[4] Our model draws on previous theoretical work. Relative to Guo & Dong He (2017), we abstract from the risk-free asset and model investor payoffs through consumption-based utility, while retaining the gain–loss logic of their framework. In the spirit of Constantinides (1990) and Abel (1990), utility is evaluated relative to a habitual reference level for consumption, which we transpose to consumption-hedging needs. Unlike the finfluencer model of Pedersen (2022), our results do not rely on beliefs or information channels in networks, but on preferences shaped by consumption motives.

The utility process is defined recursively by

$$U_t = \left[(1-\beta)\left(C_t^{real}\right)^{1-\rho} + \beta(\mathcal{M}_t(U_{t+1}) + b_t G_t)^{1-\rho}\right]^{\frac{1}{1-\rho}}, \qquad \rho \neq 1, \qquad (2)$$

with the usual log limit for $\rho = 1$. The aggregator combines current real consumption, $C_t^{real} = c_t W_t^{real}$, with a continuation value. This continuation value consists of the certainty equivalent of future utility, $\mathcal{M}_t(U_{t+1})$, and a gain–loss component G_t weighted by b_t. The certainty-equivalent operator is

$$\mathcal{M}_t(X) = (E_t[X^{1-\gamma}])^{\frac{1}{1-\gamma}}, \qquad \gamma \neq 1,$$

with log limit $\mathcal{M}_t(X) = \exp(E_t[\ln X])$ if $\gamma = 1$. The parameter $\beta \in (0,1)$ is the subjective discount factor, $\rho > 0$ is the inverse elasticity of intertemporal substitution, and $\gamma > 0$ is relative risk aversion in the certainty-equivalent operator. The parameter $\alpha \in [0,1]$ determines the agent's preference for the good of asset S produces.

The behavioral component evaluates deviations of next period's real consumption, $\tilde{C}_{t+1} = (1-c_t)W_t^{real}R_{t+1}^{real}$, from a reference level $C_{t+1}^{Ref} = \chi(1-c_t)W_t^{real}$, where $\chi > 0$ is the benchmark level for next-period consumption. We write

$$G_t = E_t\left[v\left(\tilde{C}_{t+1} - C_{t+1}^{Ref}\right)\right].$$

Real wealth is $W_t^{real} := \frac{W_t}{P_t}$ with $P_t := p_{S,t}^{\alpha} p_{M,t}^{1-\alpha}$, resembling the Cobb-Douglas indirect utilty function. The value function $v(x)$ captures the kink between gains and losses:

$$v(x) = x\mathbf{1}_{\{x \geq 0\}} + \kappa x \mathbf{1}_{\{x<0\}}, \qquad \kappa \geq 1. \qquad (3)$$

The parameter κ measures loss aversion, while $b_t \geq 0$ governs the overall weight of the gain–loss motive. In the stationary representation below, the scaled gain–loss weight is denoted by b.

2.2 Markets, Prices, and Real Wealth

The consumption goods have price dynamics $p_{j,t+1} = p_{j,t}\Pi_{j,t+1}$, where $\Pi_{j,t+1}$ is gross goods-price growth. Each risky asset has a gross return $R_{j,t+1}$. The one-period shock vector $\left(\Pi_{S,t+1}, \Pi_{M,t+1}, R_{S,t+1}, R_{M,t+1}\right)$ is i.i.d. multivariate lognormal, with nonzero correlation between goods-price growth and stock returns.

The agent's gross nominal portfolio return is

$$R_{p,t+1} = \omega_t R_{S,t+1} + (1 - \omega_t) R_{M,t+1}. \qquad (4)$$

Nominal wealth evolves according to $W_{t+1} = (1 - c_t) W_t R_{p,t+1}$. In real terms, the gross real portfolio return is $R_{t+1}^{real}(\omega_t) = \frac{R_{p,t+1}}{\Pi_{t+1}}$, where $\Pi_{t+1} = \frac{P_{t+1}}{P_t} = \Pi_{S,t+1}^{\alpha} \Pi_{M,t+1}^{1-\alpha}$. Real wealth therefore evolves as

$$W_{t+1}^{real} = (1 - c_t) W_t^{real} R_{t+1}^{real}(\omega_t). \qquad (5)$$

2.3 Portfolio Choice and Economic Channels

Under the homogeneity of the recursive aggregator and the piecewise linear gain–loss term, the portfolio choice separates from the consumption choice. We summarize the portfolio-dependent continuation payoff by the real portfolio valuation objective

$$\begin{aligned} \mathcal{J}_t(\omega; \alpha, b) \quad &:= \mathcal{M}_t\left(R_{t+1}^{real}(\omega)\right) + bE_t\left[v\left(R_{t+1}^{real}(\omega) - \chi\right)\right], \\ \omega_t^* \quad &\in \underset{\omega \in [0,1]}{\operatorname{argmax}} \, \mathcal{J}_t(\omega; \alpha, b). \end{aligned} \qquad (6)$$

Under the market numéraire normalization $p_{M,t} \equiv 1$ and $\Pi_{M,t+1} \equiv 1$,

$$R_{t+1}^{real}(\omega) = \frac{\omega R_{S,t+1} + (1 - \omega) R_{M,t+1}}{\Pi_{S,t+1}^{\alpha}}.$$

The objective $\mathcal{J}(\omega; \alpha, b)$ evaluates portfolios by their contribution to future real consumption.

The Appendix derives the log-linear channel representation used below. Let μ_S, μ_M, μ_π denote the means of log stock returns and log goods-price growth, and let $\sigma_S^2, \sigma_M^2, \sigma_\pi^2, \sigma_{SM}, \sigma_{S\pi}, \sigma_{M\pi}$ denote the corresponding variances and covariances.

For $b = 0$ and $\gamma > 1$, an interior optimum is in primitives

$$\boxed{\omega_0^* = \underbrace{\frac{\mu_S - \mu_M}{(\gamma - 1)(\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM})} - \frac{\sigma_{SM} - \sigma_M^2}{\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM}}}_{\text{financial channel}} + \underbrace{\alpha \frac{\sigma_{S\pi} - \sigma_{M\pi}}{\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM}}}_{\text{consumption-hedge channel}}} \qquad (7)$$

The first component is the financial demand for asset S: it increases in the expected relative log return of S and decreases in relative financial risk. The second component is consumption-hedging demand. If asset S's return covaries more strongly with its consumption price than the market numéraire asset M's return, then $\sigma_{S\pi} - \sigma_{M\pi} > 0$, and the agent demands more of asset S.

Expected stock returns and return risk constitute the financial channel, whereas the covariance between stock returns and the price of the consumption good constitutes the consumption-hedge channel.

For $b > 0$, ω^* is characterized by the full gain–loss FOC in Equation (13). The Appendix Table E.1 shows that the two channels are not knife-edge: for $\alpha > 0$, consumption-price exposure matters through real-return risk when $\gamma > 1$ or through gain–loss valuation when $\rho = \gamma = 1$ and $b > 0$. Gain–loss adds no third channel; it revalues the real payoff relative to χ, so consumption hedging disappears only for $\alpha = 0$ or $\rho = \gamma = 1,\ b = 0$.

Insert Figures 1

Figure 1 illustrates the full $b > 0$ solution numerically for the treated group, using the parameter vector $(\alpha, b, \kappa, \chi, \gamma) = (0.465, 0.273, 2.396, 0.980, 14.945)$, and the exogenous parameter vector $(\beta, \rho_{\text{EIS}}, \mu_r, \sigma_r, \mu_\pi, \sigma_\pi, R_M) = (0.99, 0.8, 0.07, 0.18, 0.02, 0.04, 1.05)$.[5] Panel A shows the central consumption-hedge mechanism: demand for asset S increases when its return is more strongly correlated with the price of the goods produced by S. Panel B shows the resulting portfolio share of the growth asset S over time. Panel C simulates the change in the growth path when the strength of the consumption hedge, ρ_{R_S,π_S}, increases. Panel D shows the change in the growth path when b increases. Further below, the latter two panels motivate our identification strategy. Comparing two firms, one of which has a closer consumer link via its end-consumer products, resembles a comparison between a low and a high ρ_{R_S,Π_S}. Similarly, comparing periods during which consumers are pulled to such firms in contrast to periods when this is not the case, corresponding to a comparison of high versus low β-components.

Model prediction: An agent who plans future consumption of a particular good will allocate a larger portfolio weight to the stock linked to its anticipated expenditures. FOMO strengthens this consumption link. In contrast, an agent who detaches herself from

[5] Parameter values stem from a calibration of the model outlined further below in Section 5.2.

consumption and focuses solely on financial matters will act exclusively on the basis of risk and expected return.

3. Data and Variables

To empirically test the theoretical model's prediction, we need a setting in which finfluencers and their audiences are sufficiently incentive-aligned and in which alternative explanations—such as speculation, news, fundamentals, attention, or stock visibility—do not confound. We begin with the selection of the dependent variable, the portfolio weight ω, then turn to explanatory variables and controls.

3.1 Finfluencers

We operationalize ω with stock recommendations, $Recommendation$, defined as the explicit overweighting of a stock relative to a passive broad-market numéraire. Because no existing database provides these recommendations, we construct the dataset through a dedicated collection and classification pipeline. The empirical setting must satisfy three conditions. First, the audience must encounter stock recommendations in a consumption framing in which FOMO is present. Second, users must have wealth and access to investable markets. Third, finfluencers must make credible, rational, and carefully considered recommendations that remain aligned with their audience and do not involve fraud. These restrictions align the recommendation setting with the rational model outlined above.

TikTok serves as the ideal platform for this study because of its strong consumption framing (Gibson et al., 2025; Hogsnes & Grønli, 2025; Vombatkere et al., 2024; Yang et al., 2025), where stock recommendations must compete with lifestyle and entertainment content. Many financially inexperienced users encounter stock recommendations on TikTok not because of an explicit interest in finance but because of the TikTok algorithm, which presents stock recommendations as part of broader consumption-oriented content. We retrieved all relevant data in May 2025 and collected TikTok short videos published between January 2023 and March 2025. This timeframe captures a mature financial-content ecosystem while avoiding distortions from the COVID-19 period. The focus is on finfluencers operating in developed markets, particularly Australia, Europe, and North America where consumer behavior, wealth, and access to retail and investable consumer brands align with the assumptions of our model.

To identify finfluencers, three independent researchers conducted hashtag-based and keyword-based searches. They combined separate searches for finance-related hashtags and the corresponding non-hashtag terms across several European languages.[6] For each query, the researchers manually screened every video displayed by the platform. When they identified a potentially relevant video, they examined the creator's profile and reviewed all publicly accessible historical short videos, including videos available in the archived collection. This profile-level review reduced reliance on the platform's ranking of individual search results and allowed the researchers to identify relevant content beyond the videos initially displayed.

To capture creators who occasionally published stock- or investment-related recommendations without primarily presenting themselves as stock-market creators, the researchers also extended the search to adjacent topics.[7] In addition, they used country-specific web searches and online lists mentioning finfluencers to identify further candidate accounts and screened the corresponding platform profiles and videos using the same procedure. They discontinued the search when an additional search round yielded no previously unidentified eligible videos. To ensure observable audience engagement, we retained only publicly accessible short videos with at least fifty recorded views. To reduce researcher-specific classification errors, the three researchers independently collected, watched, filtered, and labeled all publicly available videos posted by the identified finfluencers. Four months after the initial collection, we rechecked the availability of the profiles and videos and found no evidence that subsequent content removal had selectively affected the initial sample.

The analysis focuses on explicit stock recommendations that match the theoretical model's mechanism. Popular sentiment classifiers, which often capture general bullish or bearish

---

[6] The search terms included "#trading," "#stocks," "#stocks," "#investment," "#invest," "wealth," and "#finance," together with their translations into German, French, Spanish, Portuguese, Italian, Dutch, Polish, Swedish, Danish, Norwegian, and Finnish.

[7] The adjacent topics included cryptocurrency, Bitcoin, real estate, and business.

sentiment, do not capture the deliberate investment judgments studied here.[8] The sample therefore excludes ETFs, which relate to the numéraire, derivatives, which often reflect speculative motives, and real estate, which combines investment with entrepreneurial activity. The sample also excludes affiliate and pump-and-dump content because both create misaligned incentives. Further exclusions cover sentiment-based or superficial content, news without a clear investment judgment, neutral mentions, general market commentary, and references without investment implications.

A short video qualifies as an explicit stock recommendation if it includes a direct buy/sell recommendation, a positive evaluation with actionable implications, or a visual ticker or firm display accompanied by an investment endorsement (e.g., “strong buy” or “good pick”). Three researchers independently classified the videos, cross-checked their results, and merged the classifications to finalize the sample.

Our statistical analysis requires stocks with different combinations of analyst and finfluencer recommendations. We therefore add stocks recommended made by analysts but not by finfluencers. We do not select these stocks arbitrarily, because systematic differences between the groups could bias the comparison. Instead, we restrict the comparison set to stocks that finfluencers could plausibly have recommended. A market-wide control group would include many small, illiquid, or less visible firms with a negligible ex-ante recommendation probability and would therefore confound finfluencer effects with broader differences in visibility and attention. We therefore select stocks that are similar in firm size, analyst coverage, country composition, and other relevant characteristics. Table 1 shows the descriptive statistics of the dataset from analysts and financial influencers. Perfect balance remains difficult because the largest and most visible firms are often well-known brands and thus particularly likely to attract finfluencer attention. Nevertheless, the selection procedure substantially reduces observable differences between the groups.

[8] A common approach is to label bullish/bearish sentiment signals as long/short positions (Ballinari & Behrendt, 2021; Bartov et al., 2018; Chen et al., 2014; Cookson et al., 2024; Giannini et al., 2018; Kakhbod et al., 2023). In contrast, we examine explicit investment recommendations—an approach that is conceptually closer to Jegadeesh et al. (2004) or Jegadeesh & Kim (2010).

3.2 Dependent Variables

For each stock s, recommender $i \in \{\text{Analyst}, \text{Finfluencer}\}$, and period t, we construct the recommendation count:

$$RecCount_{s,t,i} = \sum_{\tau \in t} \sum_{r \in \mathcal{R}_i} \text{Recommendation}_{s,\tau,r},$$

where $\text{Recommendation}_{s,\tau,r} = 1$ if a recommender issues a recommendation for stock s on date τ, and zero otherwise. We later include the log of exposure days as an offset in regression models to scale exposure time and yield incidence rate ratios.

3.3 Explanatory Variables and Controls

We operationalize exposure to the consumption-hedge channel using the variable *ConsumerLink*. We construct this measure based on Interbrand's *Best Global Brands* ranking.[9] To obtain established consumer brands with sufficient economic relevance for planned consumption, we impose a minimum brand-value threshold of USD 50 million. *ConsumerLink* equals one when a firm satisfies this criterion and zero otherwise. The variable therefore identifies stocks that may hedge changes in the future cost of consumption goods or services produced by the respective firm. Because brand rankings may not fully capture consumption-hedge exposure, particularly for diversified firms with broad product portfolios, we construct *ConsumerLink* (LLM) as a complementary measure. This measure uses an LLM-based classification of firm descriptions provided by LSEG. The LSEG descriptions summarize each firm's activities, products, and target markets, allowing the classification to identify consumer goods and services that may not be associated with a distinct or globally recognized brand. The LLM assigns a value of one when a firm supplies goods or services to end consumers and zero otherwise. *ConsumerLink* and *ConsumerLink* (LLM) exhibit a positive but modest correlation ($\rho = 0.110$, $p < 0.001$), consistent with the idea that the two measures capture distinct dimensions of consumption-hedge exposure (see Appendix D).

[9] Interbrand, Best Global Brands. Interbrand assesses brand value as the present value of future expected brand-driven economic gains by (i) determining the economic profit of the business, (ii) isolating the brand's contribution to this profit using the Role of Brand Index, and (iii) reflecting the sustainability of these returns through a brand-strength-based risk discount. The methodology is certified in accordance with ISO 10668. See Interbrand, Best Global Brands Methodology.

$ConsumerPull$ operationalizes the time variation in the strength of the consumption-hedge channel. The variable equals one during the four weeks preceding Christmas and one week before Black Friday and zero otherwise. Alternatively, we define $ConsumerPull\ (CSI)$ using a weekly median split of the consumer sentiment index in the recommender's country[10], with values above the median indicating a high-consumption state.

$Finfluencer$ is an indicator equal to one for finfluencer recommendations and zero for analyst recommendations. We retrieve analyst recommendations from LSEG and define them as explicit, firm-specific buy/sell recommendations.

