Sentiment and Consumption in the Digital Age*

Yi Huang† Sichuang Xu‡ Yaya Yu§ Jing Zhou¶ Yujun Zhou‖

This version: April 2026

Abstract

Using high-frequency transaction-level data from a leading Chinese e-commerce platform, we construct text-based sentiment indices and match them to individual online and offline consumption records. We document three stylized facts. First, an *allocation effect*: sentiment reallocates spending between the online and offline channels rather than altering aggregate expenditure. Second, a *dynamic effect*: a rise in sentiment immediately raises online consumption, offset by a nearly equivalent drop in offline spending, before reversing the following week. Third, *heterogeneity*: online consumption is more sensitive to sentiment for less durable goods, among affluent consumers, and in developed regions. To reconcile these facts, we build a two-channel consumption-saving model in which a non-fundamental sentiment shock enhances the comparative advantage of online shopping, through greater attractiveness, convenience, and lower search frictions. Overall, our work underscores a reallocation effect of sentiment and shows that the partial-equilibrium force is largely muted, while the general-equilibrium force, which typically operates through price changes, is key for sentiment to drive aggregate consumption fluctuations.

JEL Codes: D12 E21 E71 L81

*We thank comments from Pengfei Wang, Jianfeng Yu, Danyang Xie, Jiayin Hu, the participants of Workshop on AI in Finance and Digital Economy, HKUST(GZ), the 2nd Workshop of the East and South-East Asian Macroeconomic Society Meeting at Kuala Lumpur and seminars at PBC School of Finance, Tsinghua University. Jing Zhou thanks for the financial support of National Natural Science Foundation of China (No. 72303039). The authors acknowledge and appreciate the supports from the Digital Economy Open Research Platform (www.deor.org.cn). All data was sampled and desensitized, and was remotely analyzed on the Ant Open Research Laboratory in an Ant Group Environment which is only remotely accessible for empirical analysis. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks.

†Bank for International Settlements and CEPR. Email: Yi.Huang@bis.org.

‡The Chinese University of Hong Kong, Shenzhen. Email: xusichuang@cuhk.edu.cn.

§Ant Group. Email: yuyaya.yyy@antgroup.com.

¶School of Economics, Fudan University; Email: jingzhou@fudan.edu.cn.

‖Department of Finance, Tsinghua University; Email: zhouyujun23@mails.tsinghua.edu.cn.

1 Introduction

The rapid growth of e-commerce has led to unprecedented levels of online consumer spending. A substantial body of existing research documents that positive consumer sentiment significantly boosts aggregate consumption (e.g., Ludvigson, 2004; Benhabib and Spiegel, 2018, and others). In the digital age, however, a new open question remains: does heightened sentiment boost online and offline consumption simultaneously? In other words, when consumers have access to both online and offline shopping channels, does heightened sentiment drive household expenditure toward one channel over the other, or across both?

Leveraging large-scale, transaction-level data that encompassess both online and offline expenditures, this paper attempts to address the above question and to systematically investigate the relationship between consumer sentiment and consumption behavior. Our analysis utilizes weekly proprietary data from Ant Financial Services Group and Alibaba Group, a preeminent Chinese e-commerce ecosystem. The dataset comprises approximately 7 million weekly online transactions for roughly 100,000 Chinese consumers across 322 cities throughout 2019.

This granular dataset offers three distinctive advantages. First, *high frequency*: the weekly frequency allows us to capture dynamic responses to consumer sentiment shocks, which are often transitory and short-lived; such high-frequency dynamics are typically obscured in extant literature that relies on lower-frequency data. Second, *identification strategy*: the granularity of the data enables the inclusion of a comprehensive set of individual-by-month fixed effects, effectively controlling for slow-moving fundamental changes at the individual, city, or macroeconomic levels within a narrow window, ensuring that our results isolate responses to high-frequency sentiment fluctuations rather than fundamental shifts. Third, *informational richness*: the dataset records both online and offline consumption, allowing for a direct comparison of the impact of sentiment across channels; furthermore, the inclusion of detailed transaction and demographic characteristics facilitates heterogeneity analysis across multiple dimensions.

To measure consumer sentiment, we construct a city-week-level sentiment index based on Weibo posts, employing a Large Language Model (LLM) via OpenAI's GPT-4o-mini API. This approach represents a frontier methodology in the literature. In contrast to the traditionally keyword-based lexicon methods, which struggle to capture subtle semantic nuances in text, our LLM-based measure goes beyond mere binary (positive/negative) classifications. It enables us to distinguish whether an emotional state is driven by substantive fundamental changes or by purely exogenous "mood."

Crucially, the consumer sentiment examined in this paper refers specifically to non-fundamental sentiment—that is, sentiment *not* directly driven by factors related to personal circumstances, product attributes, platform-specific factors, or broader economic conditions. By isolating this dimension, our empirical analysis provides novel micro-level empirical evidence for theoretical works exploring the real economic impacts of non-fundamental sentiment.

Our first finding reveals that heightened sentiment significantly increases online consumption,

yet it does not necessarily increase aggregate consumption. Specifically, in response to a positive sentiment shock, household's aggregate expenditure does not rise significantly, whereas online consumption rises substantially – accompanied by a nearly equivalent and significant decline in offline consumption. This pattern holds consistently across transaction value, frequency, and variety, suggesting that heightened consumer sentiment leads to a crowding-out effect of online consumption on offline consumption.

Furthermore, this crowding-out effect is economically substantial: according to our baseline estimation, a one-standard-deviation positive shock to sentiment increases online payment amounts by 0.38% and reduces offline payment amounts by 0.35%. This represents an expenditure reallocation of approximately \$114–\$203 billion from the offline to the online channel within a single week – an amount equivalent to roughly 0.16% to 0.26% of China's 2019 Gross Domestic Product.

Our second set of findings reveals a pronounced intertemporal volatility in the response of both online and offline consumption. The reallocation effect described above exhibits a rapid reversal in the subsequent week: online consumption decreases significantly, alongside a corresponding expansion in offline consumption. This reversal is consistent across transaction value, frequency, and variety. Consequently, fluctuations in sentiment amplify the volatility of online and offline consumption. Furthermore, we identify a distinct asymmetry in these dynamics: the dynamic pattern is more pronounced during periods of low sentiment than periods of high sentiment.

Our third set of findings concerns the heterogeneous responses of online and offline consumption to sentiment shocks, which vary along the dimensions of product, consumer, and regional characteristics. The key results are as follows: In terms of product attributes, the dynamic pattern documented above is significantly more pronounced for non-durable goods than for durable goods. With respect to consumer characteristics, this pattern is more salient among wealthier consumers. Finally, regarding regional attributes, the effect is more prominent in regions with fewer financial constraints and greater access to offline shopping facilities.

To rationalize these stylized facts, we develop a two-channel consumption-saving model incorporating consumer sentiment. In this framework, agents allocate consumption across online and offline channels. Their purchasing decisions depend not only on fundamental factors – such as prices and idiosyncratic matching values – but also on search costs and a "pure" sentiment component drived from the shopping experience itself.

Ceteris paribus, the online channel is characterized by lower search costs and enhanced shopping utility, owing to the convenience and appealing design of digital platforms. The model features stochastic meeting opportunities and random matching, which define a reservation threshold: consumers execute a purchase only when the perceived value exceeds this threshold and otherwise defer to the next opportunity.

Within this framework, an increase in consumer sentiment raises the utility derived from the shopping experience in both channels. However, the lower search cost and superior shopping utility of the online channel amplify its comparative advantage over offline alternatives, thereby strengthen-

ing the positive effect of sentiment on online consumption relative to offline consumption. Furthermore, the model can account for the observed dynamic reversal and heterogeneity in the empirical results.

Literature Review Existing empirical studies on sentiment and consumption have primarily focused on the impact of sentiment on aggregate consumption. For instance, Carroll, Fuhrer, and Wilcox (1994) report that consumer sentiment forecasts future changes in aggregate household consumption expenditures. Similarly, Ludvigson (2004) demonstrate that consumer confidence has modest independent predictive power for future growth in aggregate consumer spending. Subsequent research has sought to establish a causality link by exploiting exogenous events that trigger shifts in consumer sentiment. Notably, Benhabib and Spiegel (2018) find that agents in states aligned with the sitting President's political party, and that this optimism translates into higher future consumption growth. Similarly, Gillitzer and Prasad (2018) utilize individual-level survey data and postcode-level automobile purchases in Australia to demonstrate that partisan bias-rather than shifts in economic fundamentals-causally affects real consumption. Further evidence of non-fundamental sentiment shocks is provided by Brodeur (2018), who compares successful and failed terror attacks to show that increased media coverage and heightened consumer pessimism, rather than physical damage, drive reductions in economic activity. More recently, Lagerborg, Pappa, and Ravn (2022) use fatalities from mass shootings as instrumental variable to identify autonomous changes in consumer confidence unrelated to economic fundamentals, revealing that negative non-fundamental sentiment shocks induce substantial and persistent declines in real consumption.

A defining characteristic of these seminal studies is their reliance on low-frequency aggregate data (e.g., quarterly or monthly), which inherently limits their capacity to capture the rapid, dynamic transmission of sentiment shocks. Our paper departs from this literature in two significant dimensions. First, we leverage high-frequency consumption data to analyze the immediate impact of sentiment shocks. Second, we introduce and emphasize a channel-reallocation effect, examining the heterogeneous responses across online and offline consumption sectors.

Our paper contributes to the burgeoning literature on the digital economy. Prior research has established that e-commerce reshapes market outcomes through two primary channels: price competition and product variety. Specifically, Cavallo (2017) and Dolfen et al. (2023) document that e-commerce intensifies price transparency and competition, while Dolfen et al. (2023) and Couture et al. (2021) find that digital platforms enhances consumption variety – effects that collectively translate into welfare gains for consumers.

Within this literature, price incentives and geographical friction have received consideration in the literature. Regarding price, Goolsbee (2000) finds that, after controlling for observable characteristics, individuals in high-tax areas are significantly more likely to purchase online. A related study by Houde, Newberry, and Seim (2023) analyzes how "nexus" tax laws distort Amazon's investments in its distribution network. Regarding the role of geography, Forman, Ghose, and

Goldfarb (2009) show that consumers in areas with limited offline retail options derive greater benefits from online access. More recently, Bounie, Camara, and Galbraith (2025) employ billions of geolocated, high-frequency credit card transactions from France to reestimate gravity models, assessing the distance effect in online versus offline shopping.

More recently, using U.S. micro-level household data from multiple sources, Malmendier and Shen (2024) document a "scarring" effect: consumers who have experienced higher lifetime unemployment rates display persistent pessimism and significantly lower consumption levels, even after controlling for current income, wealth, and demographic characteristics. Our paper also builds on this perspective and focuses on the influence of consumer sentiment on online and offline comsumption, a dimension that remains under-explored in the exisitng literature.

Our paper also contributes to the broader literature on the drivers of of household spending. This body of research has often focused on fundamental shocks, such as changes in household income or wealth. For instance, Jappelli and Pistaferri (2010) provide a comprehensive review of the theoretical foundations and empirical approaches for estimating consumption responses to income changes. Similarly, Paiella and Pistaferri (2017) decompose wealth effects by separating anticipated from unanticipated changes. More recently, this literature has increasingly leveraged high-frequency, granular consumption data. Gelman et al. (2023), for example, use transaction-level data to analyze how aggregate consumption responds to changes in gasoline prices. Along similar lines, Ganong and Noel (2019) , Ganong et al. (2024) and Ganong et al. (2025) employ de-identified high-frequency bank account data to examine expenditure responses to various income shocks. Furthermore, Garmaise, Levi, and Lustig (2024) utilize micro-level credit card and bank transactions to demonstrate that U.S. consumer reduce their spending in response to negative macroeconomic news – even when such news is false – behaving as if they were financially constrained despite unchanged fundamentals. Like these recent contributions, our study utilizes high-frequency consumption data. However, it aligns most closely with the latter work by focusing on consumption response to sentiment shocks that are orthogonal to the fundamental factors traditionally investigated in the literature.

Our paper also relates to the theories where sentiment drives real economic fluctuations despite being orthognal to fundamentals. This strand of literature posits that, in the presence of various frictions, sentiment– despite being orthogonal to fundamental factors–can drive fluctuations in real business cycles. Prominent theoretical contributions include Angeletos and La'O (2010), Barsky and Sims (2012), Angeletos and La'O (2013), Benhabib, Wang, and Wen (2015), Benhabib, Liu, and Wang (2016) and Huo and Takayama (2024). Furthermore, Angeletos, Collard, and Dellas (2018) provide both theoretical and empirical analyses on the role of confidence shocks in business cycles. More recently, Maxted (2023) develops a macro-finance model incorporating sentiment shocks and diagnostic expectations, demonstrating their significant impact on asset prices and real economic activity. While these theoretical works have advanced our understanding of the mechanisms underlying non-fundamental sentiment, they have largely lacked detailed micro-level empirical validation – a gap the present paper seeks to bridge.