$Controls$ includes variables operationalizing the financial channel, additional preferences, and attention. Since analyst recommendations are generally associated with price changes and trading volume, we include $Dividend\ Yield$, $Price\text{-}to\text{-}Book\ Ratio$, $Beta$, and $Momentum$, while $ESG\ Scores$ capture nonfinancial preferences (Buxbaum et al., 2023; Grinblatt et al., 2023; Lalwani, 2025). Analysts may also focus on attention-grabbing "glamour" stocks and the short-term confirmation of investor beliefs (Bilinski et al., 2019; Hribar & Mcinnis, 2012; Jegadeesh et al., 2004; Lai, 2004; Mullainathan & Shleifer, 2005). We therefore control for attention using stock-level $Google\ Trends$, $Trading\ Volume$, $Volatility$, and $Market\ Capitalization$, as well as the VIX for market-wide attention.

For several robustness checks, we further collect stock returns, risk-adjusted alphas following (Kakhbod et al., 2023), option-implied volatility, and financial-event measures from LSEG, together with TikTok engagement metrics, including $Follower$, $Likes$, and $Comments$ counts. Because the distance of analyst can convey informational advantages (Bae et al., 2008), we construct the variable $Home\ Bias$, which equals one when the recommender and the recommended firm are located in the same country, while $log(Distance)$ measures the geographic distance between their respective capital cities. From Datastream, we obtain changes in revenue, $\Delta Revenue$, and changes in operating margin, $\Delta\ Operating\ Margin$.

[10] IPSOS Primary Consumer Sentiment Index via LSEG.

4. Identification Strategy and Empirical Model

The empirical challenge is that exposure to the consumption-hedge channel is not randomly assigned across stocks. Firms that produce familiar consumer goods may also be larger, more liquid, and more widely followed (Barber & Odean, 2008; Huberman, 2001; Jegadeesh et al., 2004). A positive correlation between $Recommendations$ and $ConsumerLink$ therefore does not, by itself, identify the consumption-hedge channel, because it may reflect correlated firm characteristics.

As an identification strategy, we first control for observable differences in firm size, attention, trading activity, and fundamentals (Barber & Odean, 2008; Lin & Liu, 2018; Meshcheryakov & Winters, 2022). Of course, these controls cannot eliminate bias from unobserved firm characteristics. We therefore also exploit time variation in $ConsumerPull$. Under the assumption that unobserved differences between consumption-hedge and other stocks are stable over time, comparing recommendation rates across high- and low-consumption periods removes such time-invariant heterogeneity. Time-varying shocks may nevertheless affect recommendation activity more broadly. For example, news, market conditions, or shifts in investor attention may simultaneously increase recommendations of particular stocks (Cookson et al., 2024; Hribar & Mcinnis, 2012). To account for these shocks, we introduce analysts as a comparison group. Analysts are exposed to many of the same firm-level and market-wide developments (Bae et al., 2008; Kothari et al., 2016), but might respond to the consumption-hedge channel in the same way as finfluencers.

The resulting triple-difference design compares recommendation rates across three dimensions: stocks with and without consumption-hedge exposure, high- and low-consumption periods, and finfluencers and analysts. The identifying coefficient is the interaction $ConsumerLink_s \times ConsumerPull_t \times Finfluencer_r$. We use $RecRate$ rather than $RecCount$ to account for differences in the number of days identified by the alternative $ConsumerPull$ measures and thus ensure consistent interpretations across these measures.

The triple difference measures whether finfluencers have a higher likelihood to recommend consumption-hedge stocks during high-consumption periods than analysts. Stock-type differences absorb persistent characteristics of consumer-facing firms, time variation removes stable unobserved heterogeneity, and the analyst comparison accounts for shocks that affect recommendation activity across recommender groups. Identification therefore does not rely on consumption-hedge stocks being recommended more often on average. It relies on their recommendation rate increasing disproportionately among finfluencers, relative to analysts, precisely when the consumption-hedge channel is predicted to be stronger.

We estimate Negative Binomial regressions to accommodate overdispersion in recommendation counts. $RecCount$ is retained as the dependent variable, while the log of exposure days enters as an offset, such that the model estimates recommendation incidence rates per exposure day.

We regress

$$\log\left(\frac{\mathbb{E}[\text{RecCount}_{srt} \mid X_{srt}]}{days_t}\right) = \beta_{\text{DDD}}(\text{ConsumerLink}_s \times \text{ConsumerPull}_t \times \text{Finfluencer}_r) + \text{lower-order terms} + \gamma' X_{srt}. \qquad (8)$$

with stock s a recommender r recommends in consumption period t. All specifications use robust standard errors clustered at the stock level.

The identifying assumption is that, absent the consumption-hedge channel, the change in the $ConsumerLink$-versus-non-$ConsumerLink$ recommendation gap between high- and low-$ConsumerPull$ periods would have evolved similarly for finfluencer and analysts. Equivalently, any omitted factor capable of generating the triple interaction spuriously would have to be simultaneously $ConsumerLink$-specific, $ConsumerPull$-specific, and $Finfluencer$-specific. After careful consideration, we cannot think of any plausible factor that fulfills this criterion.

Nevertheless, to test the identification assumption, we first check the internal validity of the estimate. To do this, we investigate whether the triple interaction is robust to alternative operationalizations of $ConsumerLink$ and $ConsumerPull$, whether no corresponding effects occur in placebo tests, and whether there are no spillovers between the recommendations of finfluencers and analysts. We then examine alternative behavioral and financial explanations for the overall pattern. Post-recommendation returns, risk-adjusted performance, and option-implied volatility provide insight into whether finfluencers' incentives are aligned with their audience's interests and whether the results are consistent with FOMO. Finally, we examine heterogeneous effects to determine whether the consumption-hedge channel exists across different finfluencer types or is focused on specific groups and audience structures.

5. Results

5.1 Descriptive Statistics

Table 1 summarizes our dataset. The sampled TikTok channels have a combined 4.73 million followers, and their stock recommendation short videos have drawn 37.5 million views, 1.33 million likes, and 34,700 comments. This scale shows that finfluencer stock recommendations extend well beyond a narrow audience.

Insert Table 1

After applying the strict eligibility filters, the sample contains 544 finfluencer recommendations and 21,428 analyst recommendations. Analysts exhibit a substantially higher daily recommendation rate: 1.4%, compared with 0.1% for finfluencers. Overall, 93,4% of finfluencers' recommendations are buy recommendations, which mirrors the well-documented positivity bias among analysts (De Bondt & Forbes, 1999; Lim, 2001).

Moreover, the triple difference design is based on 3,672 observations, reflecting the 2×2 combinations of the binary $ConsumerLink$ and $ConsumerPull$ indicators with 918 stocks. $ConsumerLink$ measures firms exposed to the consumption-hedge mechanism, with the brand-based measure capturing a focused set of 10 stocks, and the broader LLM-based measure $ConsumerLink\ (LLM)$ identifying 154 stocks. $ConsumerPull$ distinguishes observations during the pre-Christmas and Black Friday period from those outside this period. The alternative $ConsumerPull\ (CSI)$ classifies each stock-period according to whether the change in consumption sentiment among the recommenders viewing the stock is above or below its median, with these CSI changes exhibiting a standard deviation of 1.16 percentage points.

5.2 Calibrating the Theoretical Model

We now examine the empirical implications of the model in Equation (1). The key structural object is $\rho_{\Pi,R}$, which captures how firm outcomes covary with stock returns. We proxy it using the correlation of firm returns with changes in revenue and operating margin. Appendix Figure A.1 shows that firms with higher values of these proxies also tend to receive more finfluencer mentions. We next calibrate the model in Equation (6) to assess whether it can reproduce observed recommendation behavior with economically plausible parameters. The calibration uses quarterly data and rolling-window estimates that reflect the frequency of accounting disclosures. Because the outcome is the probability that a stock is recommended, the estimates are interpreted as selection probabilities rather than portfolio weights. Technical details are provided in Appendix E.

Table 2 reports stable parameter estimates across specifications. The estimated loss-aversion parameter is $\kappa = 2.388$, close to values in the prospect-theory literature (Tversky & Kahneman, 1991). The reference parameter, $\chi = 0.980$, implies a reference point about 2% below the previous-period level, while γ lies in the range of 14.945 to 15.345 and α between 0.392 and 0.465. Overall, the calibration shows that the model can rationalize observed finfluencer recommendations and is consistent with $ConsumerLink$ forming part of their recommendation rationale. Because this result depends on functional-form, calibration, and proxy assumptions, we next test the model's main directional prediction in a reduced-form triple difference design. To examine the mechanism in Equation (7), we use the weekly recommendation rate, $RecRate$, which better matches the video feed horizon of finfluencer content.

Insert Table 2

5.3 Identifying the Consumption-Hedge Channel

Table 3 reports the baseline triple-difference estimates. The coefficient of primary interest, $ConsumerLink \times ConsumerPull \times Finfluencer$, captures whether the relative recommendation rate for consumption-linked stocks increases more strongly for finfluencers than for analysts when the consumption-hedge channel is active. Using the baseline $ConsumerLink$ measure, the triple interaction ranges from 0.667 to 0.702 across Columns (2)–(4) and is statistically significant at conventional levels throughout. Since the model is estimated with a log link and exposure days enter as an offset, exponentiating these coefficients yields incidence-rate ratios. The estimates therefore imply that activation of the consumption channel raises the corresponding relative recommendation-rate contrast by approximately 94.8% to 101.8%. Importantly, the magnitude of the triple interaction remains stable when firm-level controls and their interactions with $ConsumerPull$ are introduced, indicating that the result is not driven by observable differences in firm characteristics or heterogeneous responses to these characteristics across consumption states.

The right panel provides a complementary test using $ConsumerLink$ (LLM). The estimated triple interaction is smaller, ranging from 0.430 to 0.506. After the inclusion of controls, the coefficients in Columns (7) and (8) are statistically significant at the 10% level and imply increases in the relative recommendation-rate contrast of approximately 65.9% and 65.4%, respectively. The alternative text-based classification produces economically meaningful but less precisely estimated effects than the baseline $ConsumerLink$ measure.

Insert Table 3

The remaining coefficients help distinguish the triple-difference effect from broader differences in recommendation behavior. Finfluencer is strongly negative across all specifications, reflecting a substantially lower baseline recommendation rate for finfluencers than for analysts. At the same time, the positive $ConsumerPull \times Finfluencer$ interaction indicates that finfluencer recommendation activity rises relative to analyst activity during high- $ConsumerPull$ periods more generally. The positive triple interaction goes beyond this aggregate change: it identifies an additional increase specifically for consumption-linked stocks. The $ConsumerLink \times Finfuencer$ interaction is also positive and significant under the baseline $ConsumerLink$ classification, suggesting a stronger unconditional orientation of finfluencers toward such stocks, whereas this pattern is not present under the $ConsumerLink\,(LLM)$. Moreover, the standalone $ConsumerLink$ coefficient becomes markedly smaller once firm characteristics are controlled for. Taken together, these results identify the consumption-hedge channel as the disproportionate increase in finfluencer recommendations of consumption-linked stocks during periods of high $ConsumerPull$, net of broader differences in recommendation activity.

5.4 Robustness to Identification Treats

The triple difference estimate isolates the consumption-hedge channel only under the identifying assumption stated above. A remaining concern is that the calendar-based

ConsumerPull measure may coincide with seasonal changes in firms' information environments or other forms of time-varying investor attention that are unrelated to the proposed mechanism.

Time variation of the consumption-hedge channel

The preceding specifications already mitigate this concern through firm-level controls, interacted controls, and the triple-difference structure. To further separate the consumption-hedge mechanism from calendar or seasonal effects, we replace the indicator-based $ConsumerPull$ measure with $ConsumerPull\ (CSI)$, which captures changes in country-level consumer sentiment around the recommendation date.

Table 4 reports the corresponding estimates. The triple interaction remains positive and highly statistically significant across all specifications in which it is identified. For the baseline $ConsumerLink$ measure, the coefficients range from 37.187 to 44.222; using $ConsumerLink\ (LLM)$, they range from 33.498 to 37.140. Because $ConsumerPull\ (CSI)$ is measured as a continuous log change, the magnitudes of these coefficients are not directly comparable with those in Table 3. Evaluated at a one-standard-deviation increase in $ConsumerPull\ (CSI)$ of 0.011, however, the estimates imply an increase in the relevant relative recommendation-rate contrast of approximately 44.6% to 62.7%. The implied effects range from 50.5% to 62.7% under the baseline $ConsumerLink$ measure and from 44.6% to 50.5% under $ConsumerLink\ (LLM)$.

Insert Table 4

The sentiment-based specification therefore yields the same qualitative conclusion using variation that is independent of the calendar definition underlying the baseline $ConsumerPull$ measure. Although the standardized effect sizes are somewhat smaller than those obtained with the indicator-based measure, their magnitudes should not be compared mechanically because the two variables capture different forms and scales of variation in consumption pulling. More importantly, the triple interaction remains positive and statistically strong under both $ConsumerLink$ definitions. The combined evidence from Tables 3 and 4 therefore supports a robust consumption-hedge channel across alternative measurements of both the consumption linkage of the recommended stock and the time-varying strength of consumption demand.

We randomly assign the treatment indicator and re-estimate the triple interaction; Table 5 presents a baseline scenario in which all triple difference effects vanish once $ConsumerPull$ is randomly assigned. Across 1,000 iterations, the placebo estimates are lower in 88% and 67% of cases, respectively. These reductions are substantial given that $ConsumerPull$ is not defined by a discrete external intervention with a precisely observable treatment date, such as a regulatory change, but instead captures theoretically derived periods of heightened consumer attention. The placebo results therefore provide strong support for the chosen consumption time frames, while a more narrowly defined $ConsumerPull$ measure would likely strengthen the estimated effects without changing our qualitative conclusion.

Insert Table 5

Finfluencer–Analyst variation in consumption-hedge channel

Event Effects

Besides controls, interacted controls, and fixed effects, we use the comparison between finfluencer and analysts to cancel out common unobservables that affect the same stock at the same time, such as attention, perception, or fundamentals. Still, the comparison suffers if institutional constraints make analysts less responsive to fundamental information events, blurring the distinction between consumption and the financial channel. We test this conjecture in Table 6 with a negative binomial regression that contrasts $Recommendation$ activity per stock during periods of high versus low $ConsumerPull$ and around earnings announcements versus periods without earnings announcements.

Insert Table 6

In the pooled specification in Columns (1)-(2), recommendation activity declines in $ConsumerPull$ as measured by the Christmas and Black Friday period (-0.143, $p < 0.01$) while the financial state shows increases (0.885, $p < 0.01$). This pattern is consistent with analyst behavior dominating the aggregate.

Once we distinguish finfluencers from analysts, the interaction coefficient on $Finfluencer \times ConsumerPull$ is positive and statistically significant (1.074, $p < 0.01$), whereas the coefficient on $Finfluencer \times FinanceState$ is negative (-0.453, $p < 0.01$). Consequently, finfluencers differ from analysts as they respond to a 193% increase in the expected number of recommendations in periods of consumption framing; however, in periods of financial information, they actually reduce those by 36%.

By contrast, analysts respond strongly to earnings events: the coefficient of 0.79 ($p < 0.01$) implies a 120% increase in the expected number of recommendations, holding other variables constant. Therefore, their muted response during consumption periods reflects a difference in recommendation behavior rather than a binding institutional constraint, which is the premise underlying the third difference in our identification strategy.

Peer Effects

Our identification strategy requires no spillover between analysts and finfluencers. We therefore test whether analyst and finfluencer recommendations affect each other. Table 7 estimates a fixed-effects negative binomial model that links current recommendation activity to lagged cumulative recommendations from each group, interacted with recommender type.

We predict the number of recommendations in each week using recommendation activity over the preceding seven days. Columns (1) and (3) initially suggest that prior analyst activity predicts subsequent analyst recommendations, prior finfluencer activity predicts subsequent finfluencer recommendations, and each group also appears to respond to the other. This cross-group association disappears once the specifications include controls. In Columns (2) and (4), analyst recommendations respond only to prior analyst activity, while finfluencer recommendations respond only to prior finfluencer activity. One additional recommendation in the preceding week increases the expected number of recommendations in the following week by 45% for analysts ($\beta = 0.371$, $p < 0.01$) and by 311% for finfluencers ($\beta = 1.413$, $p < 0.01$).

The uncontrolled cross-group relationship, therefore, reflects common determinants, including fundamentals and attention. Once these factors are accounted for, the two groups follow distinct recommendation processes. Finfluencer recommendations are thus not merely copies of analyst recommendations. At a broader level, this segmented pattern is consistent with the peer-effects and stock market participation literature in finance (Ali-

Rind et al., 2023). Consistent with prior evidence that found localized peer effects (Betzer & Harries, 2022; Bursztyn et al., 2014; Seasholes & Wu, 2007), Table 7 shows limited cross-group transmission, suggesting that spillovers are unlikely to threaten our identification.

Insert Table 7

Cultural variation of the consumption-hedge channel

The differences between finfluencers and analysts may extend beyond time periods and stock characteristics. If finfluencers cater to a consumption-hedge channel rather than the financial channel alone, cultural and geographic consumption patterns should also shape their recommendations. We test this prediction using a recommendation-level regression.

Table 8 reports fixed-effects negative binomial regressions of recommendation activity on geographic and cultural proximity. Recommendation rates decline with geographic distance for both groups, but significantly less so for finfluencers. The estimated distance elasticity is approximately -0.70 for finfluencers, compared with -1.18 for analysts; the interaction coefficient of 0.48 ($p < 0.10$) implies that a 1% increase in distance in kilometers reduces finfluencer recommendations by about 0.48 percentage points less than analyst recommendations. Finfluencers also display a substantially stronger preference for firms located in their home country. The interaction coefficient of 4.45 ($p < 0.05$) implies that the home-country recommendation-rate ratio is estimated to be $e^{4.45} \approx 85.6$ times larger for finfluencers than for analysts. Likewise, the coefficient of 1.48 for English-speaking markets ($p < 0.10$) implies an incidence-rate ratio of $e^{1.48} \approx 4.39$. These estimates indicate that, relative to analyst recommendations (Hao et al., 2025; Huberman, 2001), finfluencer recommendations are less sensitive to physical distance but more strongly associated with home-country affiliation and cultural familiarity.

Insert Table 8

5.5 Robustness of the Financial Channel

We already applied a stringent data-filtering pipeline to ensure that the finfluencers in our sample are sufficiently incentive-aligned with their audiences and that their recommendations, therefore, follow the decision problem formalized in Equation (2). We now further assess this assumption empirically. The model requires finfluencers to make economically meaningful, intertemporal, FOMO-related stock recommendations, rather than act as generic influencers or engage in fraud or scamming. We therefore examine recommendation performance to evaluate whether the data support the assumed incentive alignment.

Returns from Recommendations

Table 9 compares finfluencer and analyst recommendations for stocks covered by both groups. Analysts exhibit strong short-horizon performance, with annualized returns of 46.3% before risk adjustment and 38.9% after controlling for common risk factors; both estimates are highly statistically significant. Finfluencer recommendations show no statistically significant short-run outperformance, either in absolute terms or relative to analysts. Their performance improves, however, as the holding period lengthens. Over a one-quarter horizon, finfluencer recommendations generate an annualized Fama–French six-factor alpha of 6.68%. Depending on the specification, their performance exceeds that of analyst recommendations by 6.42% to 8.91%. These results support that finfluencer recommendations are oriented toward longer holding periods. They also provide no evidence that the finfluencers in our sample systematically deceive their target groups.

Insert Table 9

Table 10 extends the analysis to the full stock sample, including firms covered only by analysts. At the short-term day trading interval, analyst performance ranges from an annualized CAPM alpha of 38.02% to a raw return of 47.45%, whereas finfluencer recommendations again produce no statistically significant returns. The pattern reverses over longer horizons. At the quarterly horizon, finfluencer recommendations generate annualized alphas ranging from 6.68% under the Fama–French six-factor model to 10.16% in the least restrictive specification. These estimates are larger than those for analyst recommendations, although the differences are not consistently statistically significant. The broader sample evidence is that analyst derive part of their advantage from stocks that finfluencers do not cover. Overall, the strong returns do not point to “pump-and-dump” schemes or fraudulent practices; rather, they reflect that our sample of finfluencer stock recommendations is underpinned by aligned interests and a catering relationship between finfluencers and their audience.

Insert Table 10

Implied Risk from Recommendations

The alpha estimates provide a detailed view of recommendation performance, but they may remain sensitive to the sample holding period and general limits of factor models. We therefore complement this analysis by examining the option-implied volatility. Table 11 provides direct evidence on the implied volatility finfluencer- and analyst-recommended stocks in our entire sample have across the option moneyness spectrum and two maturities.

Insert Table 11

Stocks that finfluencers recommend exhibit substantially elevated implied volatility across the entire surface. At the one-month horizon, implied volatility at the 20-delta put is 34.83% for finfluencer stocks versus 24.95% for analyst stocks, a difference of 9.88% volatility points ($p < 0.01$). Comparable differentials at-the-money ($\Delta = 9.96$) and for the 20-delta call ($\Delta = 10.05$) show that the gap extends across the moneyness range. The quarterly horizon reproduces this pattern.

Decomposing the surface into tail-specific components reveals an asymmetric structure. Left-tail risk is similar across groups, at 2.47% and 2.64%, respectively ($\Delta = -0.17\%$, not statistically significant). In contrast, finfluencer stocks exhibit moderately greater right-tail risk, with 20-delta call implied volatility exceeding ATM implied volatility by an additional 0.47 (0.54) volatility percentage points ($p < 0.10$) at the one-month (one-quarter) horizon.

Overall, finfluencers recommend more volatile stocks with greater upside potential than analysts, consistent with more speculative, lottery-like stock selection. This difference disappears among $ConsumerLink$ stocks, which are themselves more volatile and offer greater upside potential. Within this group, higher volatility reduces the likelihood of a finfluencer recommendation relative to an analyst recommendation, whereas finfluencers remain more responsive to volatility among non-$ConsumerLink$ stocks. This pattern aligns with the model, in which higher σ_S reduces exposure when consumption-hedging benefits are present. At the same time, the stronger role of upside potential among $ConsumerLink$ stocks supports the model's FOMO channel. Investors are especially reluctant to miss gains in stocks linked to future consumption.

5.6 Heterogeneity

We have described the average consumption-hedge channel effect and will now examine its heterogeneity among different finfluencer. A generalized linear mixed model (GLMM) with a negative binomial distribution relates engagement per follower ($Views$, $Likes$, $Comments$) to $ConsumerLink$, with finfluencer-level random intercepts absorbing unobserved creator heterogeneity.

Table 12 reports coefficients of $ConsumerLink$ as measured by both brand value (Columns 1–3) and the LLM-based classification (Columns 4–6). Both are economically small and statistically insignificant, so consumption signals alone do not uniformly enhance engagement. Serving the consumption-hedge channel appears only reasonable once we account for the finfluencer's size. We differentiate between large-audience finfluencer if follower size exceeds its median of 19,100, and small-audience finfluencer otherwise.[11]

Large-audience finfluencer benefits disproportionately from featuring consumer brand stocks, despite exhibiting substantially lower baseline engagement per follower (-1.649, $p < 0.01$ for views). The interaction coefficients are positive and significant across all three engagement dimensions: 0.476 ($p < 0.05$) for views, 0.441 ($p < 0.10$) for likes, and 0.488 ($p < 0.05$) for comments per follower. Economically, associating investment content with prominent consumer brands raises engagement by roughly 60% relative to large-audience finfluencers' own baseline ($e^{0.476} \approx 1.61$), consistent brand recognition reinforcing algorithmic visibility at scale.

Small-audience finfluencers in Columns (4)–(6) operate from a higher engagement baseline (1.503, $p < 0.01$ for views) and respond more strongly to the broader LLM-based measure of consumer-facing firms than to brand recognition alone. The positive interaction for comments per follower (0.131, $p < 0.05$), together with directionally positive effects for views and likes, suggests that their narrower audiences engage more with firms whose products are relevant to end consumers, even when those firms lack highly recognizable

[11] The influencer marketing literature commonly distinguishes between *micro* and *macro influencers* based on audience size (Leung, Gu and Palmatier, 2022; Cheng and Zhang, 2025). We refrain from using the corresponding terms *micro-* and *macro-finfluencer* to avoid confusion with a potential distinction based on micro- versus macroeconomic content. Instead, *small-audience* and *large-audience finfluencer* refer exclusively to differences in follower size.

brands. By contrast, prominent consumer brands appear more important for large-audience finfluencers.

Insert Table 12

Overall, our findings suggest that the consumption-hedge channel is heterogeneous with respect to finfluencer size. Large-audience finfluencers focus on consumer brands, while small-audience finfluencers pick up the less obvious consumption hedges—two complementary pathways through which catering to the consumption-hedge channel translates into finfluencer success. This size-dependent pattern aligns with broader evidence that large-audience influencers primarily optimize for reach, whereas small-audience influencers generate higher engagement returns through closer audience relationships (Beichert et al., 2024; Leung, Gu, & Palmatier, 2022).

6. Conclusion

We provide first evidence that finfluencers provide a distinct form of financial advice rather than extending the role of traditional financial analysts. Unlike analysts, whose audiences typically seek financial information and evaluate recommendations based on fundamentals, finfluencers communicate with broad audiences embedded in consumer and lifestyle-oriented media environments. We find that finfluencing reaffirms the fundamental economic role of investment in supporting intertemporal consumption. Finfluencers portray stocks primarily as instruments that secure and enable conceivable consumption, following the motto “invest in what you will buy.”

In our theoretical framework, an agent demands stocks because ownership provides a hedge against future consumption costs when forward contracts on the desired product are missing. This consumption-hedging motive creates an additional channel that differs from fundamental, informational, and attention-driven investment demand. Our empirical evidence supports this mechanism. Using TikTok stock recommendations, we find that consumer framing increases recommendations of consumer-goods firms, generates higher engagement, and is associated with strong long-horizon performance. We also document important heterogeneity across finfluencers. Large-audience finfluencers tend to emphasize widely recognized brands and reach more followers, whereas small-audience finfluencers tend to focus on consumer-oriented, lesser-known stocks, but generate stronger engagement. Moreover, finfluencer recommendations are more closely associated with cultural proximity than with traditional financial information events, highlighting a communication mechanism that differs from that of analysts.

Our findings suggest that evaluating finfluencers solely by the skill to generate alpha provides an incomplete picture, with important implications for financial regulation. Finfluencer recommendations also address intertemporal consumption risks. Relying solely on conventional performance criteria that treat financial returns as an end in themselves therefore overlooks their broader economic role. Finance research and financial regulation should consequently adopt a broader framework for evaluating finfluencer activity.

Our findings also contribute to understanding why short-form video platforms can broaden participation in financial markets. By favoring narratives with persistent relevance over information that becomes quickly outdated, these platforms allow finfluencers to frame investing through durable themes such as consumption aspirations. Such content can provide accessible entry points into finance, particularly for younger audiences and individuals with limited financial literacy. Although finfluencing also encompasses pump-and-dump schemes, opportunistic sponsorships, and fraudulent advice (Rockwell & Conklin, 2026; Zülch et al., 2026), these practices do not define the phenomenon as a whole. Platforms such as TikTok have the potential to make the benefits of capital markets accessible to broader segments of society.

Several questions remain for future research. Although we focus on stock recommendations, the consumption-hedge channel may also shape longer-term investment behavior and financial decisions beyond stock investing. Identifying these additional outcomes could help researchers assess the societal welfare implications of finfluencing and, more generally, of digital financial communication.

References

Abel, A. B. (1990). Asset Prices Under Habit Formation And Catching Up With The Joneses. *The American Economic Review*, *80*(2), 38–42. https://www.jstor.org/stable/2006539

Ali-Rind, A., Boubaker, S., & Lajili Jarjir, S. (2023). *Peer effects in financial economics: A literature survey*. https://doi.org/10.1016/j.ribaf.2022.101873

Bae, K. H., Stulz, R. M., & Tan, H. (2008). Do local analysts know more? A cross-country study of the performance of local analysts and foreign analysts. *Journal of Financial Economics*, *88*(3), 581–606. https://doi.org/10.1016/J.JFINECO.2007.02.004

Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies*, *16*(1), 101–143. https://doi.org/10.1093/rfs/16.1.0101

Ballinari, D., & Behrendt, S. (2021). How to gauge investor behavior? A comparison of online investor sentiment measures. *Digital Finance*, *3*(2), 169–204. https://doi.org/10.1007/s42521-021-00038-2

Barber, B. M., & Odean, T. (2008). All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies*, *21*(2), 785–818. https://doi.org/10.1093/rfs/hhm079

Bartov, E., Faurel, L., & Mohanram, P. S. (2018). Can Twitter Help Predict Firm-Level Earnings and Stock Returns? *The Accounting Review*, *93*(3), 25–57. https://doi.org/10.2308/accr-51865

Beichert, M., Bayerl, A., Goldenberg, J., & Lanz, A. (2024). Revenue Generation Through Influencer Marketing. *Journal of Marketing*, *88*(4), 40–63. https://doi.org/10.1177/00222429231217471

Betzer, A., & Harries, J. P. (2022). How online discussion board activity affects stock trading: the case of GameStop. *Financial Markets and Portfolio Management*, *36*, 443–472. https://doi.org/10.1007/s11408-022-00407-w

Bilinski, P., Cumming, D., Hass, L., Stathopoulos, K., & Walker, M. (2019). Strategic distortions in analyst forecasts in the presence of short-term institutional investors. *Accounting and Business Research*, *49*(3), 305–341. https://doi.org/10.1080/00014788.2018.1510303

Bursztyn, L., Ederer, F., Ferman, B., & Yuchtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, *82*(4), 1273–1301. https://doi.org/10.3982/ECTA11991

Buxbaum, M., Schultze, W., & Tiras, S. L. (2023). Do analysts' target prices stabilize the stock market? *Review of Quantitative Finance and Accounting*, *61*(3), 763–816. https://doi.org/10.1007/s11156-023-01164-1

Chen, H., De, P., Hu, Y., Hwang, B.-H., Faccio, M., Kim, S., Liberti, J., Lou, D., Loughran, T., Polak, I., Xu, J., & Jefferies, seminar participants at. (2014). Wisdom of Crowds: The

Value of Stock Opinions Transmitted Through Social Media. *The Review of Financial Studies*, *27*(5), 1367–1403. https://doi.org/10.1093/RFS/HHU001

Constantinides, G. M. (1990). Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy*, *98*(3), 519–543. https://doi.org/10.1086/261693

Cookson, J. A., Lu, R., Mullins, W., & Niessner, M. (2024). The social signal. *Journal of Financial Economics*, *158*, 103870. https://doi.org/10.1016/j.jfineco.2024.103870

De Bondt, W. F. M., & Forbes, W. P. (1999). Herding in analyst earnings forecasts: evidence from the United Kingdom. *European Financial Management*, *5*(2), 143–163. https://doi.org/10.1111/1468-036X.00087

Dennis, P., & Mayhew, S. (2002). Risk-Neutral Skewness: Evidence from Stock Options. *The Journal of Financial and Quantitative Analysis*, *37*(3), 471. https://doi.org/10.2307/3594989

Epstein, L. G., & Zin, S. E. (1991). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. *Journal of Political Economy*, *99*(2), 263–286. https://doi.org/10.1086/261750

Epstein, L. G., & Zin, S. E. (2013). Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. In *World scientific handbook in financial economics series* (Vol. 57, pp. 207–239). https://doi.org/10.1142/9789814417358_0012

FINRA & CFA Institute, C. (2023). *Gen Z and Investing: Social Media, Crypto, FOMO, and Family*. https://rpc.cfainstitute.org/research/reports/2023/gen-z-investing

Giannini, R., Irvine, P., & Shu, T. (2018). Nonlocal disadvantage: An examination of social media sentiment. *Review of Asset Pricing Studies*, *8*(2), 293–336. https://doi.org/10.1093/rapstu/rax020

Gibson, D. J., Tung, A., Chen, K., Hoffman, M., & Castaneda, L. (2025). From Budgeting to Day Trading: A Content Analysis of Financial Advice on {TikTok}. *Journal of Financial Counseling and Planning*, *36*(2), 315–324. https://doi.org/10.1891/JFCP-2023-0132

Grinblatt, M., Jostova, G., & Philipov, A. (2023). Analyst Bias and Mispricing. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4413833