Finally, our work relates to the evolving empirical literature on measuring sentiment. Early studies employed local sunshine as a proxy for mood or sentiment. For instance, Cortés, Duchin, and Sosyura (2016) demonstrate that fluctuations in local sunshine systemically influence the lending decisions of loan officers. Subsequent work shifted toward text-based approaches. Shapiro, Sudhof, and Wilson (2022), for example, apply lexicon-based methods to measure news sentiment through textual analysis, showing how media-driven sentiment affects economic expectations. Thanks to advances in natural language processing, a burgeoning literature has begun to leverage large language models (LLMs) to construct more sophisticated sentiment measures, especially in financial economics. Examples include Jha, Liu, and Manela (2025), who employ LLMs to extract financial sentiment from a vast corpus of published books, and Kirtac and Germano (2025), who illustrate that LLMs possess a superior capacity for understanding contextual information, which improves the precision of sentiment extraction in financial texts. Similarly, Chang et al. (2023) utilize ChatGPT to derive sentiment from earnings call transcripts and assess its effect on investor trading behavior, while Gao, Xiong, and Yuan (2026) use the DeepSeek reasoning model to infer institutional equity market beliefs from Chinese analyst reports. Like these studies, our paper employ an LLM-based approach to construct a sentiment measure. However, unlike these studies focusing on financial or investor sentiment, our analysis center on consumer sentiment and its implications for consumption behavior.

The remainder of this paper is structured as follows. Section 2 presents the stylized facts regarding the relationship between consumer sentiment and both online and offline consumption, supported by extensive empirical analysis. Section 3 develops a theoretical framework incorporating online and offline consumption channels to account for these observed regularities. Finally, Section 4 concludes.

2 Stylized Facts

2.1 Description of the Consumption Dataset

Our analysis uses proprietary data from Alipay, a leading digital payment platform in China, founded in 2004. Alipay enables seamless wallet-free payments at physical merchants, facilitating offline transactions, while also integrating digital payments into routine online activities, such as e-commerce and food delivery. Consequently, Alipay captures a substantial and broadly representative share of consumer payment behavior across both offline and online environments. As noted by Chen, Qian, and Wen (2023), Alipay processed 1.5 billion payment transactions and approximately 812 million sales orders in 2017. Mojon et al. (2025) further document that by 2021, Alipay's e-wallet had become the dominant payment method for e-commerce transactions, commanding an 83% market share in China. Taken together, these characteristics suggest that Alipay consumption data are sufficiently extensive and representative to provide a comprehensive measure

of both online and offline consumption activities in China.

This study was remotely conducted on the Ant Open Research Laboratory[1]in an Ant Group Environment. All data was sampled and desensitized, and was analyzed on the Ant Open Research Laboratory. The laboratory is a sandbox environment where the authors can only remotely conduct empirical analysis and individual observations are invisible. The main regression variables include basic variables, investment variables, and consumption variables.

Our consumption sample, provided by Alipay, covers all 52 weeks of 2019. We select 2019 as the sample period to isolate our findings from the unprecedented disruptions to consumer behavior during the COVID-19 era (Chen, Qian, and Wen, 2021). The dataset comprises a randomly drawn sample of 100,000 active Alipay users, tracking their weekly consumption along with individual demographic characteristics—specifically, age, gender, and city of residence. For each user, the data capture both online and offline payment amounts and the number of transactions across 17 distinct product categories.

The dataset is highly granular, containing transaction-level details for both online and offline purchases and covering more than 18 million transactions. For each transaction, the dataset records payment amount, transaction timestamp, product category, and channel (online or offline). These data allow us to construct weekly measures of each consumer's online and offline payment amounts ($Amount_{i,c,f,t}$) and transaction volume ($Transactions_{i,c,f,t}$).[2] We also measure online and offline consumption variety ($Variety_{i,c,f,t}$) as the number of distinct product categories in which a consumer records at least one transaction by channel f during a given week. Product categories span a broad set of goods and services (e.g., groceries, clothing, dining, travel, entertainment, and electronics), providing comprehensive coverage of household consumption.

Individual-level demographic information—including age, gender, and residential city—is linked to each consumer, enabling rich heterogeneity analysis. The sample is representative of active Alipay users nationwide, with stratified random sampling ensuring coverage across cities of varying sizes and regions. To protect privacy, all data were sampled, desensitized, and analyzed within the Ant Open Research Laboratory, a sandbox environment in which the authors conduct empirical analysis remotely and individual-level observations are not directly visible. No personally identifiable information is accessible, and all results are reported in aggregate form.

The weekly frequency of the dataset enables the study of high-frequency consumption dynamics and responses to sentiment shocks, which are often short-lived and difficult to capture with traditional monthly or quarterly data. The inclusion of both online and offline channels allows for direct comparison of consumer behavior across digital and physical retail environments.

[1]https://www.deor.org.cn/labstore/laboratory

[2]In our notation, i denotes the consumer, c the consumer's city of residence, f the consumption channel ($f = 0$ for offline and $f = 1$ for online), and t the week.

2.2 Measurement of Sentiment

In this section, we detail the construction of our consumer sentiment measure using textual analysis. We begin by compiling a comprehensive text corpus in which individuals express their attitudes toward consumption. To spatially and conceptually align this sentiment measure with our consumption microdata, we restrict our focus to posts authored by Chinese consumers—specifically, geotagged posts from Weibo, a leading microblogging platform in China where users publicly share status updates and personal thoughts. We collect these posts over the 52 weeks of 2019, as described below.

A critical step in this process is defining posts that explicitly convey attitudes toward consumption, ensuring that our measure captures consumption-related sentiment rather than unrelated topics. To address this challenge, we implement a two-step topic-modeling procedure. We briefly summarize the two steps below, with further details provided in Appendix A.1.

First, we employ unsupervised topic modeling to identify a comprehensive set of consumption-related keywords. We begin by assembling a seed corpus of 100,000 Weibo posts published during March 2019 that contain at least one of six initial consumption keywords. We then apply both Latent Dirichlet Allocation (LDA) and a BERT-based topic model to this corpus, extracting the top 50 keywords generated by each topic model. Following a manual review of this expanded set of keywords, we retain 58 keywords that strictly capture consumption activities. Leveraging this refined keyword dictionary, we retrieve all geotagged Weibo posts published throughout 2019. To ensure data integrity, we implement two filters to eliminate advertisements and algorithmic spam: (i) we compile a dictionary of common promotional phrases to screen out commercial content, and (ii) we exclude high-frequency accounts—defined as users averaging more than one post per week. This selection protocol yields a final sample of approximately a hundred thousand geographically tagged posts that explicitly convey attitudes toward consumption.

Second, we classify the sentiment of each Weibo post using a large language model (LLM). Relative to conventional sentiment-classification approaches, the LLM-based approach provides distinct methodological advantages. First, it evaluates words within their broader semantic context rather than in isolation. This addresses a fundamental limitation of traditional lexicon-based methods, which typically ignore syntactic structure and can consequently mischaracterize the underlying polarity of the text. Second, because the employed LLMs are pre-trained on extensive, open-domain corpora, they generalize robustly across a diverse array of topics. This flexibility enables the model to capture cultural nuances and idiomatic expressions, as well as to identify complex linguistic features—such as sarcasm, irony, and internal contradictions—that frequently invert a statement's apparent sentiment. Collectively, these capabilities substantially enhance classification accuracy, allowing us to move beyond coarse proxies and construct a highly precise measure of the sentiment expressed in each post.

Specifically, we use OpenAI's GPT-4o-mini API with a carefully designed zero-shot prompt to perform two tasks simultaneously: (i) label the sentiment of each post as positive or negative, and

(ii) flag posts that discuss product quality. We exclude posts discussing product quality before constructing the sentiment measure. This restriction helps ensure that our sentiment measure captures attitudes toward consumption rather than product-specific quality concerns. Appendix A.6 provides the detailed prompt design.

Finally, after classifying the sentiment of each Weibo post, we construct city–week sentiment, $Sentiment_{c,t}$, as the average sentiment score across all posts originating from city c during week t. Appendix A.3 details this construction. The resulting sentiment measure exhibits substantial spatial variation, which provides the critical identifying variation for our empirical analysis. Figure 1 maps the cross-sectional distribution of both mean sentiment and sentiment volatility at the city level. The pronounced geographic heterogeneity in both the level and dispersion of sentiment provides ample spatial variation for our estimation. Additionally, Appendix A.4 establishes the empirical validity of our sentiment measure by benchmarking it against established macroeconomic sentiment indices and well-known consumption-related measures.

Our text-based sentiment measure offers distinct advantages over traditional survey-based consumer confidence indices. Conventional survey-based measures typically aggregate ordinal responses across heterogeneous questions, obscuring the economic interpretation of both the magnitude and direction of intertemporal index changes. Furthermore, these surveys frequently rely on subjective questions regarding general business conditions; consequently, cross-sectional variation may simply capture heterogeneous interpretations of ambiguous terminology rather than genuine shifts in underlying economic expectations. This structural ambiguity introduces classical measurement error and generates spurious volatility. In contrast, our approach directly extracts real-time, unprompted sentiment from consumers' spontaneous social media expressions, effectively circumventing the measurement errors and framing biases inherent in conventional survey design.

2.3 Empirical Design

We estimate the impact of sentiment on online and offline consumption using the following specification:

$$\begin{aligned}\log(c_{i,c,f,t}) = \beta_1 \times \text{Sentiment}_{c,t-1} + \beta_2 \times \text{Sentiment}_{c,t-1} \times \text{Online}_{i,c,f,t} \\ + \gamma X_{c,t-1} + \text{FE}s + u_{i,c,f,t}\end{aligned} \tag{1}$$

where i indexes the consumer, c the city of residence, $f \in \{0, 1\}$ the consumption channel (where $f = 0$ denotes offline and $f = 1$ denotes online), and t the week. The dependent variable, $\log(c_{i,c,f,t})$, is the natural logarithm of channel-specific consumption by consumer i residing in city c in week t, measured by (i) payment amount ($\text{Amount}_{i,c,f,t}$), (ii) number of transactions ($\text{Transactions}_{i,c,f,t}$), and (iii) consumption variety ($\text{Variety}_{i,c,f,t}$). The key explanatory variable, $\text{Sentiment}_{c,t-1}$, denotes the sentiment level in city c during week $t-1$. $\text{Online}_{i,c,f,t}$ is a dummy variable equal to one if the dependent variable $\log(c_{i,c,f,t})$ corresponds to online consumption and zero otherwise. FEs include

week, consumer, consumer×online, consumer×month, and consumer×month×online fixed effects. $X_{c,t-1}$ is the vector of city–week controls, including temperature, humidity, wind speed, sunshine, and air quality index.

By stacking each consumer's online and offline consumption observations across weeks, this empirical specification allows us to directly estimate and statistically test the differential response of online versus offline consumption to changes in sentiment. Specifically, our primary coefficients of interest are β_1 and β_2. The coefficient β_1 captures the direct effect of sentiment on offline consumption, while the interaction coefficient β_2 measures the incremental effect of sentiment on online consumption relative to offline consumption.

We then estimate the impact of sentiment on overall consumption using the following specification:

$$\log(c_{i,c,t}) = \beta_1 \times \text{Sentiment}_{c,t-1} + \gamma X_{c,t-1} + \text{FE}s + u_{i,c,t} \tag{2}$$

where $\log(c_{i,c,t})$ is the natural logarithm of overall consumption by consumer i residing in city c in week t, measured by (i) the sum of online and offline payment amounts, (ii) the sum of online and offline transaction counts, and (iii) the sum of online and offline consumption variety. FEs include week, consumer, and consumer×month fixed effects.

To estimate Equation (2), we construct a consumer-week panel in which each observation captures a consumer's overall consumption and local sentiment for a given week. The primary coefficient of interest, β_1, measures the effect of sentiment on overall consumption.

2.4 Empirical Results

We estimate the effect of consumer sentiment on consumption using the empirical specifications described in Section 2.3. We organize the results into the following stylized facts.

Fact 1. *(Reallocation Effect) While consumer sentiment exhibits a neutral overall effect on consumption, elevated sentiment is associated with an increase in contemporaneous online consumption and a simultaneous reduction in offline consumption. In the subsequent period, online consumption declines while offline consumption rises.*

We first examine the effect of sentiment on overall consumption by estimating Equation (2). As reported in Panel A of Table 1, consumer sentiment has no statistically significant effect on overall payment amount, number of transactions, and consumption variety. We then investigate the differential effects of sentiment across consumption channels by estimating Equation (1). Panel B reports these estimates. Across all columns, the coefficient on $Sentiment_{c,t-1}$—which captures the direct effect of sentiment on offline consumption—is negative and statistically significant, implying that positive sentiment is associated with reduced offline expenditure. In contrast, the coefficient on the interaction term, $Sentiment_{c,t-1} \times Online_{i,c,f,t}$, is positive, statistically significant, and larger

in magnitude than the coefficient on $Sentiment_{c,t-1}$. Taken together, these results indicate that positive sentiment induces a reallocation of expenditure toward the online channel.[3]

In the classical framework of the permanent income model, the Euler equation implies that unanticipated shocks shift marginal utility and thus affect consumption growth. Guided by this theoretical insight and following Jappelli and Pistaferri (2010), we modify the empirical specification detailed in Section 2.3 by specifying the dependent variable as the first difference of the natural logarithm of consumption, and the explanatory variable as the first difference in sentiment. This specification captures the impact of a one-unit positive sentiment shock on consumption growth. Table 2 demonstrates that these estimates are consistent with our baseline results in Table 1. Taken together, the findings indicate that a positive sentiment shock significantly lowers offline consumption growth in the following week while simultaneously raising online consumption growth, implying a short-run reallocation of expenditure from the offline to the online channel.