Guan, S. S. (2023). *The rise of the finfluencer. NYUJL & Bus., 19, 489.*

Guo, J., & Dong He, X. (2017). Equilibrium asset pricing with Epstein-Zin and loss-averse investors. *Journal of Economic Dynamics and Control*, *76*, 86–108. https://doi.org/10.1016/j.jedc.2016.12.008

Haase, F., Rath, O., Krauß, J., & Schoder, D. (2026). The Role of Finfluencers in Shaping Crowd Sentiment. *Business & Information Systems Engineering*, *68*(3), 559–582. https://doi.org/10.1007/s12599-025-00947-1

Hao, H., Liu, C., & Pang, S. (2025). Geographical proximity, cultural familiarity and financial information production. *Journal of Empirical Finance*, *80*, 101576. https://doi.org/10.1016/j.jempfin.2024.101576

Hogsnes, M., & Grønli, T.-M. M. (2025). A commercial playground: investigating TikTok influencers' commercial content. *Information, Communication & Society*, 1–21. https://doi.org/10.1080/1369118X.2025.2588348

Hribar, P., & Mcinnis, J. (2012). Investor Sentiment and Analysts' Earnings Forecast Errors. *Management Science*, *58*(2), 293–307. https://doi.org/10.1287/mnsc.1110.1356

Huberman, G. (2001). Familiarity breeds investment. *The Review of Financial Studies*, *14*(3), 659–680. https://doi.org/10.1093/rfs/14.3.659

Jegadeesh, N., Kim, J., Krische, S. D., & Lee, C. M. C. (2004). Analyzing the Analysts: When Do Recommendations Add Value? *The Journal of Finance*, *59*(3), 1083–1124. https://doi.org/10.1111/J.1540-6261.2004.00657.X

Jegadeesh, N., & Kim, W. (2010). Do Analysts Herd? An Analysis of Recommendations and Market Reactions. *The Review of Financial Studies*, *23*(2), 901–937. https://doi.org/10.1093/RFS/HHP093

Kakhbod, A., Kazempour, S. M., Livdan, D., & Schuerhoff, N. (2023). Finfluencers. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4428232

Kothari, S. P., So, E., & Verdi, R. (2016). Analysts' Forecasts and Asset Pricing: A Survey. *Annual Review of Financial Economics*, *8*(Volume 8, 2016), 197–219. https://doi.org/10.1146/ANNUREV-FINANCIAL-121415-032930/CITE/REFWORKS

Lai, R. K. (2004). A Catering Theory of Analyst Bias. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.548582

Lalwani, V. (2025). Finfluencer recommendations. *Economics Letters*, *255*, 112511. https://doi.org/10.1016/j.econlet.2025.112511

Leung, F. F., Gu, F. F., & Palmatier, R. W. (2022). Online influencer marketing. *Journal of the Academy of Marketing Science*, *50*(2), 226–251. https://doi.org/10.1007/s11747-021-00829-4

Lim, T. (2001). Rationality and Analysts' Forecast Bias. *The Journal of Finance*, *56*(1), 369–385. https://doi.org/10.1111/0022-1082.00329

Lin, T. C., & Liu, X. (2018). Skewness, Individual Investor Preference, and the Cross-section of Stock Returns. *Review of Finance*, *22*(5), 1841–1876. https://doi.org/10.1093/ROF/RFX036

Lindström, B., Bellander, M., Schultner, D. T., Chang, A., Tobler, P. N., & Amodio, D. M. (2021). A computational reward learning account of social media engagement. *Nature Communications*, *12*(1), 1311. https://doi.org/10.1038/s41467-020-19607-x

Liu, Z., Zou, L., Zou, X., Wang, C., Zhang, B., Tang, D., Zhu, B., Zhu, Y., Wu, P., Wang, K., & Cheng, Y. (2022). Monolith: Real Time Recommendation System With Collisionless

Embedding Table. *CEUR Workshop Proceedings*, *3303*. https://arxiv.org/abs/2209.07663v2

Mardon, R., Cocker, H., Daunt, K., & Kozinets, R. (2026). Spheres of influence: How influencers impact society. *Journal of Business Research*, *215*, 116339. https://doi.org/10.1016/j.jbusres.2026.116339

Meshcheryakov, A., & Winters, D. B. (2022). Retail investor attention and the limit order book: Intraday analysis of attention-based trading. *International Review of Financial Analysis*, *81*, 101627. https://doi.org/10.1016/j.irfa.2020.101627

Mixon, S. (2011). What Does Implied Volatility Skew Measure? *The Journal of Derivatives*, *18*(4), 9–25. https://doi.org/10.3905/JOD.2011.18.4.009

Mölders, M., Bock, L., Barrantes, E., & Zülch, H. (2025). Understanding finfluencers: Roles and strategic partnerships in retail investor engagement. *Journal of Business Research*, *198*, 115462. https://doi.org/10.1016/j.jbusres.2025.115462

Mullainathan, S., & Shleifer, A. (2005). The Market for News. *American Economic Review*, *95*(4), 1031–1053. https://doi.org/10.1257/0002828054825619

Pandey, P. K., Chaturvedi, P., Agnihotri, D., Bajpai, N., & Tripathi, V. (2026). From reels to returns: the influence of female financial influencers on investment decisions. *Journal of Financial Services Marketing*, *31*(1), 21. https://doi.org/10.1057/s41264-026-00346-x

Pedersen, L. H. (2022). Game on: Social networks and markets. *Journal of Financial Economics*, *146*(3), 1097–1119. https://doi.org/https://doi.org/10.1016/j.jfineco.2022.05.002

Przybylski, A. K., Murayama, K., DeHaan, C. R., & Gladwell, V. (2013). Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior*, *29*(4), 1841–1848. 10.1016/j.chb.2013.02.014

Qian, K., & Jain, S. (2024). Digital Content Creation: An Analysis of the Impact of Recommendation Systems. *Management Science*, *70*(12), 8668–8684. https://doi.org/10.1287/mnsc.2022.03655

Rockwell, C., & Conklin, M. (2026 forthcoming). Regulating Investment Advice in the Age of Social Media: Finfluencers, SEC Regulation, and SEC v. Constantinescu. *DePaul Law Review*, *73*.

Seasholes, M. S., & Wu, G. (2007). Predictable behavior, profits, and attention. *Journal of Empirical Finance*, *14*(5), 590–610. https://doi.org/10.1016/j.jempfin.2007.03.002

Securities and Exchange Commission, Investor Advisory Committee. (2024, December 10). *Recommendation of the SEC Investor Advisory Committee regarding the protection of investors in their interactions with finfluencers*. SEC https://www.sec.gov/files/approved-finfluencer-recommendations-20241210.pdf

Stocktwits. (2025, September 29). *Stocktwits and CMT association launch global platform for technical analyst professionals*. https://stocktwits.com/newsarticles/business/others/stocktwits-and-cmt-association-launch-global-platform/chDucQRR3JI

TikTok. (2025, September 5). *TikTok's community in europe tops 200 million*. https://newsroom.tiktok.com/tiktok-community-in-europe-tops-200-million?lang=en-150

TikTok. (2026a, January 23). *Announcement from the new TikTok USDS joint venture LLC*. https://newsroom.tiktok.com/announcement-from-the-new-tiktok-usds-joint-venture-llc?lang=en

TikTok. (2026b, May 19). *Built on TikTok: Celebrating the thriving small business community*. https://newsroom.tiktok.com/built-on-tiktok-celebrating-the-thriving-small-business-community?lang=en

Tversky, A., & Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. *The Quarterly Journal of Economics*, *106*(4), 1039–1061. https://doi.org/10.2307/2937956

Vombatkere, K., Mousavi, S., Zannettou, S., Roesner, F., & Gummadi, K. P. (2024). TikTok and the Art of Personalization: Investigating Exploration and Exploitation on Social Media Feeds. *Proceedings of the ACM Web Conference 2024*, 3789–3797. https://doi.org/10.1145/3589334.3645600

Yang, J., Zhang, J., & Zhang, Y. (2025). Engagement That Sells: Influencer Video Advertising on TikTok. *Marketing Science*, *44*(2), 247–267. https://doi.org/10.1287/mksc.2021.0107

Zhang, Y., Liu, C., & Lang, C. (2025). How luxury fashion brands leverage TikTok to captivate young consumers: an exploratory investigation using video analytics. *Journal of Marketing Analytics*, *13*(1), 128–144. https://doi.org/10.1057/s41270-023-00276-w

Zülch, H., Mölders, M., & Hoffmann, C. P. (2026). Navigating the rise of finfluencers: a multidimensional quality framework for strategic partnerships. *Journal of Financial Services Marketing*, *31*(1), 26. https://doi.org/10.1057/s41264-025-00334-7

Figures

Figure 1: ***Simulated Consumption-hedge driven Demand and Wealth Dynamics***

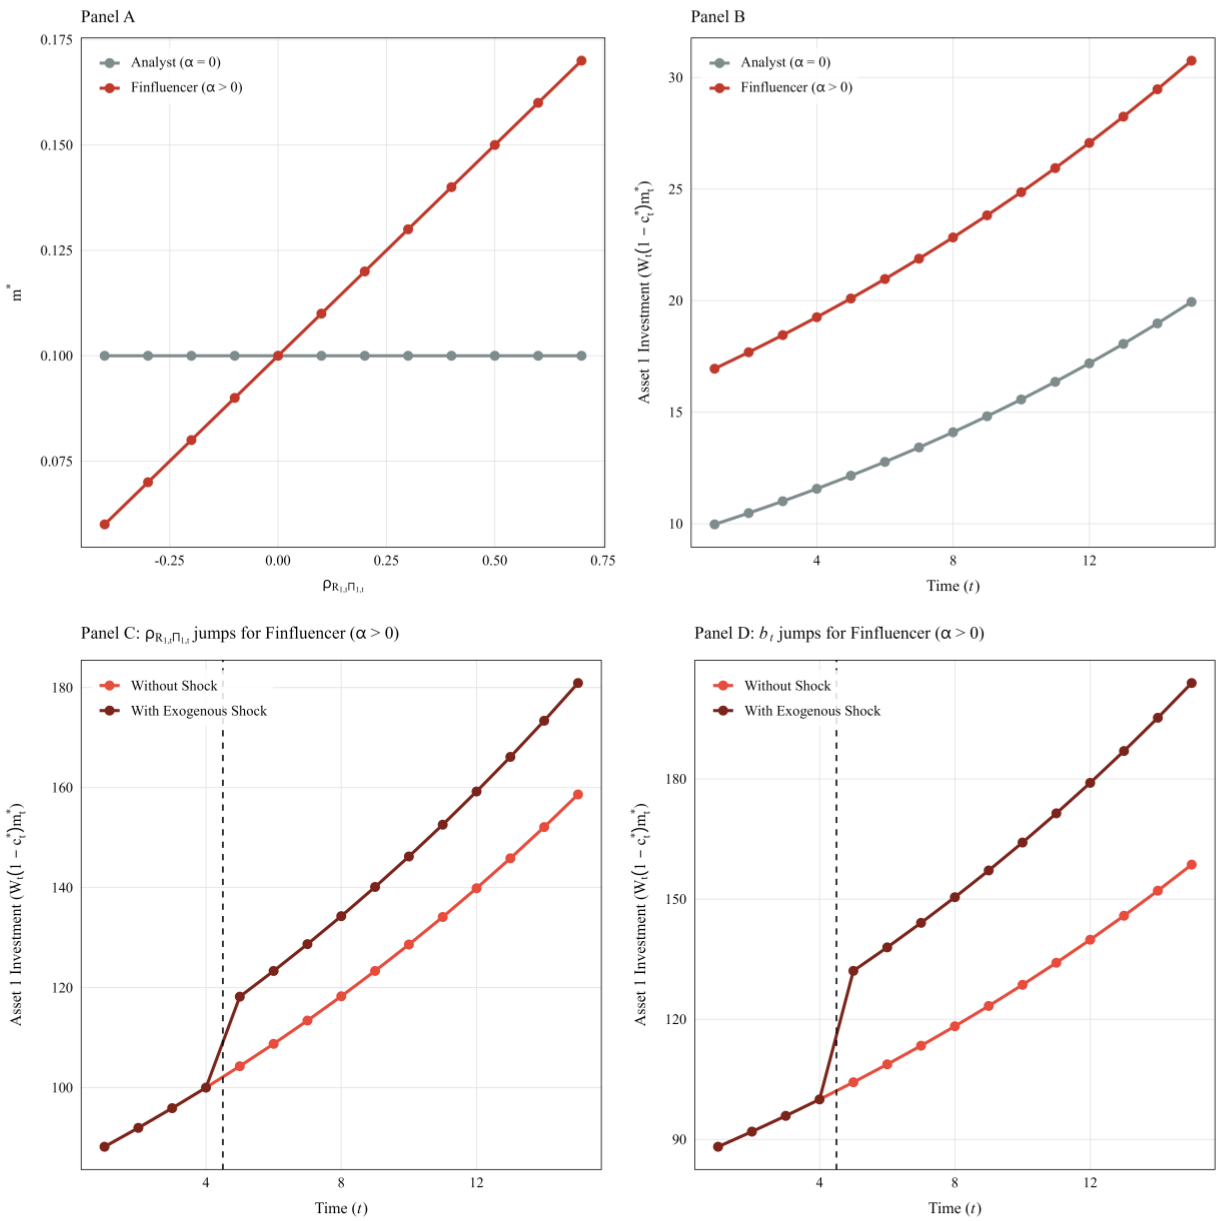

Notes: *The figure reports equilibrium trajectories from the discrete-time structural model. Panel A plots the optimal consumption-hedging portfolio weight in the risky asset S, m^*, against the correlation between returns on S and consumption-good price growth, $\rho_{R_{S,t},\Pi_{S,t}}$. Panels B–D plot absolute invested capital in S, $W_t(1-c_t^*)m_t^*$, over time t; the numéraire asset M is omitted. Panel B shows the unshocked evolution from the normalized initial condition $W_1(1-c_1^*)m_1^* = 100$. Panel C shows the adjustment to a discrete increase in $\rho_{R_{S,t},\Pi_{S,t}}$ at $t = 5$, whereas Panel D shows the adjustment to a discrete increase in the behavioral sensitivity parameter b at $t = 5$, holding $\rho_{R_{S,t},\Pi_{S,t}}$ constant. Gray denotes the benchmark specification with $\alpha = 0$; red denotes the consumption-focused specification with $\alpha > 0$. Panels C and D display only the consumption-focused specification: the lighter red line denotes the no-shock path and the darker red line the shocked path. This specification uses the posterior means from the operating-margin proxy in Table 2: $(\alpha, b, \kappa, \chi, \gamma) = (0.465, 0.273, 2.396, 0.980, 14.945)$. The exogenous parameter vector is $(\beta, \rho_{EIS}, \mu_r, \sigma_r, \mu_\pi, \sigma_\pi, R_M) = (0.99, 0.8, 0.07, 0.18, 0.02, 0.04, 1.05)$. The legend labels "Analyst" and "Finfluencer" are shorthand for model types rather than direct empirical group assignments.*

Tables

Descriptive Statistics

Table 1: ***Descriptive Statistics***

				Recommendations					Recommenders			Stocks	
	N	Mean	SD	Q05	Median	Q95		All	Finfluencers	Analysts	All	Finfluencer-	Analystonly
Dependent													
Recommendation Count (RecCount)													
Total	21,972	1	0	1	1	1	Total	10,772	98	10,674	918	252	882
Total — Finfluencers	544	1	0	1	1	1	Total — Finfluencers	98	98	0	252	252	0
Total — Analysts	21,428	1	0	1	1	1	Total — Analysts	10,674	0	10,674	882	0	882
Total Buy/Hold — Finfluencers	508	0.9338	0.2488	0.0000	1.0000	1.0000	Total Buy/Hold — Finfluencers	88	88	0	243	243	0
Total Sell — Finfluencers	36	0.06618	0.24880	0.00000	0.00000	1.00000	Total Sell — Finfluencers	21	21	0	20	20	0
Recommendation Rate (													
Daily	3,672	1.40%	2.20%	0.00%	0.13%	5.13%							
Daily— Finfluencers	1,836	0.09%	0.39%	0.00%	0.00%	0.39%							
Daily— Analysts	1,836	2.70%	2.47%	0.00%	2.56%	6.44%							
Explanatory													
Finfluencer	3,672	0.5000	0.5001	0.0000	0.5000	1.0000	Finfluencer	98	98	0	252	252	0
ConsumerLink	3,672	0.01198	0.10880	0.00000	0.00000	0.00000	ConsumerLink	314	28	286	11	10	10
ConsumerPull	3,672	0.5000	0.5001	0.0000	0.5000	1.0000	ConsumerPull	1,183	34	1,149	721	70	705
ConsumerLink (LLM)	3,672	0.6318	0.4824	0.0000	1.0000	1.0000	ConsumerLink (LLM)	7,581	82	7,499	580	154	577
ConsumerPull (CSI)	3,672	0.001939	0.011630	-0.005545	0.007050	0.009807	ConsumerPull (CSI)	0.002553	-0.010510	0.002673	0.001814	-0.007994	0.002804
Controls							**Matching Controls**						
ln(Market Value)	3,672	9.982	1.582	7.874	10.020	12.250	ln(Market Value)	Δ 1.091***	11.33	10.23	Δ 0.04954**	10.45	10.4
ln(Trading Volume)	3,672	7.366	1.624	4.685	7.461	9.861	ln(Trading Volume)	Δ 1.279***	9.098	7.819	Δ -0.006159	8.268	8.274
Additional Controls							**Continents**						
Beta (3 months)	3,672	0.8747	0.4850	0.2452	0.7992	1.8110	Europe	4,078	28	4,050	310	54	297
Conditional volatility	3,672	0.02109	0.01246	0.01159	0.01813	0.04111	America	6,751	91	6,660	603	191	581
Dividend per share	3,672	1.973	2.972	0.000	1.191	6.248	Miscellaneous	149	9	140	7	7	6
Dividend yield	3,672	2.315	2.283	0.000	1.830	6.657	***Channel***						
Earnings per share	3,672	11.47	155.70	0.00	3.29	18.28	Followers	Total 4.73m; Mean 48.2k; Median 7.85k					
ESG score	3,672	66.99	15.11	36.41	70.74	86.75	Video views	Total 37.5m; Mean 200k; Median 37.9k					
Google Trends	3,672	38.31	23.23	0.00	41.00	73.39	Video likes	Total 1.33m; Mean 7.12k; Median 802					
Momentum (1 month)	3,672	0.01974	0.05462	-0.03026	0.01299	0.08757	Comments	Total 34.7k; Mean 185.6; Median 59					
Price-to-book value	3,672	3.4410	25.2900	0.3246	2.6090	16.2400	Video length (seconds)	Total 11.8k s; Mean 63.02 s; Median 57 s					

Notes: *This table summarizes the sample at the recommendation (left panel), recommender (middle panel), and stock (right panel) levels. The left panel reports RecCount, Recommendation rate, and the triple difference sample. Means, standard deviations, medians, and selected quantiles are reported, as applicable. Variable definitions and construction are provided in Appendix Table B.1. The middle panel reports the numbers of finfluencers and analysts, both overall and by subgroup, and summarizes the matching of control stocks. It also reports differences in group means (Δ) and the corresponding Welch-test p-values; *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. The right panel reports the number of unique stocks, overall and by subgroup, as well as creator- and channel-level characteristics.*

Structural Model Calibration

Table 2: ***Structural Model Calibration***

Parameter	(1) Revenue growth $\widehat{\Pi}_{s,t} \approx \Delta$Revenue	(2) Change in operating margin $\widehat{\Pi}_{s,t} \approx \Delta$Operating Margin
α	0.392 [0.071, 0.780]	0.465 [0.103, 0.860]
b	0.279 [0.111, 0.545]	0.273 [0.108, 0.536]
κ	2.388 [1.186, 4.654]	2.396 [1.205, 4.389]
χ	0.980 [0.597, 1.503]	0.980 [0.622, 1.452]
γ	15.345 [9.777, 23.857]	14.945 [9.887, 23.105]
λ	3.396 [2.457, 4.699]	3.465 [2.481, 4.698]
LOOIC	636.9	635.6

Notes*: This table reports posterior means of the structural-model parameters estimated by Bayesian inference. Columns (1) and (2) proxy firm-level goods-price inflation, $\Pi_{S,t}$, using quarterly revenue growth and the quarterly change in operating margin, respectively. Both specifications are estimated on the same sample using quarterly observations from January 2023 through March 2025. Bracketed values are 95% equal-tailed posterior credible intervals. LOOIC denotes the leave-one-out information criterion. Section 2 sets out the theoretical model; Appendix B documents estimation of the empirical counterparts to its theoretical objects; and Appendix C provides the full calibration and model-implementation procedure.*

Main Triple-Difference Evidence

Table 3: ***Triple-Difference Estimates and Consumption-Hedge Channel***

	ConsumerLink				ConsumerLink (LLM)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ConsumerLink × ConsumerPull × Finfluencer		0.702**	0.667**	0.672**		0.430	0.506*	0.503*
		(0.307)	(0.303)	(0.303)		(0.269)	(0.272)	(0.270)
ConsumerLink × Finfluencer	2.156***	1.699***	1.822***	1.823***	0.054	-0.071	-0.206	-0.211
	(0.144)	(0.110)	(0.088)	(0.087)	(0.196)	(0.182)	(0.162)	(0.163)
ConsumerPull × Finfluencer		0.986***	1.026***	1.035***		0.702***	0.700***	0.720***
		(0.133)	(0.133)	(0.128)		(0.227)	(0.233)	(0.228)
ConsumerLink × ConsumerPull		-0.838**	-0.809**	-0.828**		-0.193**	-0.225***	-0.168**
		(0.378)	(0.394)	(0.392)		(0.079)	(0.080)	(0.079)
ConsumerLink	0.620***	0.905***	-0.345***	-0.345***	0.126**	0.195***	0.183***	0.168***
	(0.055)	(0.049)	(0.125)	(0.127)	(0.051)	(0.051)	(0.044)	(0.044)
ConsumerPull		-0.123***	-0.118***	0.954***		-0.003	0.023	1.228***
		(0.035)	(0.042)	(0.365)		(0.069)	(0.074)	(0.366)
Finfluencer	-3.731***	-3.917***	-4.046***	-4.048***	-3.689***	-3.798***	-3.865***	-3.867***
	(0.103)	(0.094)	(0.069)	(0.069)	(0.147)	(0.137)	(0.129)	(0.131)
Controls			✓	✓			✓	✓
Interacted Controls				✓				✓
Observations	3672	3672	3672	3672	3672	3672	3672	3672
R^2	0.251	0.254	0.290	0.292	0.246	0.250	0.288	0.290
BIC	13312.8	13285.1	12750.5	12800.5	13399.6	13365.7	12773.3	12826.2

Notes: *This table reports triple-difference estimates from Negative Binomial regressions of RecCount on consumption-hedge stocks, as measured by ConsumerLink in Columns (1)–(4) and ConsumerLink (LLM) in Columns (5)–(8), respectively. ConsumerPull equals one if the recommendation is published before a defined consumption event. Finfluencer equals one for TikTok finfluencer recommendations and zero for analyst recommendations. The logarithm of exposure days is included as an offset, such that the model estimates recommendation rates per exposure day. Appendix Table B.1 provides variable descriptions. Coefficients are reported as logarithms of incidence rate ratios. Standard errors are robust, clustered by firm, and reported in parentheses. Controls and controls interacted with ConsumerPull are indicated in the table. *, **, and *** denote significance at the 10%, 5%, and 1% levels..*

Table 4: ***Supplementary Triple-Difference Estimates and Consumption-Hedge Channel***

	ConsumerLink				ConsumerLink (LLM)			
Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
ConsumerLink × ConsumerPull (CSI)× Finfluencer		37.187***	42.087***	44.222***		33.498***	33.547***	37.140***
		(10.133)	(12.695)	(13.212)		(10.300)	(11.626)	(11.592)
ConsumerLink × Finfluencer	1.402***	1.337***	1.231***	1.231***	-0.012	-0.034	-0.097	-0.105
	(0.177)	(0.180)	(0.162)	(0.164)	(0.192)	(0.195)	(0.174)	(0.175)
ConsumerPull (CSI)× Finfluencer		-25.401***	-27.503***	-29.997***		-25.223***	-29.421***	-32.082***
		(6.977)	(9.645)	(8.354)		(7.841)	(10.111)	(8.311)
ConsumerLink × ConsumerPull (CSI)		-2.532	0.418	-1.791		-0.769	0.030	-2.665
		(9.298)	(7.594)	(4.770)		(4.318)	(3.880)	(3.259)
ConsumerLink	0.505***	0.507***	0.054	0.053	0.137**	0.134**	0.107**	0.109**
	(0.075)	(0.074)	(0.066)	(0.065)	(0.053)	(0.055)	(0.048)	(0.046)
ConsumerPull (CSI)		6.094***	7.995***	-23.878***		6.496*	7.801**	-27.243***
		(2.239)	(2.045)	(8.563)		(3.927)	(3.032)	(8.829)
Finfluencer	-4.071***	-4.045***	-4.090***	-4.088***	-3.700***	-3.698***	-3.772***	-3.769***
	(0.091)	(0.092)	(0.086)	(0.087)	(0.145)	(0.150)	(0.142)	(0.142)
Controls			✓	✓			✓	✓
Interacted Controls				✓				✓
Observations	3672	3672	3672	3672	3672	3672	3672	3672
R^2	0.244	0.246	0.278	0.280	0.225	0.227	0.271	0.273
BIC	14909.9	14911.1	14355.7	14409.8	15276.9	15278.6	14504.5	14555.5

*Notes: This table reports triple-difference estimates from Negative Binomial regressions of RecCount on consumption-hedge stocks, as measured by ConsumerLink in Columns (1)–(4) and ConsumerLink (LLM) in Columns (5)–(8), respectively. ConsumerPull (CSI) measures the log-change in the Consumer Sentiment Index during the recommendation relative to the preceding period for a recommender. Finfluencer equals one for TikTok finfluencer recommendations and zero for analyst recommendations. The logarithm of exposure days is included as an offset, such that the model estimates recommendation rates per exposure day. Appendix Table B.1 provides variable descriptions. Coefficients are reported as logarithms of incidence rate ratios. Standard errors are robust, clustered by firm, and reported in parentheses. Controls and controls interacted with ConsumerPull are indicated in the table. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Placebo Timing Tests

Table 5: ***Placebo Tests for the Triple-Difference Estimate***

	ConsumerLink				ConsumerLink (LLM)			
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Constant	-3.857***	-7.586***	-7.550***	-7.556***	-3.983***	-7.806***	-7.711***	-7.268***
	(0.029)	*(0.371)*	*(0.310)*	*(0.310)*	*(0.046)*	*(0.335)*	*(0.269)*	*(0.417)*
ConsumerLink	1.362***	0.017	0.040	36.685***	0.211***	0.218***	0.195***	-0.964*
	(0.046)	*(0.131)*	*(0.128)*	*(6.845)*	*(0.059)*	*(0.061)*	*(0.057)*	*(0.528)*
Placebo ConsumerPull	0.476***	0.490***	0.397	0.404	0.502***	0.523***	0.338	-0.210
	(0.036)	*(0.038)*	*(0.507)*	*(0.508)*	*(0.069)*	*(0.071)*	*(0.447)*	*(0.776)*
Finfluencer	-3.654***	-3.654***	-3.654***	-3.654***	-3.516***	-3.516***	-3.516***	-3.516***
	(0.112)	*(0.112)*	*(0.112)*	*(0.112)*	*(0.152)*	*(0.153)*	*(0.153)*	*(0.153)*
ConsumerLink × Placebo ConsumerPull	0.174	0.197	0.153	0.400***	-0.034	-0.048	-0.010	0.784
	(0.136)	*(0.128)*	*(0.194)*	*(0.129)*	*(0.080)*	*(0.083)*	*(0.092)*	*(0.906)*
ConsumerLink × Finfluencer	1.708***	1.708***	1.708***	1.708***	-0.097	-0.097	-0.097	-0.097
	(0.121)	*(0.121)*	*(0.122)*	*(0.122)*	*(0.209)*	*(0.210)*	*(0.210)*	*(0.211)*
Placebo ConsumerPull × Finfluencer	-0.088	-0.088	-0.088	-0.088	-0.065	-0.065	-0.065	-0.065
	(0.112)	*(0.113)*	*(0.113)*	*(0.113)*	*(0.226)*	*(0.227)*	*(0.227)*	*(0.228)*
Placebo ConsumerPull × ConsumerLink × Finfluencer	-0.172	-0.172	-0.172	-0.172	-0.041	-0.041	-0.041	-0.041
	(0.165)	*(0.165)*	*(0.165)*	*(0.165)*	*(0.249)*	*(0.249)*	*(0.250)*	*(0.250)*
Controls		✓		✓		✓		✓
Interacted controls			✓	✓			✓	✓
Num. Obs.	3672	3672	3672	3672	3672	3672	3672	3672
AIC	497.8	577.6	667.6	716.8	499.2	577.4	667.5	846.7

Notes*: This table reports placebo triple-difference estimates from Negative Binomial regressions of RecCount on consumption-hedge stocks, as measured by ConsumerLink in Columns (1)–(4) and ConsumerLink (LLM) in Columns (5)–(8), respectively. ConsumerPull is randomly assigned at the calendar-week level, with half of the calendar weeks assigned to each state. Finfluencer equals one for TikTok finfluencer recommendations and zero for analyst recommendations. The logarithm of exposure days is included as an offset, such that the model estimates recommendation rates per exposure day. Appendix Table B.1 provides variable descriptions. Coefficients are reported as logarithms of incidence rate ratios. Standard errors are robust, clustered by firm, and reported in parentheses. Controls and controls interacted with ConsumerPull are indicated in the table. Results using the ConsumerPull (CSI) specification are qualitatively similar and are not tabulated. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Recommendation Activity around Events

Table 6: ***Recommendation Counts around Consumption and Earnings-Announcement Windows***

Variable	**(1)**	**(2)**	**(3)**	**(4)**
(Constant)	-2.200***	-6.067***	-1.527***	-5.221***
	(0.014)	(0.121)	(0.014)	(0.126)
Consumption Window	-0.143***	-0.137***	-0.184***	-0.165***
	(0.038)	(0.036)	(0.038)	(0.037)
Earnings Announcement Window	0.885***	0.781***	0.891***	0.790***
	(0.052)	(0.050)	(0.053)	(0.050)
Finfluencer			-3.880***	-10.672***
			(0.060)	(0.465)
Finfluencer × Consumption Window			1.212***	1.074***
			(0.171)	(0.144)
Finfluencer × Earnings Announcement Window			-0.485**	-0.453*
			(0.217)	(0.232)
Controls		✓		✓
Num.Obs.	208632	208632	208632	208632
Pseudo-R^2	0.003	0.025	0.145	0.174
BIC	148126.8	145014.0	127142.0	123051.4
AIC	148096.0	144870.5	127080.5	122764.5

Notes*: This table reports negative binomial fixed-effects estimates of weekly RecCount around consumption and earnings-announcement windows. Consumption Window equals one during the Christmas and Black Friday pre-period. Earnings-Announcement Window is an indicator if firm reported financial information. Variable definitions and construction are provided in Appendix Table B.1. Standard errors are clustered and reported in parentheses. Pseudo-R-squared, AIC, and BIC are reported as model-fit statistics. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Lagged Recommendation Activity

Table 7: ***Lagged Within- and Cross-Group Recommendation Counts***

Variable	Forward Analyst RecCount		Forward Finfluencer RecCount	
	(1)	(2)	(3)	(4)
Lagged Analyst RecCount	0.505***	0.371***	0.415***	-0.034
	(0.030)	(0.029)	(0.115)	(0.107)
Lagged Finfluencer RecCount	0.528***	-0.011	3.034***	1.413***
	(0.092)	(0.094)	(0.234)	(0.212)
Controls		✓		✓
Num.Obs.	746951	746951	746951	746951
BIC	897680	881745	46000	39954
Pseudo-R^2	0.003	0.024	0.042	0.171

Notes: *This table reports Poisson pseudo-maximum-likelihood (PPML) estimates of the association between seven-day-ahead RecCount and its seven-day lagged value for analysts and finfluencers. Variable definitions and construction are provided in Appendix Table B.1. Driscoll–Kraay standard errors, using a bandwidth of seven days, are reported in parentheses. Pseudo-R^2, AIC, and BIC are reported as model-fit statistics. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Geographic Frictions and Home Bias