Although our results indicate that sentiment has no significant impact on total expenditure, they highlight that sentiment induces a short-run reallocation of expenditure from the offline to the online channel, which is driven by differences in the shopping experience across channels, rather than by shifts in economic fundamentals such as product attributes, personal income, or good prices.

According to the point estimates reported in Table 1, combined with the sample standard deviations of sentiment shown in Table A2, a one-standard-deviation positive shock to sentiment increases online payment amount by 0.38% (0.0116×0.3277)[4] and decreases offline payment amount by 0.35% (0.0106×0.3277)[5], which just offsets the increase in online consumption, supporting our argument for the reallocation effect.

While the magnitude of these estimates might appear modest, their economic significance is substantial given the massive nationwide consumption base and the high-frequency nature of our week-level dataset. To explicitly illustrate this, we perform a simple back-of-the-envelope calculation extrapolating from our well-identified consumption responses to positive sentiment shocks to an aggregate national response.

As outlined above, from our baseline regression detailing the differential response of online versus offline consumption to a positive sentiment shock, we find that a one-standard-deviation positive shock increases online payment amount by 0.38% and decreases offline payment amount by 0.35%. Applying these estimates to the estimated nationwide weekly consumption base—which totals approximately \$114 billion online and \$202.9 billion offline[6]—translates to roughly \$429.4

[3] In Appendix A.8.4, we conduct an additional robustness check using the average payment amount per transaction as the dependent variable, which provides consistent evidence.

[4] In Column (1) of Panel B of Table 1, the combined coefficient of $Sentiment_{c,t-1}$ and the interaction term $Sentiment_{c,t-1} \times Online_{i,c,f,t}$ is 0.0116, representing the estimated effect of sentiment on online consumption. As shown in Table A2, the standard deviation of $Sentiment_{c,t-1}$ is 0.3277.

[5] In Column (1) of Panel B of Table 1, the coefficient of $Sentiment_{c,t-1}$ is 0.0106, representing the estimated effect of sentiment on offline consumption.

[6] Using the official RMB/USD exchange rate as of December 31, 2019, from the National Bureau of Statistics of China, the average weekly online and offline expenditure per randomly drawn Alipay user in our sample is \$81 and \$145, respectively. Extrapolating this to the total Chinese population in 2019 yields aggregate weekly consumption

$\sim$ \$710.5 million in expenditure being reallocated from the offline to the online channel within a single week, following a nationwide, one-standard-deviation positive sentiment shock. Annualizing the weekly expenditure reallocated from the offline to the online channel following a one-standard-deviation positive sentiment shock yields an aggregate shift equivalent to $0.16\% \sim 0.26\%$ of China's 2019 Gross Domestic Product. This effect is substantial, amounting to approximately $17\%(= 0.38/2.3)$ of the annual sentiment-driven boost to consumption estimated in Benhabib and Spiegel (2018).

It is worth noting that we control for a rich set of fixed effects: week-level, consumer-level, consumer-month-level, and online-consumer-month-level. The combination of high-frequency weekly microdata and this saturated fixed-effect structure allows us to partial out the direct effects of slow-moving economic fundamentals—such as household income, local price levels, and broader macroeconomic trends. Furthermore, this empirical design absorbs potential indirect channels through which sentiment might endogenously alter these structural factors. Our approach also inherently mitigates concerns of reverse causality, specifically the threat that local consumption dynamically drives city-level sentiment. Because aggregate consumption is highly persistent and adjusts slowly over the business cycle, it is unlikely to generate the high-frequency, week-to-week variation observed in our sentiment measure. Supplemental instrumental variable estimates provide further support for this identifying assumption. Therefore, our results capture high-frequency movements in consumer sentiment, and the short-run response is driven by variations in consumer sentiment rather than by other fundamental factors.

Fact 2. *(Dynamics and Asymmetry) The effects of sentiment on online and offline consumption are primarily driven by periods of low sentiment, indicating asymmetric impacts. Moreover, elevated sentiment increases the volatility of online consumption while reducing the volatility of offline consumption.*

Table 3 explores potential asymmetries in the impact of sentiment on consumption. To investigate these dynamics, we partition the full sample at the median of the sentiment measure and estimate Equation (1) separately for each subsample. Columns (1), (3), and (5) present the estimates for periods of below-median sentiment, whereas columns (2), (4), and (6) report the corresponding results for periods of above-median sentiment.

As reported in columns (1) and (2), a one-standard-deviation positive sentiment shock is associated with an approximately 0.47% (-0.0142×0.3277) decline in offline payment amount within the below-median subsample, a magnitude substantially larger than the marginal 0.07% ($-0.0022 \times$

estimates of roughly \$113 billion by online channel and \$203 billion by offline channel. It is important to note that although Alipay's extensive market penetration ensures its broader user base is highly representative of the Chinese population, our specific empirical sample comprises a random draw of active consumers, who inherently exhibit higher transaction frequencies. Consequently, this simple extrapolation serves as an illustrative upper bound for the true national aggregate.

0.3277) reduction estimated for the above-median subsample.[7] Moreover, this same shock induces an incremental 0.855% (0.0261×0.3277) increase in online payment amount relative to the offline channel during periods of below-median sentiment, compared to only a 0.35% (0.0106×0.3277) differential increase during periods of above-median sentiment.[8] Estimates using transaction volume and consumption variety as dependent variables exhibit consistent asymmetric patterns.

Taken together, these results imply that the decline in offline consumption and the short-run shift from offline to online channels associated with elevated sentiment are larger when sentiment is relatively low. This evidence supports an asymmetric relationship between sentiment and consumption. Appendix A.8.3 provides additional discussion of this asymmetry.

Furthermore, we investigate how sentiment influences consumption volatility. Specifically, we adopt a rolling-window approach to calculate consumption volatility. For each consumer in a given week, we calculate the standard deviation of their consumption over the subsequent one to eight weeks. We then normalize this standard deviation by the consumer's full-sample mean consumption to account for individual-specific scale differences, yielding a standardized measure of consumption volatility. Next, we estimate a modified version of the specification detailed in Section 2.3, taking these measures of consumption volatility as the dependent variables. Table 4 reports the results. Panel A demonstrates that elevated sentiment is associated with increased volatility in total consumption. Panel B decomposes this result by channel and reveals that elevated sentiment is associated with reduced volatility in offline consumption but heightened volatility in online consumption.

Next, we consider three dimensions of heterogeneity: characteristics of goods, consumers, and regions.

Fact 3. *(Heterogeneity by Goods Characteristics) The effects of sentiment on online and offline consumption are predominantly concentrated among nondurable goods and product categories with a lower online consumption percentage.*

We examine heterogeneity in the effects of sentiment on consumption by product durability. As detailed in Section 2.1, the consumption dataset provided by Alipay records each consumer's weekly payment amount and transaction volume across 17 distinct product categories. This level of granularity allows us to identify the product categories most sensitive to sentiment fluctuations. Following Dautović, Hau, and Huang (2024), we aggregate the 17 categories into three broad classifications: durables, nondurables, and others. Durables comprise housing and electronic goods, whereas nondurables encompass daily necessities, outdoor activities, clothing, and pets. We exclude the remaining categories from this specific heterogeneity analysis. Appendix A.8.5 provides a more granular, category-level analysis, including those product categories omitted from this section.

[7] In columns (1) and (2) of Table 3, the estimated coefficients on $Sentiment_{c,t-1}$ are -0.0142 and -0.0022, respectively.

[8] In columns (1) and (2) of Table 3, the estimated coefficients on $Sentiment_{c,t-1} \times Online_{i,c,f,t}$ are 0.0261 and 0.0106, respectively.

Based on this classification scheme, we calculate each consumer's weekly payment amount and transaction volume separately for durable and nondurable goods. We then estimate Equation (1) using channel-specific consumption of durables and nondurables as alternate dependent variables. Figure 3 illustrates that positive sentiment induces an increase in online consumption and a corresponding decrease in offline consumption for nondurable goods. Conversely, the estimated effects of sentiment on durable goods consumption are statistically indistinguishable from zero. Overall, these findings indicate that the short-run reallocation of consumption from the offline to the online channel is predominantly driven by nondurable consumption.

Furthermore, we explore product-level heterogeneity by online consumption percentage. Using the consumption dataset provided by Alipay, we compute the ratio of online payment amount to total payment amount for each of the 17 product categories. Based on this metric, we partition the categories into two groups characterized by a low versus a high online consumption percentage. We then aggregate weekly online and offline payment amounts, along with transaction volume, for each consumer separately within these two distinct categorical groupings. This data structure enables us to assess whether the sensitivity of consumption to sentiment shocks varies systematically with the underlying online consumption share of each product category. As illustrated in Figure 4, the estimates demonstrate that sentiment exerts a substantially larger impact on consumption within categories exhibiting a lower share of online consumption—that is, goods with a greater reliance on offline channels.

Fact 4. *(Heterogeneity by Consumer Characteristics) Online consumption responds to consumer sentiment more strongly among richer consumers (measured by past consumption level, past cash flow, age). Online consumption responds to consumer sentiment more strongly among consumers who are more emotional.*

Along the dimension of consumer wealth, several pieces of evidence consistently show that consumption responds to consumer sentiment more strongly among richer consumers.

First, we proxy for individual wealth using a proprietary metric of past consumption level, estimated by Alipay via an internal algorithm that synthesizes comprehensive consumer data. We sort consumers into quintiles based on their past consumption level and divide them into two distinct groups: the bottom two quintiles form the low-wealth group, and the top two quintiles form the high-wealth group, with the middle quintile excluded from this analysis. Re-estimating Equation 1 for each group, we plot the coefficients on $Sentiment_{c,t-1}$ and $Sentiment_{c,t-1} \times Online_{i,c,f,t}$, along with their 95% confidence intervals, in Figure 5. The results reveal that both the decline in offline consumption and the short-run reallocation toward online channels following elevated sentiment are heavily concentrated among wealthier consumers. For affluent consumers (top wealth quintile), the online consumption response to sentiment is 1.5 times larger than that for low-wealth consumers.

Second, we corroborate this finding using an alternative wealth proxy: past cash flow. Specifically, Alipay provides the aggregate cash transferred into each consumer's account throughout 2018

(the year preceding our sample period). Splitting consumers into high and low groups based on their past cash flow, the results in Figure 6 further support that consumers with higher past cash flows exhibit a substantially stronger consumption response to sentiment.

Third, we partition consumers into three age groups: 30 or younger, 31–40, and 41 or older. Re-estimating Equation 1 for consumers in each age group, we present the results in Figure 7. The estimates indicate that the decline in offline consumption and the short-run shift from offline to online channels associated with elevated sentiment are concentrated among the 31–40 age demographic, consistent with our preceding results.

Lastly, we examine distributional effects using the unconditional quantile regression framework following Borgen (2016). For a given quantile τ, we replace the dependent variable in Equation 1 with its recentered influence function (RIF). The RIF is defined as the unconditional τ-quantile plus an influence-function adjustment, which depends on whether an observation falls below that quantile and on the nonparametric estimate of the outcome density evaluated at τ. We estimate this density using a Gaussian kernel with a rule-of-thumb bandwidth. The estimated coefficients capture the marginal effect of sentiment on the unconditional τ-quantile of the consumption distribution. By repeating this estimation across quantiles from the 25^{th} to the 75^{th}, we trace how the impact of sentiment evolves across the consumption distribution. As illustrated in Figure 8, the magnitudes of both the sentiment-driven decline in offline consumption and the short-run reallocation toward online channels are larger at higher quantiles of the overall consumption distribution.

Along the dimension of emotional susceptibility, we employ three distinct proxies: (i) regional sentiment volatility; (ii) local public attention to a patriotic campaign, which heightens emotional arousal; and (iii) local public attention to a major consumer-fraud episode, which dampens emotional susceptibility. Across all three measures, our analysis consistently shows that consumption responds to consumer sentiment more strongly among consumers who are more emotional.

First, we capture emotional susceptibility using regional sentiment volatility. For a given city and week, we compute the variance in sentiment scores across all posts; higher variance signifies greater sentiment dispersion among residents, implying elevated emotional susceptibility. Splitting the sample at the median city-week sentiment volatility, we estimate Equation 1 separately for each subsample. Estimated results shown in Table 5 provide evidence that the effects of sentiment on online and offline consumption are more pronounced among consumers residing in cities exhibiting higher sentiment volatility.