Table 8: ***Recommendation Counts and Geographic and Cultural Distance***

Variable	**(1)**	**(2)**	**(3)**	**(4)**
(Constant)	15.877***	16.854***	9.906***	10.745***
	(1.486)	(1.604)	(1.181)	(1.251)
log(Distance)	-1.127***	-1.150***	-1.153***	-1.178***
	(0.145)	(0.155)	(0.147)	(0.157)
English-Speaking Market	-0.872*	-0.990	-0.870	-0.989
	(0.521)	(0.626)	(0.544)	(0.656)
Common Language	0.049	0.103	0.152	0.225
	(0.517)	(0.602)	(0.581)	(0.678)
Home Country	-6.076***	-6.308***	-6.319***	-6.583***
	(1.247)	(1.322)	(1.278)	(1.365)
Non-U.S. Listing	-6.179***	-6.332***		
	(0.749)	(0.845)		
Finfluencer		-8.820***		-7.093***
		(2.440)		(2.059)
Finfluencer × log(Distance)		0.462*		0.479*
		(0.252)		(0.259)
Finfluencer × English Speaking		1.475*		1.476*
		(0.759)		(0.790)
Finfluencer × Common Language		-0.017		-0.160
		(0.757)		(0.875)
Finfluencer × Home Country		4.172**		4.446**
		(2.012)		(2.053)
Finfluencer × Non-U.S. Listing		1.865*		
		(1.057)		
Num.Obs.	756	756	720	720
Pseudo-R^2	0.100	0.155	0.062	0.120
BIC	1967.8	1889.9	1582.2	1519.5
AIC	1940.0	1834.3	1559.3	1473.7

Notes*: This table reports negative binomial fixed-effects estimates of the association between RecCount, the dependent variable, and alternative distance measures. Variable definitions and construction are provided in Appendix Table B.1. Standard errors are clustered by firm and reported in parentheses. Pseudo-R-squared, AIC, and BIC are reported as model-fit statistics. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Returns in the Common-Stock Sample

Table 9: ***Finfluencer and Analyst Recommendation Returns for Commonly Recommended Stocks***

Horizon	Group	Return	CAPM alpha	FF3 alpha	FF5 alpha	FF6 alpha	N
1 Day	Finfluencers	34.52 (41.47)	19.35 (42.46)	38.66 (44.45)	26.71 (41.83)	23.26 (40.31)	543
	Analysts	46.03*** (5.90)	38.90*** (5.45)	42.07*** (5.40)	42.59*** (5.44)	43.24*** (5.48)	7810
	Δ	-11.50 (41.89)	-19.55 (42.81)	-3.41 (44.78)	-15.88 (42.18)	-19.98 (40.68)	
1 Week	Finfluencers	13.78 (18.88)	8.85 (16.97)	13.13 (16.65)	16.22 (15.93)	13.77 (15.69)	543
	Analysts	17.80*** (2.48)	7.83*** (2.24)	8.00*** (2.20)	9.79*** (2.21)	11.28*** (2.21)	7810
	Δ	-4.03 (19.04)	1.02 (17.12)	5.13 (16.79)	6.43 (16.09)	2.49 (15.85)	
1 Month	Finfluencers	27.66*** (10.61)	14.33 (9.91)	15.90 (9.78)	15.09 (9.74)	14.52 (9.68)	543
	Analysts	11.46*** (1.13)	0.97 (1.02)	2.51** (1.00)	3.19*** (1.01)	4.22*** (1.01)	7720
	Δ	16.20 (10.67)	13.36 (9.96)	13.39 (9.83)	11.90 (9.80)	10.30 (9.73)	
1 Quarter	Finfluencers	10.16** (4.01)	5.18 (3.73)	7.76** (3.82)	6.73* (3.76)	6.68* (3.75)	542
	Analysts	9.94*** (0.66)	-2.67*** (0.60)	-1.15* (0.59)	-0.58 (0.60)	0.26 (0.60)	7440
	Δ	0.22 (4.06)	7.85** (3.78)	8.91** (3.87)	7.30* (3.81)	6.42* (3.80)	

Notes: *This table reports the performance of finfluencer and analyst recommendations over one-day, one-week, one-month, and one-quarter holding periods, following Kakhbod et al. (2023). Performance is measured using raw and risk-adjusted returns. The sample is restricted to long and short recommendations (excluding hold recommendations) for stocks recommended by both finfluencers and analysts. The Δ column reports differences in mean performance between finfluencers and analysts. Variable definitions and construction are provided in Appendix Table B.1. Heteroskedasticity- and autocorrelation-consistent (HAC) standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Post-Recommendation Returns

Table 10: ***Finfluencer and Analyst Recommendation Returns for Recommended Stocks***

Horizon	Group	Return	CAPM α	FF3 alpha	FF5 alpha	FF6 alpha	N
1 Day	Finfluencers	34.52 (41.47)	19.35 (42.46)	38.66 (44.45)	26.71 (41.83)	23.26 (40.31)	543
	Analysts	47.45*** (5.56)	38.02*** (5.11)	40.14*** (5.06)	39.61*** (5.09)	41.22*** (5.10)	12198
	Δ	-12.93 (41.84)	-18.67 (42.76)	-1.48 (44.74)	-12.90 (42.14)	-17.96 (40.63)	
1 Week	Finfluencers	13.78 (18.87)	8.85 (16.97)	13.13 (16.64)	16.22 (15.93)	13.77 (15.69)	543
	Analysts	17.91*** (2.40)	7.16*** (2.16)	7.05*** (2.13)	8.40*** (2.14)	9.69*** (2.12)	12198
	Δ	-4.13 (19.03)	1.69 (17.11)	6.08 (16.78)	7.82 (16.08)	4.08 (15.83)	
1 Month	Finfluencers	27.66*** (10.60)	14.33 (9.90)	15.90 (9.78)	15.09 (9.74)	14.52 (9.68)	543
	Analysts	13.94*** (1.11)	2.42** (1.00)	3.64*** (0.99)	3.89*** (0.99)	4.63*** (0.99)	12108
	Δ	13.72 (10.66)	11.91 (9.95)	12.26 (9.83)	11.19 (9.79)	9.89 (9.73)	
1 Quarter	Finfluencers	10.16** (4.01)	5.18 (3.73)	7.76** (3.82)	6.73* (3.76)	6.68* (3.75)	542
	Analysts	14.54*** (0.68)	-0.28 (0.61)	1.11* (0.61)	1.27** (0.62)	1.99*** (0.62)	11826
	Δ	-4.38 (4.06)	5.46 (3.78)	6.65* (3.87)	5.46 (3.81)	4.69 (3.80)	

Notes: *This table reports the performance of finfluencer and analyst recommendations over one-day, one-week, one-month, and one-quarter holding periods, following Kakhbod et al. (2023). Performance is measured using raw and risk-adjusted returns. The sample includes long and short recommendations (excluding hold recommendations) for all covered stocks, including stocks recommended exclusively by analysts. The Δ column reports differences in mean performance between finfluencers and analysts. Variable definitions and construction are provided in Appendix Table B.1. Heteroskedasticity- and autocorrelation-consistent (HAC) standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Option-Implied Risk

Table 11: ***Recommendation-Implied Volatility (IV) and Smile Slopes***

Horizon	Group	20Δ Put IV	ATM IV	20Δ Call IV	Put-Wing Slope	Call-Wing Slope	N
1 Month	Finfluencers	34.83*** (1.66)	32.27*** (1.57)	32.74*** (1.72)	2.47*** (0.25)	0.81*** (0.28)	486
	Analysts	24.95*** (0.32)	22.32*** (0.29)	22.69*** (0.31)	2.64*** (0.06)	0.34*** (0.05)	10791
	Δ	9.88*** (1.69)	9.96*** (1.60)	10.05*** (1.75)	-0.17 (0.26)	0.47* (0.28)	
1 Quarter	Finfluencers	34.33*** (1.57)	31.30*** (1.47)	31.65*** (1.63)	2.84*** (0.25)	0.48* (0.25)	486
	Analysts	24.97*** (0.31)	22.26*** (0.28)	22.22*** (0.30)	2.69*** (0.05)	-0.06 (0.04)	10791
	Δ	9.36*** (1.60)	9.04*** (1.50)	9.43*** (1.65)	0.15 (0.26)	0.54** (0.26)	

Notes: *This table reports option-implied volatility for stocks recommended by TikTok finfluencers and professional analysts. The first three outcome columns report annualized implied volatility, in percentage points, for 20-delta out-of-the-money puts, at-the-money options, and 20-delta out-of-the-money calls. Put-Wing Slope is the 20-delta put IV minus ATM IV; Call-Wing Slope is the 20-delta call IV minus ATM IV. Results are reported for one-month and one-quarter maturities. Finfluencer − Analyst rows report differences between the two groups. Variable definitions and construction are provided in Appendix Table B.1. Heteroskedasticity-robust standard errors are reported in parentheses. N denotes the number of recommendation-option matches. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Engagement Heterogeneity

Table 12: ***Finfluencers' Consumption-Hedge Channel and Audience Engagement***

	ConsumerLink			ConsumerLink (LLM)		
Variable	**Views per Follower**	**Likes per Follower**	**Comments per Follower**	**Views per Follower**	**Likes per Follower**	**Comments per Follower**
Constant	2.057***	-1.848***	-4.930***	0.491**	-3.239***	-6.465***
	(0.198)	(0.208)	(0.220)	(0.217)	(0.229)	(0.236)
Large-Audience Finfluencer	-1.649***	-1.473***	-1.602***			
	(0.204)	(0.215)	(0.198)			
ConsumerLink	-0.170	-0.142	-0.192	-0.046	-0.059	-0.046
	(0.168)	(0.172)	(0.150)	(0.045)	(0.046)	(0.040)
Large-Audience Finfluencer × ConsumerLink	0.476**	0.441*	0.488**			
	(0.224)	(0.229)	(0.197)			
Small-Audience Finfluencer				1.503***	1.337***	1.470***
				(0.194)	(0.204)	(0.190)
ConsumerLink				0.085	0.086	0.131**
				(0.068)	(0.068)	(0.062)
Num. Obs.	544	544	544	544	544	544
Marginal R^2	0.160	0.102	0.077	0.147	0.093	0.073
Conditional R^2	0.900	0.756	0.567	0.897	0.752	0.561
AIC	13057.2	9016.1	5600.8	13060.8	9018.7	5603.2

Notes: *This table reports negative binomial generalized linear mixed-model estimates of engagement on stock-recommendation videos. The dependent variables are views, likes, and comments per follower. The natural logarithm of the creator's follower count enters as an offset, so coefficients describe engagement rates per follower. All specifications include creator-level random intercepts. Coefficients are reported in log incidence-rate units. Variable definitions and construction are provided in Appendix Table B.1. Standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.*

Appendix

The supplementary material reports additional distributions, alternative specifications, coverage analyses, and variable definitions that complement the main results.

A Supplementary Figures

Figure A.1: ***Associations between proxies for $\rho_{R,\Pi}$ and ω***

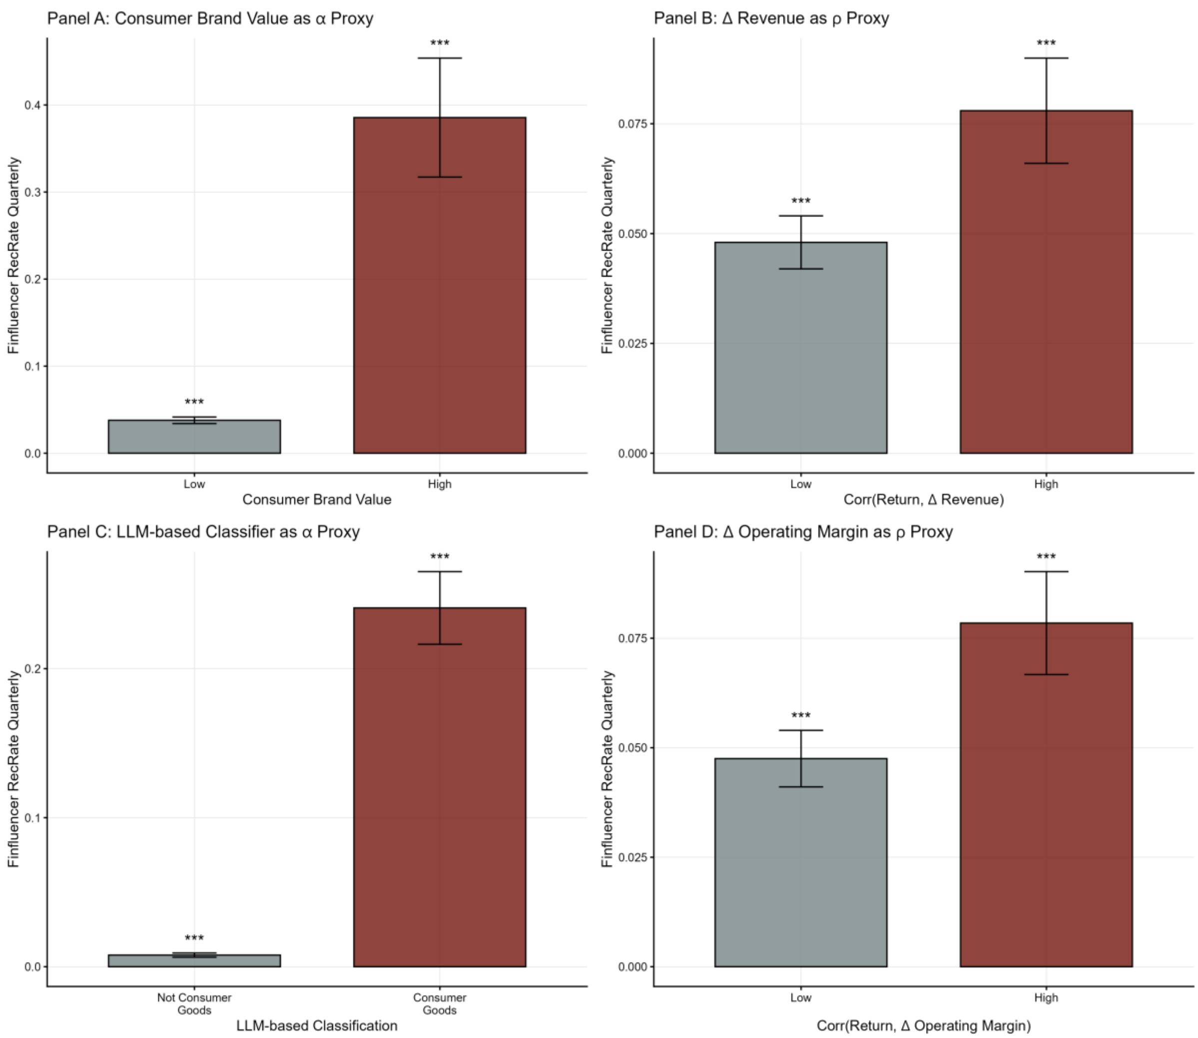

Notes: *This figure reports mean quarterly finfluencer recommendation rate (RecRate, a measure for ω) across groups defined by four firm-characteristics. Panels A and C compare firms by consumer brand value and by an LLM-based classification of consumer-product firms, respectively. Panels B and D compare firms with high versus low quarterly correlations between stock returns and changes in revenues or operating margins, respectively; these correlations proxy for $\rho_{R,\Pi}$. Bars represent group means, and asterisks indicate differences between groups that are significant at the 10%, 5%, and 1% levels; variable definitions and construction are provided in Appendix Table E.1.*

B Supplementary Tables

Table B.1: ***Variable Definitions and Data Sources***

Variable	Definition, theoretical construct, and measurement	Source
I. Reduced-Form Triple-Difference Model		
I.1 Recommendation Outcomes and Recommender Type		
$Recommendation_{r,s,t}$	Identifies an explicit, firm-specific Buy or Sell recommendation for stock s issued by recommender type r at time t. Hold assessments are collected but are used only when constructing post-recommendation performance measures; they do not define the core Buy/Sell recommendation outcome.	[1,2,8]
$RecCount_{r,s,t}$	Measures the number of recommendations for stock s issued by recommender type r during the defined period. It is constructed by summing the individual Buy/Sell recommendation indicators within recommender-type–stock–period cells.	[1,2,8]
$Finfluencer_i$	Identifies the recommender group expected to place greater weight on consumption objectives. It equals one for a TikTok finfluencer recommendation and zero for a professional analyst recommendation.	[1,2,8]
I.2 Consumption-Hedge Exposure and Consumer Pull		
$ConsumerLink_s$	Measures whether stock s is plausibly connected to economically meaningful future household consumption. It equals one when the firm's Interbrand Brand Value exceeds USD 50 billion and zero otherwise.	[4,8]
$ConsumerLink_s^{LLM}$	Measures a broader form of direct exposure to household consumption. It equals one when gpt-oss-120b classifies the firm's standardized LSEG description as primarily supplying finished goods or services to end consumers, and zero for upstream, business-to-business, financial, holding, or unclassifiable firms.	[2,8,9]
$ConsumerPull_t$	Measures calendar periods in which consumption motives and consumption-related FOMO are especially salient. It equals one during the four weeks preceding Christmas and the week preceding Black Friday and zero otherwise.	[8]
$ConsumerPull_{s,t}^{CSI}$	Measures a stock-specific consumer-sentiment state and the magnitude of the associated reference-point shift. It is constructed from the period-to-period change in the IPSOS Primary Consumer Sentiment Index; a stock-specific median separates higher- from lower-sentiment phases, while the observed change in sentiment is retained rather than collapsed into a binary indicator.	[6,8]
I.3 Financial, Preference, and Attention Controls		
$Market\ Capitalization_{s,t}$	Measures firm size and visibility. It is constructed as the natural logarithm of the firm's market value.	[3,8]
$Trading\ Volume_{s,t}$	Measures stock liquidity and trading attention. It is constructed as the natural logarithm of share trading volume.	[3,8]
$Beta_{s,t}$	Measures the stock's exposure to systematic market risk. It is measured using the three-month market beta.	[3]
$Momentum_{s,t}$	Measures recent return continuation and the financial attractiveness of the stock. It is measured as the stock's return over the preceding one-month horizon.	[3,8]

$Conditional\ Volatility_{s,t}$	Measures time-varying firm-specific return risk and attention. It is estimated as the conditional volatility from a GARCH(1,1) model.	[3,8]
$Dividends\ per\ Share_{s,t}$	Measures the firm's cash distribution per outstanding share. It is measured as dividends paid per share.	[3]
$Dividend\ Yield_{s,t}$	Measures dividend income relative to the stock price. It is measured as the firm's dividend per share divided by its market price.	[3]
$Earnings\ per\ Share_{s,t}$	Measures accounting profitability attributable to each outstanding share. It is measured as earnings per share.	[3]
$Price\text{-}to\text{-}Book\ Ratio_{s,t}$	Measures market valuation relative to accounting book value. It is measured as the stock price divided by book value per share.	[3]
$ESG\ Score_{s,t}$	Measures nonfinancial firm characteristics that may enter investor preferences. It is measured using the LSEG composite environmental, social, and governance score.	[2]
$Google\ Search\ Intensity_{s,t}$	Measures stock-level public attention. It is measured using the Google Trends search-intensity index associated with the firm or stock.	[5,8]
VIX_t	Measures market-wide risk perceptions and investor attention. It is measured using the CBOE Volatility Index.	[3]
II. Robustness, Mechanisms, and Heterogeneity		
II.1 Event and Placebo Tests		
$Consumption\ Window_t$	Measures the consumption-related event period used in the event regressions. It applies the same Christmas and Black Friday pre-period definition as ConsumerPull.	[8]
$Earnings\ Announcement\ Wi$	Measures periods of firm-specific financial information production. It equals one in the week after the earnings-announcement and zero otherwise.	[2,8]
$Placebo\ ConsumerPull_t$	Measures placebo timing unrelated to the theoretically predicted consumption window. It is constructed by assigning alternative windows of comparable length and re-estimating the triple-difference specification.	[8]
II.2 Lagged and Forward Recommendation Activity		
$Lagged\ Analyst\ RecCount_{s,t}$	Measures recent professional-analyst recommendation activity for stock s. It is constructed as log(1 + RecCount) accumulated from day t−6 through day t.	[2,8]
$Lagged\ Finfluencer\ RecCou$	Measures recent finfluencer recommendation activity for stock s. It is constructed as log(1 + RecCount) accumulated from day t−6 through day t.	[1,8]
$Forward\ Analyst\ RecCount_{s}$	Measures subsequent professional-analyst recommendation activity for stock s. It is constructed as RecCount accumulated from day t+1 through day t+7.	[2,8]
$Forward\ Finfluencer\ RecCo$	Measures subsequent finfluencer recommendation activity for stock s. It is constructed as RecCount accumulated from day t+1 through day t+7.	[1,8]
II.3 Geographic and Cultural Proximity		
$Geographic\ Distance_{i,s}$	Measures geographic frictions between recommender i and the recommended firm. It is constructed as the natural logarithm of the distance in kilometers between the recommender's home market and the firm's headquarters.	[1,2,8]

$English\ Speaking\ Market_s$	Measures familiarity with an English-language market environment. It equals one when the relevant firm market is classified as English-speaking and zero otherwise.	[8]
$Common\ Language_{i,s}$	Measures linguistic proximity between the recommender and the firm market. It equals one when the two locations share a common language and zero otherwise.	[8]
$Home\ Country_{i,s}$	Measures domestic-market familiarity and home bias. It equals one when the recommender and the recommended firm are located in the same country and zero otherwise.	[1,2,8]
$NonUS\ Listing_s$	Measures whether the recommended security is listed outside the United States. It equals one for a non-U.S. listing and zero for a U.S. listing.	[2,8]
II.4 Post-Recommendation Performance		
$Return_{s,t+h}$	Measures unadjusted performance following a recommendation. It is calculated separately for h ∈ {1D, 1W, 1M, 1Q}, corresponding to the next day, week, month, and quarter, and is reported on the paper's daily-equivalent scale.	[3,8]
$CAPM\ Alpha_{t+h}$	Measures post-recommendation performance beyond exposure to the market factor. It is estimated separately for h ∈ {1D, 1W, 1M, 1Q} using the Capital Asset Pricing Model.	[3,7,8]
$FF3\ Alpha_{t+h}$	Measures post-recommendation performance after adjustment for the market, size, and value factors. It is estimated separately for h ∈ {1D, 1W, 1M, 1Q} using the Fama–French three-factor model.	[3,7,8]
$FF5\ Alpha_{t+h}$	Measures post-recommendation performance after adjustment for the market, size, value, profitability, and investment factors. It is estimated separately for h ∈ {1D, 1W, 1M, 1Q} using the Fama–French five-factor model.	[3,7,8]
$FF6\ Alpha_{t+h}$	Measures post-recommendation performance after adjustment for the six-factor specification used in the paper. It is estimated separately for h ∈ {1D, 1W, 1M, 1Q}.	[3,7,8]
II.5 Option-Implied Risk		
$\mathrm{IV}^{\tau}_{\Delta 20,Put}$	Measures the market-implied cost of downside protection for the recommended stock. It is the closing-midpoint implied volatility of a 20-delta out-of-the-money put for maturity τ ∈ {1M, 1Q}.	[2]
IV^{τ}_{ATM}	Measures the central level of market-implied stock-return uncertainty. It is the at-the-money implied volatility for maturity τ ∈ {1M, 1Q}.	[2]
$\mathrm{IV}^{\tau}_{\Delta 20,Call}$	Measures the option-implied price of upside risk for the recommended stock. It is the closing-midpoint implied volatility of a 20-delta out-of-the-money call for maturity τ ∈ {1M, 1Q}.	[2]
$Put\ Wing\ Slope^{\tau}$	Measures downside asymmetry in the implied-volatility smile rather than a literal tail-loss probability. It is constructed as 20-delta put implied volatility minus at-the-money implied volatility for τ ∈ {1M, 1Q}; using delta points makes the measure comparable across stocks despite differences in prices and strikes.	[2,8]
$Call\ Wing\ Slope^{\tau}$	Measures upside asymmetry and lottery-like upside potential in the implied-volatility smile. It is constructed as 20-delta call implied volatility minus at-the-money implied volatility for τ ∈ {1M, 1Q}; using delta points makes the measure comparable across stocks despite differences in prices and strikes.	[2,8]
II.6 Engagement and Finfluencer Heterogeneity		

$Follower\ Count_i$	Measures the potential audience size of finfluencer i. It is recorded as the finfluencer's number of followers and enters the engagement models as a natural-log offset.	[1,8]
$Large\text{-}audience\ Finfluencer_i$	Identifies finfluencers with comparatively large audiences. It equals one when the finfluencer's follower count exceeds the analysis-sample median of 19,100 and zero otherwise.	[1,8]
$Small\text{-}audience\ Finfluencer$	Identifies finfluencers with comparatively small audiences. It equals one when the finfluencer's follower count is at or below the analysis-sample median of 19,100 and zero otherwise.	[1,8]
$Video\ Views_{i,s,t}$	Measures the reach of a stock-recommendation video. It is recorded as the number of video views and modeled as an engagement rate per follower.	[1,8]
$Video\ Likes_{i,s,t}$	Measures positive audience engagement with a recommendation video. It is recorded as the number of video likes and modeled as an engagement rate per follower.	[1,8]
$Video\ Comments_{i,s,t}$	Measures active audience engagement with a recommendation video. It is recorded as the number of video comments and modeled as an engagement rate per follower.	[1,8]
$Video\ Length_{i,s,t}$	Measures the duration of the recommendation content. It is recorded as the length of the TikTok video in seconds and reported as a descriptive channel characteristic.	[1]
III. Structural Model — Empirical Proxies		
III.1 Observed Portfolio Allocation		
$\omega^{\mathrm{obs}}_{s,j,k} \in \{0,1\}$	Measures the recommendation-based allocation decision for stock s in observed decision k of quarter j. It is constructed from the recommendation record and equals one when the decision allocates to stock s and zero when it allocates to market asset M, thereby providing the observed binary counterpart to the model's ω = 1 versus ω = 0 comparison.	[1,8]
$y_j = \sum_{k=1}^{K_j} \omega^{\mathrm{obs}}_{s,j,k}$	Measures the quarterly number of observed portfolio allocations to stock s and is the dependent count in the structural calibration. It sums the recommendation-based allocation indicators across the K_j observed decisions in quarter j and enters a zero-inflated binomial likelihood, with an additional probability of zero allocations.	[1,8]
$\hat{\omega}_j = y_j/K_j$	Measures the observed quarterly allocation rate and operationalizes the theoretical portfolio weight ω. It normalizes the recommendation-based allocation count by the number of observed decisions; the structural logit link maps the value difference between full allocation to stock s, ω = 1, and full allocation to market asset M, ω = 0, onto the corresponding allocation probability.	[1,8]
$Brand\ Value_{s,t}$	Measures the economic value of the firm's brand in billions of U.S. dollars. It is used to define the structural-model sample and, in alternative endogenous specifications, to allow the structural value difference to receive greater weight for more valuable consumer brands.	[4,8]
III.2 Quarterly Return and Consumption-Price Inputs		
$r_{s,t} = \log(R_{s,t})$	Measures the quarterly log return of focal stock s. It is constructed from the stock's Datastream return and enters the realized-return calibration directly or as the underlying series for the estimated return moments.	[3,8]
$r_{M,t} = \log(R_{M,t})$	Measures the quarterly financial payoff of market asset M. It is constructed from the S&P 500 log return and enters the realized-return calibration directly or as the underlying series for the estimated market-return moments.	[3,8]

$\pi_{s,t}$	Measures quarterly growth in the price of the consumption good associated with stock s. It is proxied alternatively by the quarterly change in the firm's operating margin or by quarterly revenue growth, allowing the calibration to use profitability- or sales-based consumption-price exposure.	[3,8]
$\pi_{M,t}$	Measures market-wide consumption-price growth for the numéraire good. It is constructed from the quarterly log change in the Consumer Price Index and enters the realized-return calibration as the inflation component of the common consumption deflator.	[3,8]
$\hat{\mu}_{R_s,t}$	Measures the expected quarterly log return of stock s. It is estimated from the one-quarter-lagged return series by assigning equal weight to the firm's full available history and its trailing eight-quarter rolling mean.	[3,8]
$\hat{\mu}_{R_M,t}$	Measures the expected quarterly log return of market asset M. It is estimated from the one-quarter-lagged S&P 500 return series by assigning equal weight to the available historical mean and the trailing eight-quarter rolling mean.	[3,8]
$\hat{\mu}_{\pi_s,t}$	Measures expected quarterly growth in the firm-specific consumption-price proxy. It applies the same equal-weight historical and trailing eight-quarter estimator to lagged operating-margin growth or, in the alternative specification, lagged revenue growth.	[3,8]
$\hat{\sigma}^2_{R_s,t}$	Measures the standalone quarterly return variance of stock s. It is obtained from the corresponding lagged first and second return moments, with a positive floor and equal-weight shrinkage toward the historical-volatility reference used in the calibration.	[3,8]
$\hat{\sigma}^2_{R_M,t}$	Measures the standalone financial risk of market asset M. It is obtained from the lagged S&P 500 first and second moments over the historical and trailing eight-quarter windows, with a positive floor and shrinkage toward the calibration reference variance.	[3,8]
$\hat{\sigma}^2_{\pi_s,t}$	Measures uncertainty in the firm-specific consumption-price proxy. It is obtained from the lagged first and second moments of operating-margin growth or revenue growth, with a positive floor and shrinkage toward the calibration reference variance.	[3,8]
$\hat{\sigma}_{R_s,R_M}$	Measures the quarterly return covariance between stock s and market asset M. It is estimated from the focal stock-return and S&P 500 return series and supplies the covariance required to characterize their relative financial risk.	[3,8]
$\hat{\sigma}_{R_s,\pi_s}$	Measures the ability of stock s to hedge changes in the price of its associated consumption good. It is estimated as the covariance between quarterly stock returns and the selected operating-margin-growth or revenue-growth proxy.	[3,8]
$\hat{\sigma}_{R_M,\pi_s}$	Measures the ability of the S&P 500 market asset to hedge the same firm-specific consumption-price shock. It is estimated as the covariance between quarterly market returns and the selected operating-margin-growth or revenue-growth proxy.	[3,8]
$\hat{\rho}_{R_s,\pi_s}$	Measures the strength and direction of the consumption-hedge relation between stock performance and the firm's consumption-price proxy. It is constructed by standardizing the stock–price covariance with the corresponding quarterly return and price-growth variances, using operating-margin growth or revenue growth as the alternative price proxy.	[3,8]

Notes. *[1] Own TikTok data collection and manual recommendation classification. [2] LSEG Workspace, including professional analyst recommendations, option-implied measures, ESG data, and financial-event information. [3] LSEG Datastream, including stock returns, S&P 500 returns, CPI inflation, accounting variables, market characteristics, and market-wide indices. [4] Interbrand Best Global Brands. [5] Google Trends. [6] IPSOS Primary Consumer Sentiment Index via LSEG. [7] Fama–French factor data; the final data-provider citation should be added to the manuscript bibliography. [8] Own construction, transformation, or calibration input. [9] gpt-oss-120b classification based on standardized LSEG company descriptions.*

C Calibration of Structural Model

The structural model maps latent utility differentials between assets S and M onto observed quarterly portfolio-allocation counts y_j. Here, j indexes a quarterly observation, and y_j is the number of allocations to asset S among K_j observed allocation decisions in that quarter. This mapping uses a zero-inflated binomial likelihood with a logit link: in the binomial component, $y_j \sim \text{Binomial}(K_j, p_j)$, where $p_j = \text{logit}^{-1}(\eta_j)$ and $\eta_j = \lambda[\mathcal{J}_j(1;\alpha,b) - \mathcal{J}_j(0;\alpha,b)]$. The likelihood allows for an additional probability of zero counts. Joint lognormality of nominal returns and the price deflator $P_t = p_{S,t}^{\alpha}$ allows closed-form evaluation of $\mathcal{J}_j(1;\alpha,b)$ and $\mathcal{J}_j(0;\alpha,b)$. These closed forms ensure numerical stability and automatic differentiation during inference.

We estimate two specifications using quarterly changes in different proxies for $\Pi_{S,t+1}$: (i) revenue changes (ΔRevenue) and (ii) operating-margin changes (ΔOperating Margin). Hamiltonian Monte Carlo with the No-U-Turn Sampler in Stan performs the estimation. Weakly informative priors center on economically grounded benchmarks: $\alpha \sim \mathcal{B}(2,2)$; $\gamma \sim \mathcal{LN}(\log 3, 0.5)$; $\kappa \sim \mathcal{LN}(\log 2.25, 0.35)$; and $\chi \sim \mathcal{LN}(0, 0.25)$. The sensitivity and weighting parameters are $\lambda \sim \mathcal{LN}(0, 0.5)$ and $b \sim \mathcal{LN}(-1, 0.5)$. Four independent Markov chains approximate the posterior, each with 2,000 iterations (1,000 warmup) for a total of 4,000 post-warmup draws. Global convergence requires a potential scale reduction factor below 1.01 and an effective sample size (ESS) exceeding 1,500 for all parameters. Model comparison uses the Leave-One-Out Information Criterion (LOOIC) on the full sample of $N = 8{,}154$ observations.