Second, we proxy for emotional susceptibility using local public attention to the "1.4 Billion Flag Guardians," a mid-2019 patriotic social media campaign in China that triggered a surge in nationalist sentiment. To capture this exogenous, emotion-boosting shock, we calculate the week-over-week change in each city's Baidu Search Index[9] for the campaign keyword, with higher values indicating greater emotional susceptibility among local consumers. Figure 12 reports the results of

[9]The Baidu Search Index is a keyword search volume analysis tool that functions similarly to Google Trends, capturing public search interest on Baidu, one of the largest search engines in China.

estimating Equation 1 separately for consumers with below-median and above-median local public attention to the campaign.

Third, we use local public attention to a prominent consumer-fraud episode in 2019, "Live Streaming False Propaganda," as an emotion-dampening proxy. Following high-profile livestreaming scandals in 2019 that exposed deceptive sales tactics, public skepticism rose, prompting more cautious, less emotional shopping behavior. We measure this shock using the week-over-week change in each city's Baidu Search Index for the event keyword; higher index values denote periods of relatively lower emotional susceptibility. Figure 13 reports the results of estimating Equation 1 separately for consumers with below-median and above-median local public attention to this fraud event.

Collectively, the last two heterogeneity analyses yield consistent results: during periods of significant shocks that heighten emotional susceptibility, the sentiment-driven decline in offline consumption and the simultaneous reallocation to online channels are markedly amplified.

Fact 5. *(Heterogeneity by Regional Characteristics) Online consumption responds to consumer sentiment more strongly among regions that are richer (measured by local GDP per capita and financial constraints) and have more access to offline shopping.*

We examine city-level heterogeneity in the effect of consumer sentiment on consumption by local economic development and financial constraints.

First, we proxy local economic development using city-level GDP per capita in 2018, the year prior to our sample period. We partition consumers into two groups based on whether GDP per capita in their city of residence falls above or below the sample median, and we re-estimate Equation 1 separately for each subsample. As depicted in Figure 11, the estimates indicate that the sentiment-driven decline in offline consumption and the short-run reallocation toward online channels are more pronounced in more economically developed cities.

Second, we examine heterogeneity across two dimensions of consumer financial constraints: traditional and digital. To proxy for traditional credit constraints, we use the number of bank branches in the consumer's city of residence as of the end of 2018. Following Sun and Yannelis (2016), who exploit state-level bank branching deregulation as an exogenous shock that relaxes credit constraints for local households, we posit that a higher density of bank branches corresponds to looser local credit conditions. We partition the sample into two groups at the median number of bank branches: consumers residing in cities with less binding traditional credit constraints (bank branches above the median) and those in cities with tighter credit constraints (bank branches below the median). Re-estimating Equation 1 separately for each subsample, Figure 9 plots the estimated coefficients on $Sentiment_{c,t-1}$ and $Sentiment_{c,t-1} \times Online_{i,c,f,t}$ alongside their 95% confidence intervals. The results indicate that both the decline in offline consumption and the differential response of online versus offline consumption to a positive sentiment shock are more pronounced among consumers in cities with greater bank branch density.

To proxy for the local digital financial environment, we use the city-level Digital Financial Inclusion Index (DFII) for 2018 developed by Guo et al. (2020). Higher DFII values denote a more mature digital-finance infrastructure and expanded access to digital credit, thereby implying less binding digital financial constraints for local consumers. We partition the sample into two groups at the median DFII and re-estimate Equation 1 separately for each subsample. Figure 10 plots the estimated coefficients on $Sentiment_{c,t-1}$ and $Sentiment_{c,t-1} \times Online_{i,c,f,t}$ along with their 95% confidence intervals. The estimates reveal that both the effect of sentiment on offline consumption and the differential response of online versus offline consumption to positive sentiment shocks are stronger among consumers residing in cities with relaxed digital financial constraints. These results are consistent with the premise that digital financial access amplifies the transmission of consumer sentiment into spending, particularly through the online channel.

Beyond economic development, we examine heterogeneity by local offline consumption opportunities, which we proxy using city-level counts of points of interest (POIs) and shopping malls at the end of 2019. We source nationwide POI microdata from Dianping—one of China's premier business-directory platforms—which provides granular geographic coordinates and categorical classifications (e.g., Sports & Recreation, Food & Beverage). For each city, we compute the aggregate number of POIs and shopping malls in 2019. To test for heterogeneous effects of sentiment, we partition the cities at the sample median for each measure of offline consumption opportunity and estimate Equation 1 separately for each subsample. As reported in Table 6, the impact of sentiment on both online and offline consumption is more pronounced in cities endowed with a higher concentration of POIs or shopping malls.

2.5 Additional Results

2.5.1 Local Projections

In this section, we present the estimation results based on local projections. Compared to the OLS regression used earlier, local projections offer two key advantages: (i) they enable the estimation of the dynamic effects of sentiment on consumption across multiple future periods; (ii) unlike vector autoregression (VAR) models, local projections do not impose global recursive restrictions and do not require specifying the full dynamic system (potentially leading to misspecification bias). Instead, local projections provide a more flexible and robust approach by estimating impulse responses directly and independently for each horizon, without assuming a particular underlying data-generating process. While the existing literature typically employs SVAR approaches, our high-frequency weekly data and the granularity of individual-level observations allow us to adopt the more direct local projection approach. Following Òscar Jordà (2005) and Jordá, Schularick,

and Taylor (2016), the detailed specification of our local projections is as follows:

$$\Delta c_{i,c,f,t+h} = \beta_{h,1} \times \Delta\text{Sentiment}_{c,t-1} + \beta_{h,2} \times \Delta\text{Sentiment}_{c,t-1} \times \text{Online}_{i,c,f,t+h} + \sum_{l=1}^{L} \Gamma_{h,l} C_{i,c,t-l} + \gamma X_{c,t-1} + \text{FE}s + u_{i,c,f,t+h}, \ h = 0, \ldots, H \quad (3)$$

where i indexes the consumer, c the city of residence, $f \in \{0, 1\}$ the consumption channel (where $f = 0$ denotes offline and $f = 1$ denotes online), and t the week. The dependent variable is defined as $\Delta c_{i,c,f,t+h} \equiv \log(c_{i,c,f,t+h}) - \log(c_{i,c,f,t+h-1})$, where $c_{i,c,f,\tau}$ is the channel-specific consumption by consumer i residing in city c in week τ. Individual consumption is measured by (i) payment amount ($\text{Amount}_{i,c,f,t}$), (ii) number of transactions ($\text{Transactions}_{i,c,f,t}$), and (iii) consumption variety ($\text{Variety}_{i,c,f,t}$). The sentiment shock is $\Delta\text{Sentiment}_{c,t-1} \equiv \text{Sentiment}_{c,t-1} - \text{Sentiment}_{c,t-2}$, where $\text{Sentiment}_{c,\tau}$ denotes the sentiment level in city c during week τ. $\text{Online}_{i,c,f,t}$ is a dummy variable that equals one if the dependent variable corresponds to online consumption and zero for offline consumption. The term $C_{i,c,t-l}$ collects l lags of consumption vectors, and $X_{c,t-1}$ is the vector of contemporaneous city-week controls, including temperature, humidity, wind speed, sunshine, and air quality index. FEs include week, consumer, consumer×online, consumer×month, and consumer×month×online components. The coefficient $\beta_{h,1}$ captures the response of offline consumption to a positive sentiment shock, while $\beta_{h,2}$ represents the differential response of online versus offline consumption to a positive sentiment shock at horizon h.

We then estimate the impact of sentiment on overall consumption using the following specification:

$$\Delta c_{i,c,t+h} = \beta_h \times \Delta\text{Sentiment}_{c,t-1} + \sum_{l=1}^{L} \Gamma_{h,l} C_{i,c,t-l} + \gamma X_{c,t-1} + \text{FE}s + u_{i,c,t+h}, \ h = 0, \ldots, H \quad (4)$$

where $\Delta c_{i,c,t+h} \equiv \log(c_{i,c,t+h}) - \log(c_{i,c,t+h-1})$ and $c_{i,c,t+h}$ is the sum of online and offline consumption by consumer i residing in city c in week $t+h$. FEs include week fixed effects, consumer fixed effects, and consumer-month fixed effects. The coefficient β_h captures the response of overall consumption to a positive sentiment shock.

Figure 2 reports impulse responses of consumption to a positive sentiment shock, with 95% confidence bands depicted as dark-blue dotted lines. The bottom row plots the impulse responses of overall consumption (i.e., the estimates of β_h in Equation 4) and shows no statistically significant response at any horizon. The middle row plots impulse responses of offline consumption (i.e., the estimates of $\beta_{h,1}$ in Equation 3) and indicates a short-term reaction: a positive sentiment shock at week t significantly reduces offline consumption at $t+1$, reverses at $t+2$, and is indistinguishable from zero thereafter. The top row plots the impulse responses of the difference between online and offline consumption (i.e., the estimates of $\beta_{h,2}$ in Equation 3) and shows a corresponding short-run reallocation: a positive sentiment shock at week t significantly shifts consumption from the offline

to the online channel at $t+1$, reverses at $t+2$, and is indistinguishable from zero thereafter.

2.5.2 Other Robustness Checks

We conduct a series of additional checks in the Appendix that reinforce our baseline findings. First, Appendix A.8.6 shows that the results are robust to alternative sentiment constructions, including retaining product-quality posts, using cloud thickness as a sentiment proxy, applying alternative sentiment classifiers (DeepSeek and a dictionary-based method), and adopting an alternative prompt design. Second, Appendix A.8.7 shows that our conclusions are unchanged under alternative consumption measures, including IHS transformations, indicator outcomes, and one-week-ahead growth rates. Third, Appendix A.8.8 shows that the results are stable across alternative samples, including excluding major shopping-festival weeks, excluding the highest GDP-per-capita cities, and aggregating consumption to the city level. Finally, Appendix A.8.9 reports instrumental-variables estimates based on two distinct instruments: the second-stage patterns are consistent with the baseline, and the first stage confirms that both instruments significantly shift sentiment.

3 A Theory on Online and Offline Consumption

To rationalize the empirical facts documented in Section 2, we develop a two-period consumption-saving model with online and offline consumption. Our model augments the standard consumption-saving framework with a parsimonious extension: the introduction of online and offline shopping channels and consumer sentiment, while preserving analytical tractability. By explicitly modeling the role of sentiment, we characterize the mechanism through which sentiment fluctuations generate heterogeneous responses in online and offline consumption.

3.1 Consumption bundle

Consider a consumer living for two periods, t and $t+1$, who receives exogenous income Y in period t. Preferences are defined over a composite consumption bundle C_t and a shopping utility component. Lifetime utility U_t is given by:

$$U_t \quad = \quad \log(C_t) + \beta \log(C_{t+1}) + u_t, \tag{5}$$

where $\beta \in (0,1)$ is the subjective discount factor.

The composite consumption bundle C_t is defined by a constant elasticity of substitution (CES) aggregator over a continuum of differentiated goods $j \in [0,1]$. The key departure from the standard setup is that the effective quality of each variety depends on the consumer's shopping channel choice. Let $\mathbf{1}_{jt} \in \{0,1\}$ be an indicator equal to 1 if the consumer purchases good j online and 0 otherwise. The matching value of good j purchased online is γ_{jt}, and the matching value of good j purchased

offline is normalized to 1. Hence:

$$C_t = \left(\int_0^1 \left[\mathbf{1}_{jt} \cdot \gamma_{jt} + (1 - \mathbf{1}_{jt}) \right] C_{jt}^{\frac{\theta-1}{\theta}} \, dj \right)^{\frac{\theta}{\theta-1}}, \tag{6}$$

where $\theta > 1$ is the elasticity of substitution across varieties and C_{jt} is the quantity of good j consumed in period t. The stochastic term $\gamma_{jt} > 0$ captures the idiosyncratic matching value the consumer obtains from purchasing variety j online, reflecting the precision of fit or alignment with personal tastes. We assume γ_{jt} is independently and identically distributed (i.i.d.) across goods and across consumers.[10] Specifically, γ_{jt} follows a log-normal distribution, $\log(\gamma_{jt}) \sim \mathcal{N}(\mu_j, \sigma_j^2)$. A higher μ_j implies that online shopping for type-j goods tends to yield a better match and, other things equal, a higher online expenditure share for that product type. We further assume $\sigma^2 < \mu$ to ensure that the distribution of online matching values is sufficiently concentrated, so that the online channel remains economically viable for a non-negligible fraction of consumers.

Similarly, the consumption bundle in period $t+1$ is given by:

$$C_{t+1} = \left(\int_0^1 \left[\mathbf{1}_{jt+1} \cdot \gamma_{jt+1} + (1 - \mathbf{1}_{jt+1}) \right] C_{jt+1}^{\frac{\theta-1}{\theta}} \, dj \right)^{\frac{\theta}{\theta-1}}, \tag{7}$$

where γ_{jt+1} shares the same distributional assumptions as γ_{jt}.

The consumer faces the following budget constraints:

$$\int_0^1 P_{jt} C_{jt} \, dj + S_t \leq Y, \tag{8}$$

$$\int_0^1 P_{jt+1} C_{jt+1} \, dj \leq R \, S_t, \tag{9}$$

where P_{jt} is the price of variety j, S_t denotes savings, and R is the exogenous risk-free gross interest rate.

To isolate the effect of consumer sentiment on shopping behavior, we assume that channel choice is not driven by differences in the fundamental characteristics of the goods. Specifically, we abstract from disparities in quality or other intrinsic attributes that might otherwise govern the selection of a shopping channel.[11] We therefore assume that, for a given product type j, prices are identical across online and offline channels (Dolfen et al., 2023; Jo, Matsumura, and Weinstein, 2024). This assumption is empirically supported by Cavallo (2017) and Cavallo (2018), who document that prices are largely synchronized across online and offline retailers. To focus on the role of sentiment shocks, we also abstract from uncertainty in period $t+1$.[12]

[10] γ_{jt} is also independent of the price P_{jt}: because γ_{jt} is i.i.d. across consumers, the seller of good j faces deterministic aggregate demand and sets P_{jt} to clear the market.

[11] In reality, the fundamental characteristics of goods and services may correlate with shopping channels, but this lies beyond the scope of our analysis.

[12] At high frequency, we do not consider changes in the risk-free rate or consumer income. We also abstract from

3.2 Shopping utility

The shopping channel decision $\mathbf{1}_{jt}$ for each variety depends on the term u_t in (5), which represents the non-pecuniary utility (or disutility) derived from the shopping experience. It is defined as a CES-type aggregator over the shopping utilities from both channels across the two periods:

$$u_t = \int_0^1 \left\{ \left[\mathbf{1}_{jt} \cdot u^o_{jt} + \mathbf{1}_{jt+1} \cdot u^o_{jt+1} \right]^\omega + \left[(1 - \mathbf{1}_{jt}) \cdot u^f_{jt} + (1 - \mathbf{1}_{jt+1}) \cdot u^f_{jt+1} \right]^\omega \right\} dj, \tag{10}$$

where u^o_{jt} (u^f_{jt}) is the online (offline) shopping utility for good j in period 't', with u^o_{jt+1} and u^f_{jt+1} defined analogously for period $t+1$. The parameter $\omega \in (0,1)$ governs the degree of complementarity in shopping utility across the two periods.

The net shopping utility for variety j in period t combines a utility and a disutility component. When shopping online, the consumer derives utility $v^o_j \eta^\alpha_t$ from the shopping experience; when shopping offline, the corresponding utility is $v^f_j \eta^\alpha_t$. The terms v^o_j and v^f_j capture channel-specific shopping utility for type-j goods.

The random variable η_t captures a consumer sentiment shock: a higher η_t means the consumer enjoys shopping more. This shock has two key features. First, it is transitory, consistent with existing studies (e.g., Benhabib, Wang, and Wen (2015); Angeletos, Collard, and Dellas (2018)). Second, it is independent of fundamentals ?it affects neither the intrinsic characteristics of goods nor the consumer's budget constraint. This purely non-fundamental nature distinguishes sentiment shocks from "news shocks" or "noise shocks," which pertain to fundamental factors and typically generate wealth effects (e.g., Barsky and Sims, 2012). The parameter $\alpha \in (0,1)$ captures concavity in shopping utility, reflecting diminishing marginal returns from additional shopping within a given product type.

The consumer also incurs a channel-specific disutility z^o_j (z^f_j) from online (offline) shopping, reflecting search effort and time costs. Offline shopping carries additional disutility from physical effort and travel costs. The net shopping utilities are therefore $(v^o_j \eta^\alpha_t - z^o_j)$ for online and $(v^f_j \eta^\alpha_t - z^f_j)$ for offline.

Furthermore, we assume that the online and offline shopping utilities for type-j goods are proportional to total expenditure $\int_0^1 P_{jt} C_{jt} dj$. The steady-state values of $\left(v^o_j - z^o_j\right)$ and $\left(v^f_j - z^f_j\right)$ capture the relative importance of the shopping experience for type-j goods in the consumer's overall shopping activity. The shopping utilities for period t are then defined as:

$$u^o_{jt} = \left(v^o_j \cdot \eta^\alpha_t - z^o_j\right) \left(\int_0^1 P_{jt} C_{jt} dj \right), \tag{11}$$

$$u^f_{jt} = \left(v^f_j \cdot \eta^\alpha_t - z^f_j\right) \left(\int_0^1 P_{jt} C_{jt} dj \right), \tag{12}$$

period-$t+1$ household income; its inclusion does not affect our results.

The shopping utilities for period $t+1$ are defined analogously:

$$u^o_{jt+1} = \left(v^o_j - z^o_j\right)\left(\int_0^1 P_{jt+1}C_{jt+1}dj\right), \tag{13}$$

$$u^f_{jt+1} = \left(v^f_j - z^f_j\right)\left(\int_0^1 P_{jt+1}C_{jt+1}dj\right), \tag{14}$$

where we abstract from η_{t+1} since our focus is on the effect of sentiment in period t.

To ensure that the consumer derives positive utility from shopping, we assume $v^o_j > z^o_j$ and $v^f_j > z^f_j$ around the steady state of η_t. Compared to the offline channel, the online channel offers a distinct advantage: online shopping provides seamless access to a vast array of goods, allowing consumers to obtain instant gratification without the physical burden of travel or exposure to inclement weather (Dinerstein et al., 2018). Accordingly, we normalize the offline experience utility to zero and impose $v^o_j > v^f_j$. Moreover, relative to the offline channel, online platforms facilitate price comparisons, customer reviews, and screening of merchant characteristics, which substantially reduces search costs (Gu and Wang, 2022; Moraga-González, Sándor, and Wildenbeest, 2022). That is, $z^f_j > z^o_j$.

The sequence of events in period t unfolds as follows:

Stage 1 (*Consumption-saving decision*): At the beginning of period t, the consumer sentiment shock η_t is realized. The representative consumer chooses total consumption expenditure $\int_0^1 P_{jt}C_{jt}dj$ to maximize lifetime utility.

Stage 2 (*Shopping and channel choice*): After the intertemporal allocation, the consumer shops for individual varieties, at which point the idiosyncratic matching shocks γ_{jt} are realized. For each variety, the consumer compares the online and offline options and solves a discrete choice problem to determine the purchase channel.

3.3 The consumer's decision on total expenditure

We model the consumer's decision problem regarding online and offline consumption as a two-stage optimization. In the first stage, the consumer solves the intertemporal consumption-savings trade-off, which determines total consumption expenditure $\int_0^1 P_{jt}C_{jt}dj$. In the second stage, the consumer allocates this expenditure across shopping channels for each variety $j \in [0,1]$.

Applying standard CES demand analysis and the Euler equation for intertemporal decision making, we have the following result:

Proposition 1. ***(Effect of sentiment on total consumption expenditure)*** *Total consumption expenditure in periods t and $t+1$ is invariant to non-fundamental sentiment.*

Proof: See Appendix B.1

Because sentiment is non-fundamental and orthogonal to any fundamental properties of goods, it does not shift total expenditure. Instead, total consumption expenditure is determined by the

intertemporal trade-off between the marginal propensity to consume and the propensity to save.

It is important to distinguish our findings from the spending surges typically observed during shopping festivals. Such spikes are generally driven by strategic discounts or intertemporal price discrimination — mechanisms that are entirely distinct from the channel we study. To eliminate the influence of these factors, we include week fixed effects and consumer-month fixed effects in our empirical specifications to absorb transitory, price-driven fluctuations. Our findings are also robust to the exclusion of major Chinese shopping festivals (Table A12).

On the one hand, this neutrality result is consistent with our empirical evidence illustrated in Table 1 (Panel A) and the impulse response of Figure 2 (third row), explaining *Fact 1*. Furthermore, it aligns with the broader empirical literature, which documents that even when sentiment significantly predicts aggregate consumption, the economic magnitude remains modest (Carroll, Fuhrer, and Wilcox, 1994; Ludvigson, 2004).

On the other hand, this result may appear to stand in tension with existing findings in the literature. On the theoretical side, Angeletos and La'O (2013), Barsky and Sims (2012) and Benhabib, Wang, and Wen (2015) show that non-fundamental sentiments can drive fluctuations in the real economy. On the empirical side, studies exploiting exogenous variation in sentiment from elections (Benhabib and Spiegel, 2018; Gillitzer and Prasad, 2018), terrorist attacks and mass shootings (Brodeur, 2018; Lagerborg, Pappa, and Ravn, 2022) document that non-fundamental sentiment shocks have real economic consequences. In contrast, we find that consumer sentiment has little impact on the total expenditure of an individual consumer. How can this apparent tension be reconciled? We argue that these existing studies, taken together, reflect a self-fulfilling mechanism: optimism induces behaviors that, in the aggregate, validate the initial optimistic beliefs. However, existing research has rarely been able to disentangle each step in this propagation chain—for instance, how a transient sentiment shock affects a single decision-maker's choices. Our high-frequency data allow us to focus precisely on the individual consumer's response, rather than on the aggregate equilibrium outcome, which would evolve much more slowly. Our findings thus capture one link in the transmission chain of sentiment, and this link reveals that because optimism does not alter the consumer's fundamental circumstances, it does not affect her total consumption expenditure.

3.4 The consumer's decision on shopping and channel: online vs. offline

Having determined total expenditure in t, the consumer allocates her budget across specific varieties through either online or offline shopping. We focus here on the intra-period shopping behavior and the consumer's choice of shopping channel.

For each type-j good, the consumer's channel choice involves a trade-off between two distinct sources of utility. The first is a consumption-side gain: purchasing online exposes the consumer to an idiosyncratic matching value γ_{jt} that governs the effective consumption of the good. A higher $\gamma_{jt} > 1$ amplifies the quantity-equivalent consumption through the CES demand function $C_{jt} =$

$[\mathbf{1}_{jt}\gamma_{jt} + (1 - \mathbf{1}_{jt})]^{\theta}(P_{jt}/P_t)^{-\theta}C_t$, whereas purchasing offline yields a baseline effective consumption normalized to unity ($\gamma_{jt} = 1$). The second is a shopping-experience differential: online shopping yields u^o_{jt}, while offline shopping yields u^f_{jt}. The total expected benefit of purchasing online in period t is therefore the sum of the instantaneous consumption gain and the expected shopping utility from both periods conditional on today's choice, minus the corresponding expected benefit of purchasing offline. The consumer therefore chooses to purchase good j online if and only if the total expected benefit exceeds that of the offline alternative, which defines an endogenous threshold γ^*_{jt}. Analogous logic governs the consumer's channel choice in period $t + 1$.

Proposition 2 formally characterizes the consumer's optimal channel choice for each type-j good.

Proposition 2. *(**Optimal shopping channel choice**) For a type-j good, the consumer's optimal shopping channel decision in period t is given by:*

$$\text{purchase} \begin{cases} \text{online}, & \text{if } \gamma_{jt} \geq \gamma^*_{jt}; \\ \text{offline}, & \text{if } \gamma_{jt} < \gamma^*_{jt}; \end{cases} \tag{15}$$

where the threshold γ^*_{jt} is endogenously determined. The consumer's optimal shopping channel decision in period $t + 1$ is analogous.

Proof: See Appendix B.2.

The intuition underlying this proposition is that if the online merchant encountered is sufficiently attractive – that is, if γ_{jt} exceeds the threshold γ^*_{jt} – the consumer purchases online; otherwise, the consumer purchases the good offline.

In what follows, we maintain the following assumption on the threshold:

Assumption 1. *The threshold satisfies $\gamma^*_{jt} > 1$ for both periods.*

Recall that γ^*_{jt} denotes the reservation matching value in period t. A threshold exceeding unity ($\gamma^*_{jt} > 1$) implies that the consumer requires an *above-average* matching value to purchase online in period t. This condition captures the inherent frictions of online shopping: consumers cannot physically inspect products, face shipping delays, and incur return costs, all of which create a systematic hurdle that the matching value must overcome. Since in our model online shopping offers higher shopping utility ($v^o > v^f$) and lower search costs ($z^o < z^f$) than offline shopping, without such a hurdle (i.e., if $\gamma^*_{jt} \leq 1$), the consumer would opt for the online channel even when the match is below average, implying that online shopping dominates offline unconditionally—a scenario that contradicts the observed coexistence of both channels in the data.[13]

[13] In our sample, online transactions account for 43% of all transactions, and online expenditure constitutes 36% of total expenditure.