D Data Preparation

D.1 Options-Implied Measures

We proxy forward-looking downside risk by the closing midpoint implied volatility of each stock's 20-delta put, retrieved from LSEG Workspace. Fixing delta compares options at a common risk-neutral moneyness rather than at a common strike or raw price move, reducing mechanical cross-stock differences driven by share-price levels and overall volatility (Mixon, 2011). Higher 20-delta-put implied volatility therefore indicates a higher market-implied cost of insuring against adverse price outcomes—that is, more expensive downside protection, rather than a literal tail-loss probability (Bakshi et al., 2003; Dennis & Mayhew, 2002).

D.2 LLM-based ConsumerLink Measure

Our alternative *ConsumerLink* measure is an LLM-based binary indicator of a firm's direct exposure to household consumption. Using standardized company descriptions, the open source LLM gptoss-120b identifies the firm's dominant business activity and assesses whether it primarily sells finished goods or services to end consumers. The indicator equals one for consumer-facing, downstream businesses and zero for firms primarily operating upstream, in business-to-business or industrial markets, in finance or holding structures, or where the description is insufficient to classify the activity reliably.

E Theoretical Model Details

This Appendix complements Section 2 by (i) formalizing the distributional assumptions, (ii) establishing the homogeneity and separation properties that underpin the closed-form solutions reported in the main text, and (iii) deriving the consumption-hedging comparative static under a log-linear approximation. Every equation below is either a derivation step, a formal condition, or an explicit closed form not present in the main text; we refer back to main-text equations rather than restating them.

E.1 Assumptions

Assumption A1 (Shock structure).

Under the market numéraire normalization $p_{M,t} \equiv 1$ and $\Pi_{M,t+1} \equiv 1$, the shock vector reduces to

$$\left(\ln\Pi_{S,t+1}, \ln R_{S,t+1}, \ln R_{M,t+1}\right) \sim N(\mu, \Sigma),$$

i.i.d. over time, with mean $\mu \in \mathbb{R}^3$ and positive-definite covariance matrix $\Sigma \in \mathbb{R}^{3\times3}$. The normalization removes the redundant nominal dimension while preserving the stochastic nature of the market numéraire asset's stock return.

Let

$$r_S = \ln R_S, \qquad r_M = \ln R_M, \qquad \pi = \ln\Pi_S.$$

Denote means by μ_S, μ_M, μ_π, variances by $\sigma_S^2, \sigma_M^2, \sigma_\pi^2$, and covariances by $\sigma_{SM}, \sigma_{S\pi}, \sigma_{M\pi}$.

Parameter restrictions.

The parameters satisfy

$$\beta \in (0,1), \qquad \rho > 0, \qquad \gamma > 0, \qquad \kappa \geq 1, \qquad \alpha \in [0,1], \qquad \chi > 0, \qquad b \geq 0.$$

We assume $\mathcal{J}(\omega) > 0$ and $\beta\mathcal{J}(\omega^*)^{1-\rho} < 1$, which ensure a positive and finite stationary recursive value.

Because real consumption, future real wealth, and the reference point are all linear in W_t^{real}, and because both the certainty equivalent and the recursive aggregator are homogeneous of degree one, any stationary optimum admits

$$U_t = \Psi^* W_t^{real}.$$

Following the homogeneity-based scaling used in recursive narrow-framing models, the time-t gain–loss weight can be written as

$$b_t = b \frac{\mathcal{M}_t(U_{t+1})}{\mathcal{M}_t\left(W_{t+1}^{real}\right)}.$$

Under the stationary representation $U_t = \Psi^* W_t^{real}$, this reduces to

$$b_t = b\Psi^*.$$

This scaling preserves homogeneity of the recursive problem and leads to the stationary gain–loss weight b used in the main text. Substituting the stationary representation into the recursive problem gives

$$\Psi^* = \max_{c\in(0,1),\, \omega\in[0,1]} \left[(1-\beta)c^{1-\rho} + \beta(1-c)^{1-\rho}\left(\Psi^* \mathcal{J}(\omega)\right)^{1-\rho}\right]^{\frac{1}{1-\rho}}.$$

The aggregator is strictly increasing in $\Psi^* \mathcal{J}(\omega)$. Hence portfolio choice separates:

$$\omega^* \in \arg\max_{\omega\in[0,1]} \mathcal{J}(\omega).$$

The consumption-saving margin is governed by ρ. For $\rho \neq 1$, $c^* = 1 - \beta^{1/\rho} \mathcal{J}(\omega^*)^{(1-\rho)/\rho}$, while for $\rho = 1$, $c^* = 1 - \beta$. See Guo & Dong He (2017) for more details.

E.2 Log-Linear Channel Representation

The exact real payoff is

$$R^{real}(\omega) = \frac{\omega e^{r_S} + (1-\omega)e^{r_M}}{e^{\alpha\pi}}.$$

Since the log of a weighted sum of lognormal returns is not normally distributed, we use the log-linear approximation

$$\ln R^{real}(\omega) \approx X(\omega) = \omega r_S + (1-\omega) r_M - \alpha\pi.$$

Under A1, $X(\omega) \sim N\left(\mu(\omega), s^2(\omega)\right)$, where

$$\mu(\omega) = \omega\mu_S + (1-\omega)\mu_M - \alpha\mu_\pi, \qquad (9)$$

and

$$\begin{aligned} s^2(\omega) \quad &= \omega^2\sigma_S^2 + (1-\omega)^2\sigma_M^2 + 2\omega(1-\omega)\sigma_{SM} \\ &\quad +\alpha^2\sigma_\pi^2 - 2\alpha[\omega\sigma_{S\pi} + (1-\omega)\sigma_{M\pi}]. \end{aligned} \qquad (10)$$

The first line of (10) is the financial channel. The second line is the consumption-hedge.

For the log-linearized payoff $e^{X(\omega)}$,

$$\mathcal{M}_\gamma(\omega) = \exp\left[\mu(\omega) + \frac{1-\gamma}{2}s^2(\omega)\right], \qquad \gamma \neq 1.$$

Define

$$\Delta\mu := \mu_S - \mu_M, \qquad A := \sigma_S^2 + \sigma_M^2 - 2\sigma_{SM},$$
$$B := \sigma_{SM} - \sigma_M^2, \qquad H_\pi := \sigma_{S\pi} - \sigma_{M\pi}.$$

Then

$$\frac{1}{2}s^{2\prime}(\omega) = B + \omega A - \alpha H_\pi, \qquad (11)$$

or, equivalently,

$$\frac{1}{2}s^{2\prime}(\omega) = \sigma_{SM} - \sigma_M^2 + \omega(\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM}) - \alpha(\sigma_{S\pi} - \sigma_{M\pi}).$$

E.3 Gain–Loss Valuation Layer

For $X(\omega) \sim N\big(\mu(\omega), s^2(\omega)\big)$, define

$$d_\omega^0 = \frac{\ln\chi - \mu(\omega)}{s(\omega)}, \qquad d_\omega = \frac{\ln\chi - \mu(\omega) - s^2(\omega)}{s(\omega)}.$$

The closed-form gain–loss term is

$$\begin{aligned} G(\omega) \quad &:= E\big[v\big(e^{X(\omega)} - \chi\big)\big] \\ &= e^{\mu(\omega)+\frac{1}{2}s^2(\omega)}[1 + (\kappa-1)\Phi(d_\omega)] - \chi[1 + (\kappa-1)\Phi(d_\omega^0)]. \end{aligned}$$

Define

$$\Lambda_\omega = 1 + (\kappa-1)\Phi(d_\omega), \qquad \Theta_\omega = \Lambda_\omega - (\kappa-1)\frac{\phi(d_\omega)}{s(\omega)}.$$

Then

$$G'(\omega) = e^{\mu(\omega)+\frac{1}{2}s^2(\omega)}[\Delta\mu\Lambda_\omega + (B + \omega A - \alpha H_\pi)\Theta_\omega]. \qquad (12)$$

In primitive covariances,

$$G'(\omega) = e^{\mu(\omega)+\frac{1}{2}s^2(\omega)}\big[(\mu_S - \mu_M)\Lambda_\omega + \big(\sigma_{SM} - \sigma_M^2 + \omega A - \alpha(\sigma_{S\pi} - \sigma_{M\pi})\big)\Theta_\omega\big].$$

Combining Equation (12) with the derivative of the certainty equivalent and dividing by the positive certainty-equivalent term gives

$$\begin{aligned} 0 \quad &= \Delta\mu + (1-\gamma)[B + \omega A - \alpha H_\pi] \\ &\quad + b e^{\frac{\gamma}{2}s^2(\omega)}[\Delta\mu \Lambda_\omega + (B + \omega A - \alpha H_\pi)\Theta_\omega]. \end{aligned}$$

Equivalently,

$$\begin{aligned} 0 \quad &= (\mu_S - \mu_M) + (1-\gamma)[\sigma_{SM} - \sigma_M^2 + \omega A - \alpha(\sigma_{S\pi} - \sigma_{M\pi})] \\ &\quad + b e^{\frac{\gamma}{2}s^2(\omega)}\big[(\mu_S - \mu_M)\Lambda_\omega + \big(\sigma_{SM} - \sigma_M^2 + \omega A - \alpha(\sigma_{S\pi} - \sigma_{M\pi})\big)\Theta_\omega\big]. \end{aligned} \tag{13}$$

Regrouping the terms into financial and consumption-hedge components yields Equation (13).

When $b = 0$, the first-order condition reduces to

$$0 = \Delta\mu + (1-\gamma)[B + \omega A - \alpha H_\pi], \tag{14}$$

or, in primitive covariances,

$$0 = (\mu_S - \mu_M) + (1-\gamma)[\sigma_{SM} - \sigma_M^2 + \omega A - \alpha(\sigma_{S\pi} - \sigma_{M\pi})].$$

This gives Equation (7). Thus, the benchmark allocation separates cleanly into a financial component and a consumption-hedge component.

E.4 Special Cases

Table E.1 summarizes how the channels change across limiting cases of the log-linearized model. The table also clarifies the role of $b > 0$. Gain–loss valuation is not necessary for a consumption-hedge channel when $\gamma > 1$, because consumption-hedge covariances already affect real-return risk. It is retained in the full model because it captures reference-dependent valuation of real consumption capacity and preserves a role for consumption exposure even when the certainty-equivalent variance channel is neutral.

Table E.1: ***Channel activity across limiting cases of the log-linearized model***

Parameter regime	Solution / characterization of ω^*	Financial channel	consumption-hedge channel	Main implication
$\alpha > 0$: consumption-hedge exposure is potentially relevant				
$\gamma > 1,\ b > 0$ (full model)	No general closed-form solution. The interior optimum solves the full gain–loss FOC.	✓	✓	Gain–loss valuation is active; financial and consumption-hedge exposures jointly affect portfolio choice.
$\gamma > 1,\ b = 0$	$\omega_0^* = \frac{\mu_S - \mu_M}{(\gamma - 1)(\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM})} - \frac{\sigma_{SM} - \sigma_M^2}{\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM}} + \alpha \frac{\sigma_{S\pi} - \sigma_{M\pi}}{\sigma_S^2 + \sigma_M^2 - 2\sigma_{SM}}$	✓	✓	Closed-form benchmark without gain–loss valuation.
$\rho = \gamma = 1,\ b > 0$	The CE variance channel is neutral; ω^* is determined through gain–loss valuation of the real payoff.	✓	✓	Consumption-hedge exposure remains relevant because gain–loss valuation evaluates real consumption capacity relative to χ.
$\rho = \gamma = 1,\ b = 0$	$\omega^* = 1$ if $\mu_S > \mu_M$; $\omega^* = 0$ if $\mu_S < \mu_M$; any $\omega \in [0,1]$ if $\mu_S = \mu_M$.	✓	χ	With neither CE-variance effects nor gain–loss valuation, the consumption-hedge allocation channel collapses.
$\alpha = 0$: no consumption-hedge exposure				
$\gamma > 1,\ b > 0$	No general closed-form solution; solve the full FOC after setting $\alpha = 0$.	✓	χ	Gain–loss valuation remains active, but all consumption-hedge covariances drop out.
$\gamma > 1,\ b = 0$ (financial benchmark)	$\omega_{\text{it}}^* = \frac{\mu_S - \mu_M}{(\gamma-1)(\sigma_S^2+\sigma_M^2-2\sigma_{SM})} - \frac{\sigma_{SM}-\sigma_M^2}{\sigma_S^2+\sigma_M^2-2\sigma_{SM}}$	✓	χ	Financial channel only.

The consumption-hedge part of $s^2(\omega)$ is

$$\alpha^2\sigma_\pi^2 - 2\alpha[\omega\sigma_{S\pi} + (1-\omega)\sigma_{M\pi}].$$

Therefore, for $\omega > 0$ and $\alpha > 0$,

$$\frac{\partial s^2(\omega)}{\partial \sigma_{S\pi}} = -2\alpha\omega < 0.$$

If

$$\sigma_{S\pi} = \mathrm{Corr}(r_S, \pi)\sigma_S\sigma_\pi,$$

then

$$\frac{\partial s^2(\omega)}{\partial \mathrm{Corr}(r_S, \pi)} = -2\alpha\omega\sigma_S\sigma_\pi < 0.$$

Thus, a stronger positive covariance between asset S's return and consumption-hedge inflation lowers log real-return risk for portfolios that hold S. For $\gamma > 1$, the certainty-equivalent component values this risk reduction positively.

With $b > 0$, the same hedge is additionally evaluated relative to the reference threshold χ. If the consumption hedge reduces the probability or severity of falling below the reference level, gain–loss valuation amplifies demand for the consumption-linked stock. This amplification is typical in calibrations where χ is above one and downside real-consumption shortfalls are salient. It is not imposed mechanically, however, because the sign and magnitude of the gain–loss response depend on the location of χ, the loss-aversion parameter κ, and the probability mass near the threshold.

E.6 Exact Model Without Log-Linearization

Without log-linearization,

$$R^{real}(\omega) = \Pi_S^{-\alpha}[\omega R_S + (1-\omega)R_M].$$

The marginal real payoff from increasing the weight on S is

$$\frac{\partial R^{real}(\omega)}{\partial \omega} = \Pi_S^{-\alpha}(R_S - R_M).$$

For $\gamma \neq 1$, the exact interior first-order condition is

$$\begin{aligned} 0 \quad &= \mathcal{M}_\gamma\left(R^{real}(\omega)\right)^\gamma E\left[\Pi_S^{-\alpha}(R_S - R_M)\left(R^{real}(\omega)\right)^{-\gamma}\right] \\ &\quad + bE\left[\Pi_S^{-\alpha}(R_S - R_M)\left(\mathbf{1}_{\{R^{real}(\omega)\geq\chi\}} + \kappa\mathbf{1}_{\{R^{real}(\omega)<\chi\}}\right)\right]. \end{aligned}$$

This exact condition has the same economic structure as the main text: the marginal financial payoff $R_S - R_M$ is converted into real consumption units by the stochastic deflator $\Pi_S^{-\alpha}$. However, $\omega R_S + (1-\omega)R_M$ is a weighted sum of lognormal variables, so the exact expectations generally do not reduce to a closed-form expression in primitive covariances. The log-linear representation is therefore used in the main text to display the channels parsimoniously.

For the exact special case $\gamma = 1$ and $b = 0$,

$$\mathcal{M}_1\left(R^{real}(\omega)\right) = \exp(E[\ln(\omega R_S + (1-\omega)R_M)] - \alpha E[\pi]).$$

The term $-\alpha E[\pi]$ is independent of ω, so the consumption-hedge allocation channel collapses. The interior condition becomes

$$E\left[\frac{R_S - R_M}{\omega R_S + (1-\omega)R_M}\right] = 0.$$

Thus, in both the exact and log-linearized models, when $\gamma = 1$ and $b = 0$, consumption prices affect real consumption capacity but cancel from the marginal portfolio-allocation condition. The consumption-hedge channel is therefore a level channel, whereas the portfolio-allocation channel is driven solely by the relative financial payoff $R_S - R_M$.