3.5 Effects of sentiment on online and offline expenditure

We begin by characterizing the online expenditure share at the variety level before aggregating to the overall online share. For each type-j good, consumers draw an idiosyncratic matching value γ_{jt} from the log-normal distribution $F_j(\cdot)$. Given the optimal threshold γ_{jt}^* derived in Proposition 2, consumers with $\gamma_{jt} \geq \gamma_{jt}^*$ purchase online, while those with $\gamma_{jt} < \gamma_{jt}^*$ purchase offline. The online expenditure share for type-j goods can therefore be expressed as the fraction of total expenditure on good j that flows through the online channel:

$$\omega_{jt}^o = \frac{\int_{\gamma \geq \gamma_{jt}^*} P_{jt}\gamma^\theta (P_{jt}/P_t)^{-\theta} dF_j(\gamma)}{\int_{\gamma \geq \gamma_{jt}^*} P_{jt}\gamma^\theta (P_{jt}/P_t)^{-\theta} dF_j(\gamma) + \int_{\gamma < \gamma_{jt}^*} P_{jt}(P_{jt}/P_t)^{-\theta} dF_j(\gamma)}. \tag{16}$$

The numerator captures online expenditure (where the matching shock amplifies effective consumption by γ_{jt}^θ), and the denominator captures total expenditure on good j across both channels. Since γ_{jt} is i.i.d. and the price term $P_{jt}(P_{jt}/P_t)^{-\theta}$ cancels out, the expression simplifies to:

$$\omega_{jt}^o = \frac{\int_{\gamma \geq \gamma_{jt}^*} \gamma^\theta dF_j(\gamma)}{\int_{\gamma \geq \gamma_{jt}^*} \gamma^\theta dF_j(\gamma) + F_j(\gamma_{jt}^*)} \tag{17}$$

where $F_j(\gamma_{jt}^*)$ is the mass of consumers that purchase offline.

From this variety-level expression, we can derive the aggregate online expenditure share. Specifically, the aggregate online expenditure share ω_t^o and its offline counterpart $\omega_t^f = 1 - \omega_t^o$ are defined as the total expenditures across all varieties through each channel, divided by total expenditure:

$$\omega_t^o = \frac{\int_0^1 \mathbf{1}_{jt} P_{jt} C_{jt} dj}{\int_0^1 P_{jt} C_{jt} dj}, \qquad \omega_t^f = \frac{\int_0^1 (1 - \mathbf{1}_{jt}) P_{jt} C_{jt} dj}{\int_0^1 P_{jt} C_{jt} dj}, \tag{18}$$

where $\mathbf{1}_{jt}$ is an indicator that equals one if the consumer purchases good j online and zero otherwise.

The cental interest of this paper is the impacts of "non-fundamental" sentiment shocks on online and offline expenditure. We turn to the following proposition.

Proposition 3. *(**Effect of sentiment on online and offline expenditure**) when consumer sentiment increases, the consumer allocates more expenditure online and less offline: $\partial \omega_{jt}^o / \partial \eta_t > 0$ and $\partial \omega_t^o / \partial \eta_t > 0$.*

Proof. See Appendix B.3. ■

The intuition is as follows. A positive sentiment shock enhances the perceived shopping utility of online purchases relative to offline purchases. This shifts the relative payoff structure facing the consumer, lowering the threshold γ_{jt}^* and thereby increasing the fraction of households that choose the online channel for each variety. Consequently, S_{jt}^o rises for every type-j good, and the aggregate

online share S_t^o increases. This channel operates purely through reallocation: because sentiment does not alter the consumer's intertemporal budget constraint or the goods' fundamental attributes, total consumption expenditure remains unchanged, consistent with Fact 1.

The implication of Proposition 3 is consistent with Panel B of Table 1. The coefficient on $\text{Sentiment}_{c,t}$ is negative, while the coefficient on the interaction term $\text{Sentiment}_{c,t} \cdot \text{Online}_{i,c,f,t+1}$ is positive. More interestingly, the magnitudes of these two coefficients imply that in response to a positive sentiment shock, the value of transactions, the number of transactions, and the variety of transactions for offline consumption all decrease, while the corresponding variables for online consumption all increase by nearly the same degree. This pattern suggests that online consumption "crowds out" offline consumption: total expenditure does not change significantly, but online expenditure rises while offline expenditure declines by a similar magnitude. This reallocation effect distinguishes our findings from the wealth-effect channel emphasized in prior literature.

3.6 Effects of sentiment on consumption dynamics and volatility

From Proposition 3, we can further demonstrate the online consumption dynamics have the following property:

Proposition 4. *(**Effect of sentiment on dynamics of online expenditure**) Following a positive sentiment shock, online expenditure exhibits an initial increase followed by a subsequent decline, whereas offline expenditure follows an inverse trajectory, decreasing initially before recovering. That is,* $\partial \omega_{t+1}^o / \partial \eta_t < 0$.

Proof. please see Appendix B.4. ■

Explaining Fact 2: The above proposition explains the divergent dynamics of the online and offline expenditure displayed in Figure 2 (first and second row). The economic intuition is two-fold: First, micro-level consumer sentiment is usually transitory rather than persistent, after a positive shock to sentiment, the sentiment is likely to go back to its average level. When the consumer sentiment reverse to the average level, she does not has strong tempt to consume more. Second, as an consumer increase her immediate consumption of online goods in response to heightened sentiment, the marginal utility of such consumption diminishes. This satiation effect, together with her budget constraint, therefore weakens the incentive for further online acquisitions. Consequently, the consumer undergoes a substitution process, shifting their marginal expenditure from online to offline channels following the initial spike in online spending.

Although there is no change in expectations of either the exogenous fundamentals at any horizon or the endogenous outcomes in the medium to long run, a short-lived non-fundamental consumer sentiment shock also produces dynamic movements and reversals in online and offline consumption. Once again, controlling a rich set of fixed effect helps us filter out the short-run impacts of sentiments on consumption and our featuring pattern that high-frequency data captures the short-run impacts

which highlights stronger reallocation effects and weaker wealth effects, which is quite different from the existing literature which dominantly study the medium or long run effects of sentiments on the macroeconomy.

This dynamic path leads directly to the following corollary, which is consistent with the empirical evidence presented in Table 4.

Lemma 1. *A transitory sentiment shock induces greater intertemporal volatility in online expenditure.*

Asymmetry: When consumer sentiment is low, a small increase is relatively significant. When consumer sentiment is low, the marginal utility gain from shopping is relatively high, the convienience and instaneous gratification from online shopping can alleviate low mood, thus a positive sentiment shock has relatively strong effect on improving online shopping. When sentiment is high, the satisfication from shopping is already high, the marginal utility gain from a bit more shopping is relatively low.

3.7 Effects of sentiment by product attributes

Next, we examine the effects of consumer sentiments on online and offline consumption across three dimensions: product attributes, consumer characteristics and regional factors. We begin by investigating how the sensitivity of online consumption to sentiment varies according to the two key product attributes: (i) the degree of durability (durable goods, non-durable goods and services); and (ii) product suitability for online channels.

The responsiveness of consumers to sentiment is closely tied to the nature of the good's utility flow. Durable, non-durable goods and services differ fundamentally in their intertemporal utility profiles. To simplify our analysis, we assume that the online and offline matching shocks, as well as prices, are time-variant, such that, $\gamma^o_{jt} = \gamma^o_{jt+1}$, $\gamma^f_{jt} = \gamma^f_{jt+1}$, and $P_{jt} = P_{jt+1}$. While non-durable goods are consumed within a single period, durable goods provide a stream of utility across both period t and $t+1$. The durability of goods affects the response of consumers to sentiment in the following way.

Proposition 5. *The response of online expenditure to sentiments is weaker for durable goods than non-durable goods.*

Proof. See Appendix B.7. ■

Explaining Fact 3: The above proposition suggests that consumers are more cautious – and are thus less swayed by sentiment – when purchasing durable goods. This theoretical result squares with the empirical result found in Figure 3. The underlying rationale is also intuitive: because durable goods generate utility over an extended horizon, the purchase decision is more heavily weighted toward fundamental long-run factors rather than transitory sentiment shocks. Consequently, the

marginal impact of sentiment is attenuated for goods with longer service lives. This finding also echoes Zeldes (1989) which suggests that less durable goods consumption is more sensitive to transitory shocks.[14]Here, services are excluded from the our sample, as they typically require on-site (or local) consumption regardless of the purchasing channel (Bounie, Camara, and Galbraith, 2025).

Further we consider the "online-friendliness" of products. Products frequently purchased online are typically characterized by lower information asymmetry, higher substitutability between online and offline variants, and lower shipping costs. In short, these varieties possess attributes that facilitate online matching. Formally, we estabilish the following:

Proposition 6. *The response of online expenditure to sentiments is stronger for goods with high online-market suitability. That is,* $\partial^2\omega_{jt}^o/\partial\eta_t\partial\mu_j > 0$.

Proof. See Appendix B.8. ■

Explaining Fact 3: The above proposition squares with Figures A9 and 4, supporting Fact 3. The reason for this result is that an increase in consumer sentiment amplifies the comparative advantage of online platforms, thereby intensifying the "crowding out" of offline transactions.

3.8 Effects of sentiment by consumer attributes

We now examine how the impact of sentiment on online consumption varies with consumer-level heterogeneity. Specifically, using our theoretical framework, we focus on two pivotal dimensions: consumer wealth and emotionality.

3.8.1 Consumer's wealth

We use the exogenous income, Y_t, as a proxy for the consumer's wealth. Our model demonstrates that the sensitivity of online expenditure to sentiment is non-uniformly distributed across wealth levels.

Proposition 7. *The response of online expenditure to sentiment is stronger among more wealthier consumers. That is,* $\partial^2\omega_{jt}^o/\partial\eta_t\partial Y > 0$.

Proof. See Appendix B.5. ■

Explaining Fact 4: The above proposition indicates that wealthier consumers exhibit a more pronounced response in online consumption to sentiment shocks. The underlying economic intuition is rooted in the declining marginal utility of consumption: low-wealth individuals face high marginal utility (low C_t), In other words, the opportunity cost of the short-run consumption is higher for

[14]Another reason for the differential response of durable goods and non-durable goods to consumer sentiment is that durable goods may be tied to loans which is more strongly subject to fundamental factors.

consumers who have more binding budget constraint. Conversely, wealthier consumers, facing less binding budget constraint, have a lower opportunity cost of acting on sentimental impulses, thus exhibiting a stronger reallocation effect. This result is consistent consistent with the precautionary motive of consumers (Acemoglu and Scott, 1994).

This proposition provides an explanation for Fact 4. While our dateset does not explicitly report wealth, we use several empirical proxies to validate this relationship, as shown in Figures 5 to 10. First, we use historical income and cash flows in consumers' Alipay account, with higher past spending or cash flows indicating larger wealth. Second, we exploit life-cycle variation in disposable income, noting that middle-aged group (31-40) typically possess higher discretionary wealth compared to younger and older cohorts. Finanlly, we consider financing constraints – specifically traditional and digital credit limit – as indicators of the effective wealth available for consumption.

3.8.2 Consumer's emotionality

To understand the impact of non-fundamental sentiment on consumer's behavior, we analyze how this impact varies with the consumer's emotional state. We define the ''emotionality" of consumer as the weight consumers place on sentiment relative to fundamentals of products. In our model, more emotional consumers are characterized by a higher parameter v. The following proposition describes how this degree of emotionality dictates the consumption response.

Proposition 8. *The response of online expenditure to sentiment is stronger among more emotional consumers. That is,* $\partial^2\omega^o_{jt}/\partial\eta_t\partial(v^o_j - v^f_j) > 0$.

Proof: See Appendix B.6.

Explaining Fact 4: This prediction is consistent with the empirical evidence in Figures 12 and 13, which show that sentiment-driven effects are heightened during periods of elevated consumer emotionality. The intuition is straightforward: emotionality causes consumers to be more easily affected by sentiment, with online expenditure exhibiting stronger response to sentiment.

3.9 Effect of sentiment by regional attributes

It is naturally expected that the regional factors also plays a role in determining a consumer's allocation toward online expenditure. In our framework, we consider the regional income and accessibility of offline shopping. The former factor is captured by the consumer's income as in the previous subsection, while the latter factor is parameterized by the search cost differential, $z^f - z^o$, which may reflect a higher density of physical markets or shopping malls. This leads to the following characterization of how regional characteristics moderate the influence of sentiment.

Proposition 9. *The response of online expenditure to sentiments is weaker in higher-income regions, in regions with weaker credit constraint, and in regions with greater accessibility to offline shopping. That is,* $\partial^2\omega^o_{jt}/\partial\eta_t\partial z^f_j < 0$.

Proof. See Appendix B.8. ■

Explaining Fact 5: The above proposition provides a theoretical foundation for the empirical patterns documented in Table 6 and Figure 11, thereby corroborating Fact 5. The intuition regarding regional income parallels that of Proposition 7. Enhanced offline access reduces the search cost gap betwen online and offline channels, making offline shopping both more attractive and convenient. Consequently, in regions with an abundance of offline options, the impact of sentiment on online expenditure is dampened. This result is consistent with the evidence reported by Forman, Ghose, and Goldfarb (2009) and Couture et al. (2021).

In regions with weaker credit constraints, liquidity is abundant, and the transimission of sentiment to consumption is less frictional, thus the reallocation effect is stronger. This result is consistent with Figure 9.

4 Conclusion

This paper provides new empirical and theoretical insights into the role of sentiment in driving consumption fluctuations within the digital economy. Leveraging high-frequency, transaction-level data from a leading Chinese e-commerce platform, our empirical analysis demonstrates that a positive non-fundamental sentiment shock triggers a significant and substantial increase in online consumption, accompanied by a comparable decrease in offline consumption. This pattern indicates a pronounced reallocation effect from offline to online channels. This reallocation reverses significantly in subsequent periods, generating excessive volatility in consumption along both channels. Moreover, the effects of sentiments are heterogeneous, varying systematically with product characteristics, consumer demographics, and regional factors.

To rationalize these empirical findings, we develop a two-channel consumption-saving model incorporating consumer sentiment. A key feature of this model is the presence of two distinct shopping channels – online and offline – with online shopping characterized by low search costs and enhanced shopping utility. A rise in consumer sentiment amplifies the perceived utility of the online experience, thereby creating a wedge that allocates household expenditure toward online platforms, even as total aggregate expenditure remains unchanged in response to the non-fundamental shock. The model further account for several additional empirical patterns, including dynamic reversals, asymmetric responses, and the observed heterogeneity across product categories and consumer attributes.

Overall, our analysis – based on high-frequency micro-level data – underscores the critical role of channel heterogeneity in the transmission of sentiment to consumption. To the best of our knowledge, this study is the first to document the reallocation effect between online and offline channels, an effect that remains undetectable in aggregate-level analyses. While our findings are based on a substantial sample of Chinese consumers, the generalizability of these results to other

institutional contexts remains an important avenue for future research.

Figures and Tables

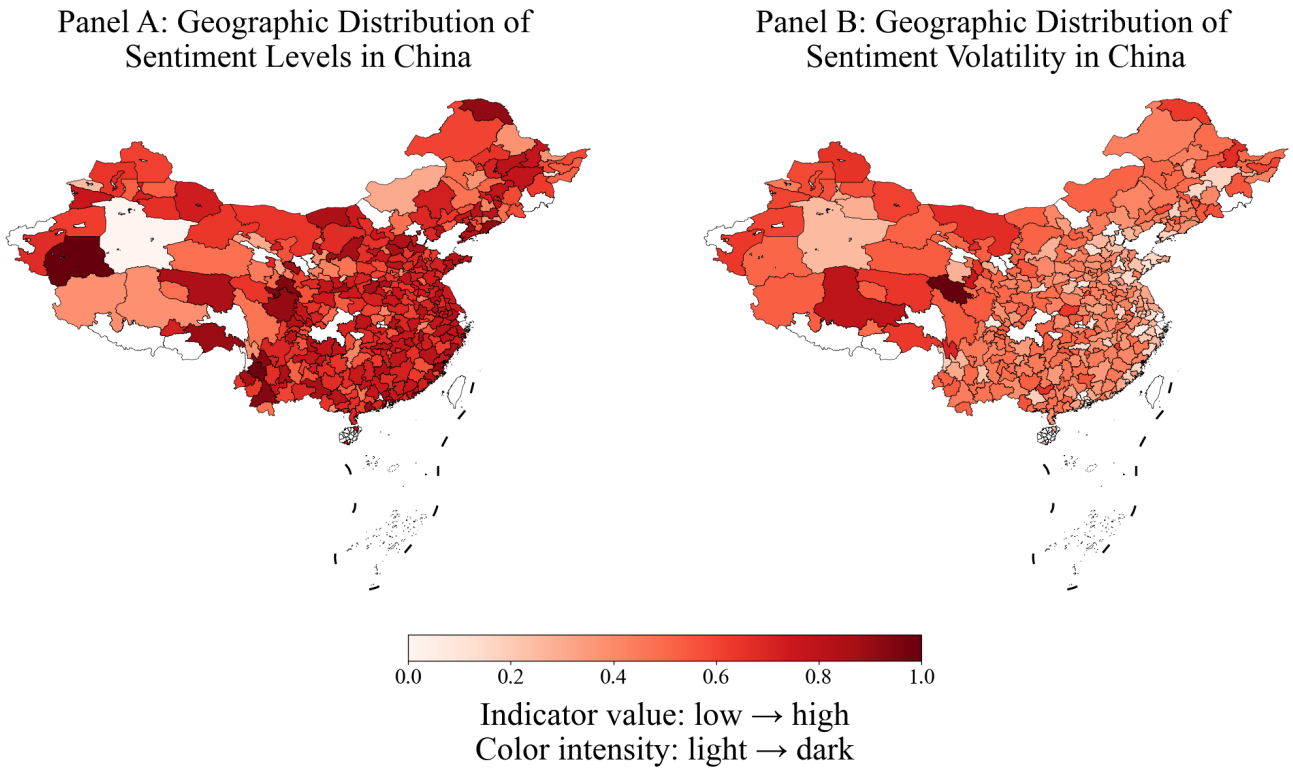


Figure 1: Geographic Distribution of Sentiment in China

Notes: This figure illustrates the cross-city distribution of sentiment. Panel A shows the distribution of average sentiment levels across cities, while Panel B displays the distribution of sentiment volatility. City-level average sentiment is calculated as the mean of weekly sentiment over 52 weeks for each city, then normalized to the [0,1] range before plotting the distribution map. City-level sentiment volatility is computed as the dispersion of weekly sentiment over 52 weeks for each city and also normalized to the [0,1] range. Color intensity reflects the magnitude of each measure, with darker red indicating higher values. White areas represent cities with no observations or a value of zero.

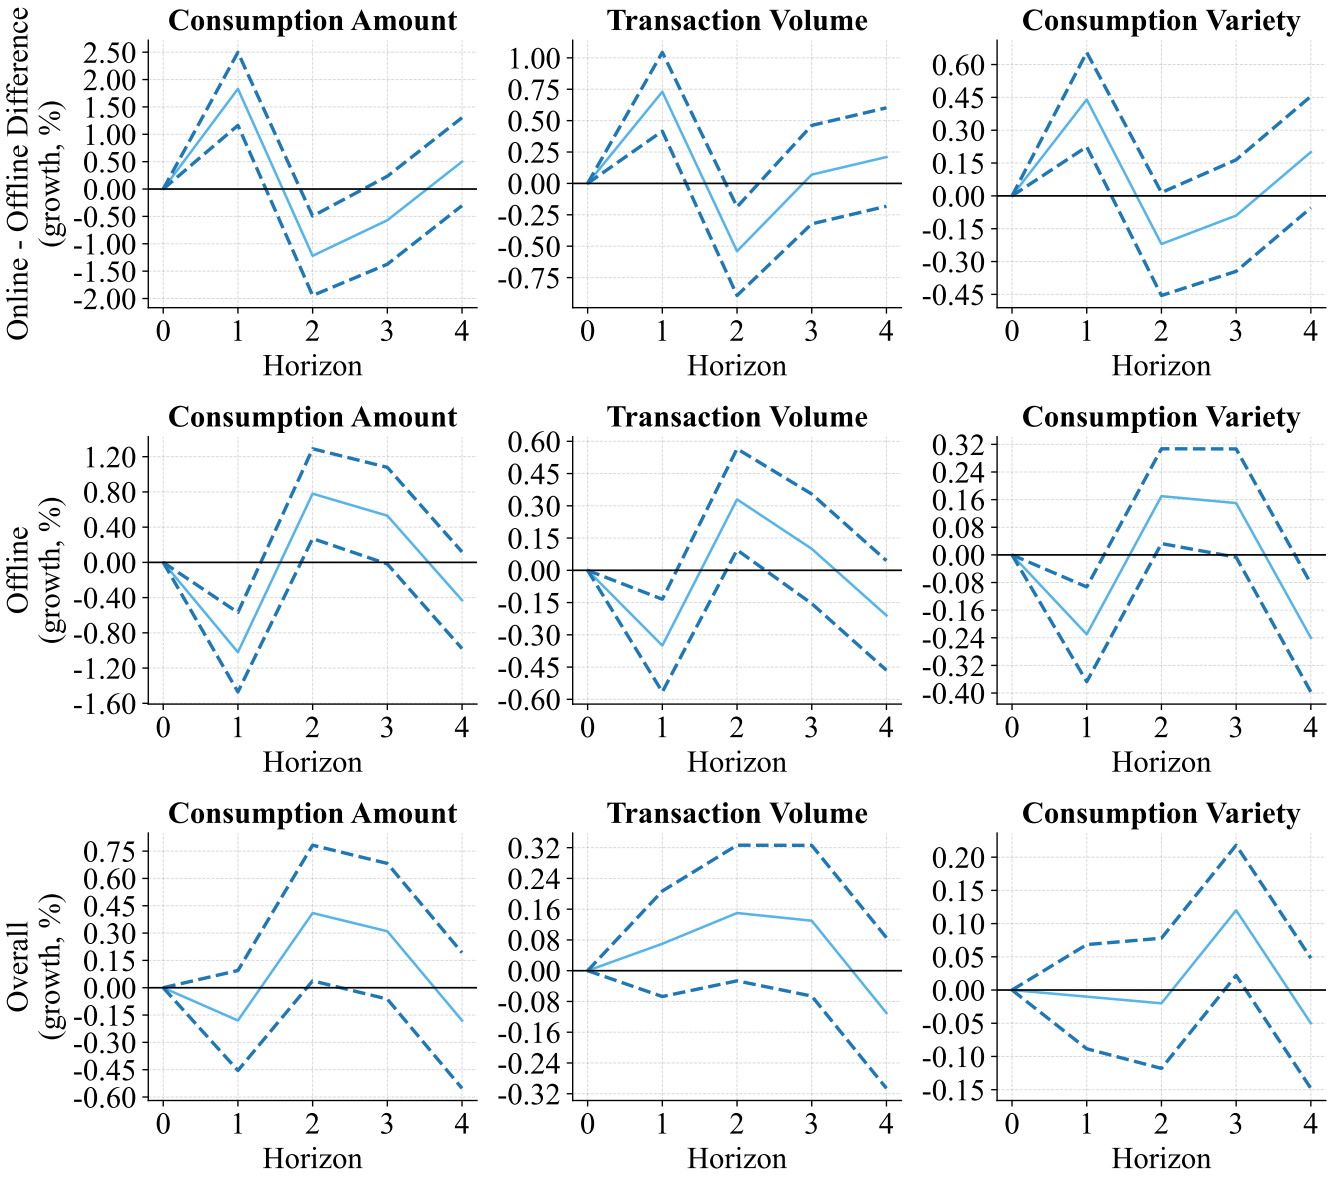


Figure 2: Impulse Responses of Consumption to a Positive Sentiment Shock

Notes: This figure plots the impulse responses of the online–offline differential, offline, and overall consumption (in growth rate expressed in percentage points) to a positive sentiment shock, estimated using the local projection method described in Section 2.5.1. Each row contains three panels, with dependent variables corresponding to growth in consumption amount, transaction volume, and consumption variety, respectively. The solid blue lines represent the point estimates of the coefficients $\hat{\beta}_{h,2}$ (online–offline differential; Equation 3), $\hat{\beta}_{h,1}$ (offline; Equation 3), and $\hat{\beta}_h$ (overall; Equation 4). The dark-blue dotted lines indicate 95% confidence intervals.

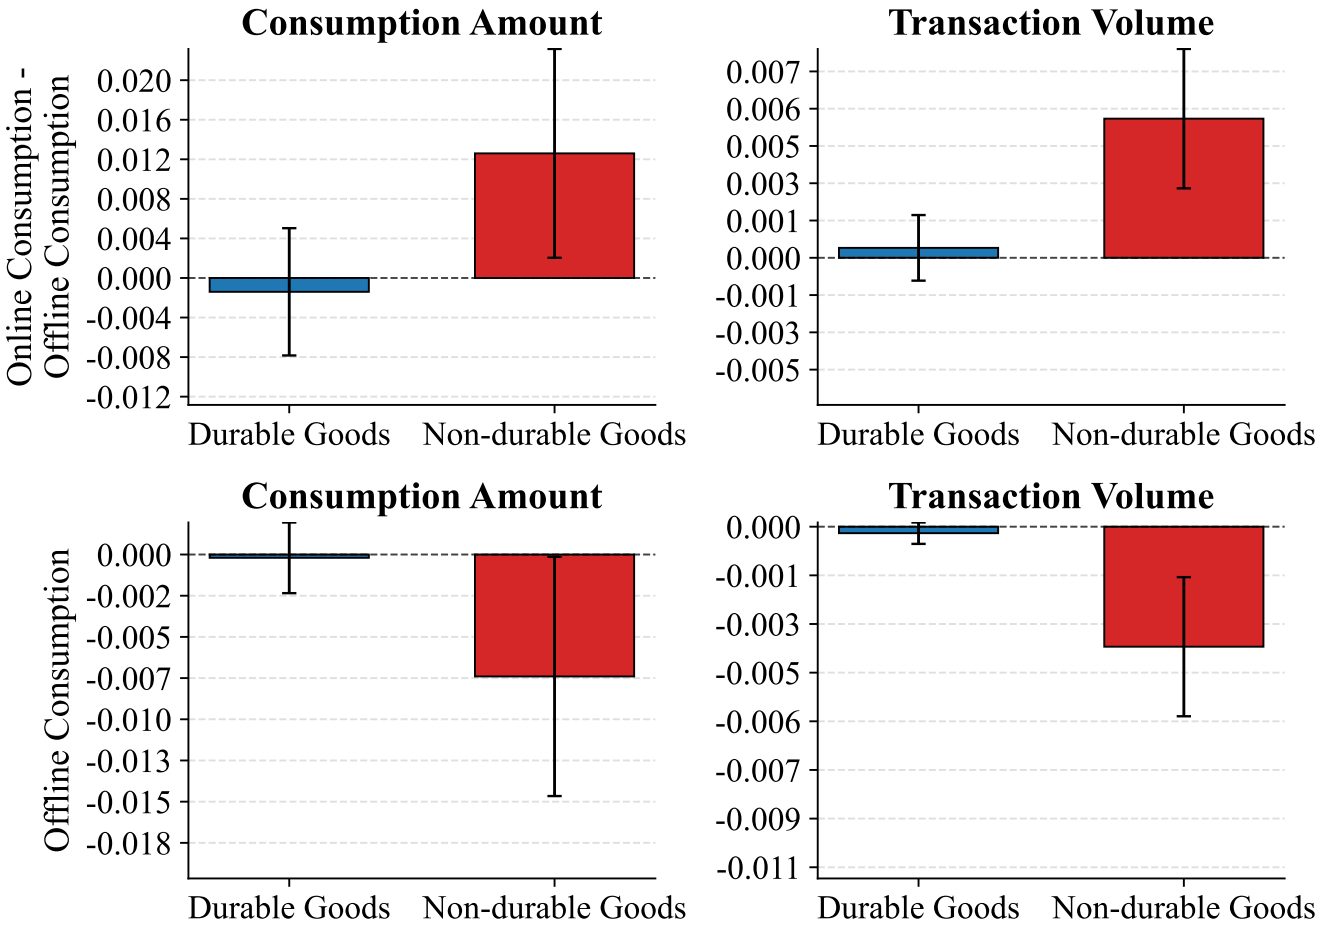


Figure 3: Product-Level Heterogeneity: Durables vs. Nondurables

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption by product durability, estimated using ordinary least squares regression (Equation 1). The dependent variable is the natural logarithm of consumption amount in the left panels and the natural logarithm of transaction volume in the right panels. In the top panels, the bars corresponding to "Durable Goods" and "Non-durable Goods" show the impact of sentiment on the difference between offline and online consumption of durables and non-durables, respectively. In the bottom panels, the bars corresponding to "Durable Goods" and "Non-durable Goods" show the impact of sentiment on the offline consumption of durables and non-durables, respectively. Durable goods include the house and electronic items, while non-durable goods include daily necessities, outdoor activities, clothing, and pets. All estimates are reported with 95% confidence intervals.

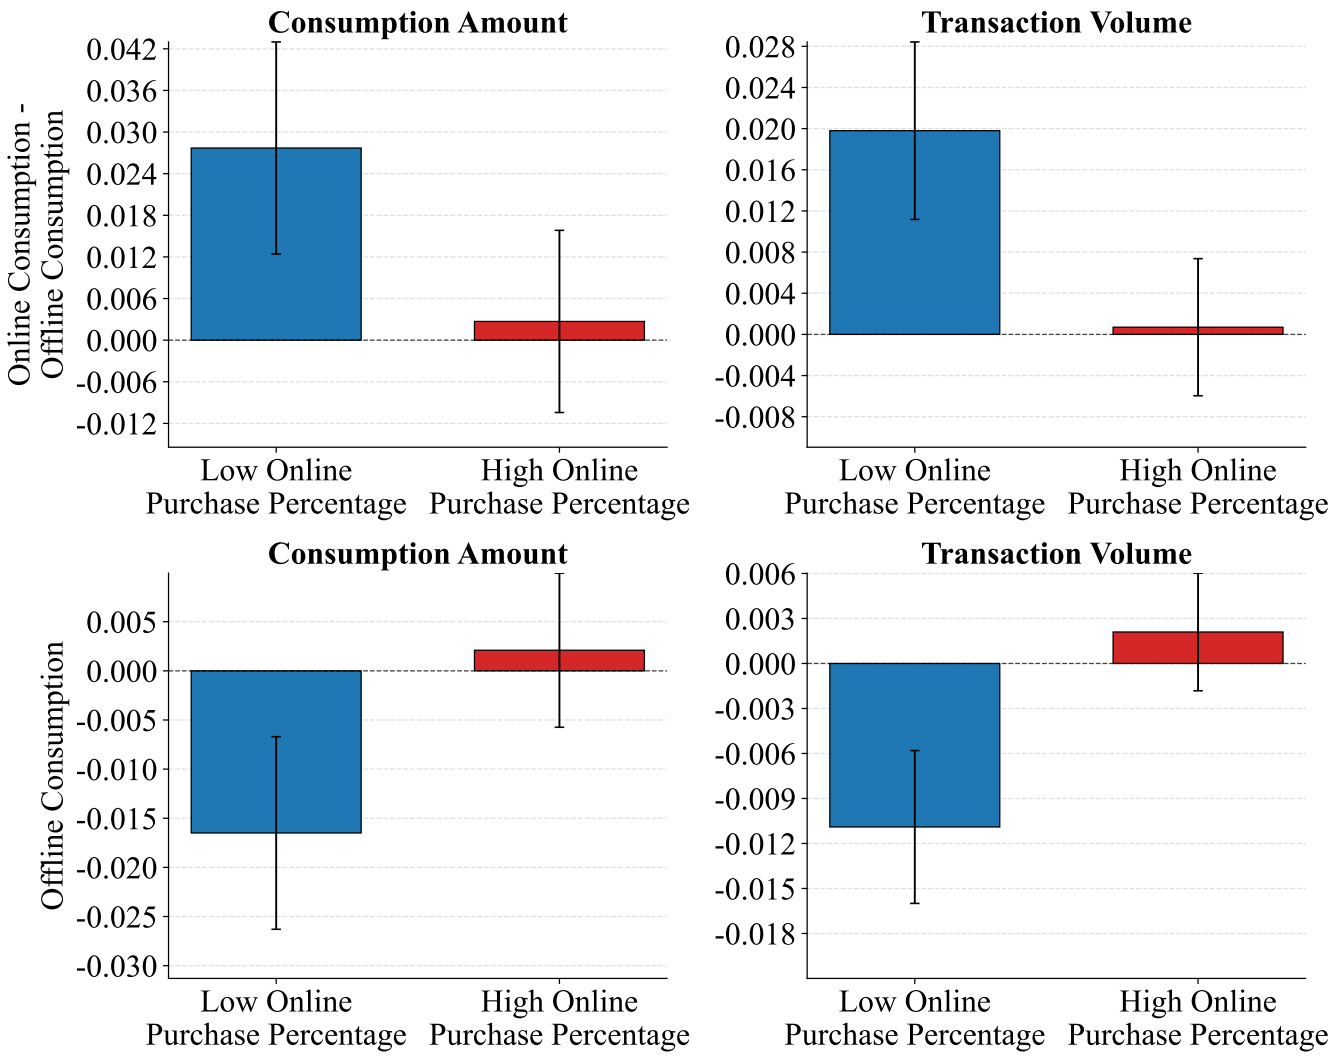


Figure 4: Impacts of Sentiment on Consumption under Different Online Consumption Percentage

Notes: This figure presents the heterogeneous effects of sentiment on consumption across different categories of consumption goods, estimated by ordinary least squares regression. Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment on consumption of product categories with lower percentages of online consumption relative to overall consumption, and the red bar on the right shows the effect of sentiment on consumption of product categories with higher percentages of online consumption relative to overall consumption. Error bars denote 95% confidence intervals.

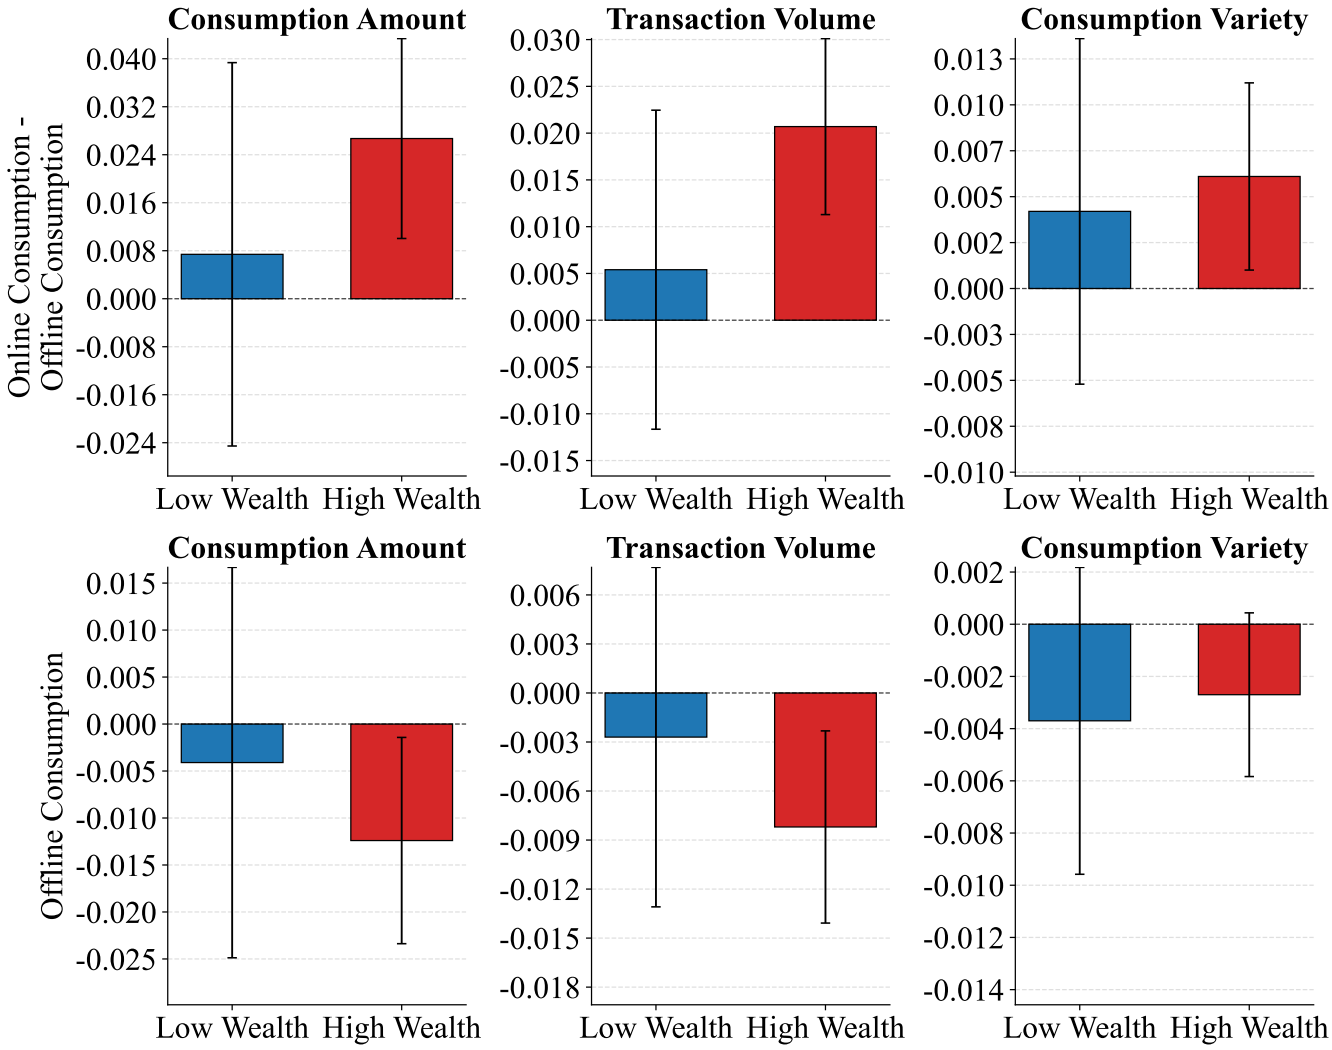


Figure 5: Consumer-Level Heterogeneity by Wealth

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across consumer wealth levels, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left reports the effect of sentiment for low-wealth consumers, while the red bar on the right reports the effect of sentiment for high-wealth consumers. Consumer wealth is a proprietary metric estimated by Alipay using an internal algorithm that synthesizes a comprehensive set of consumer data. Error bars represent 95% confidence intervals.

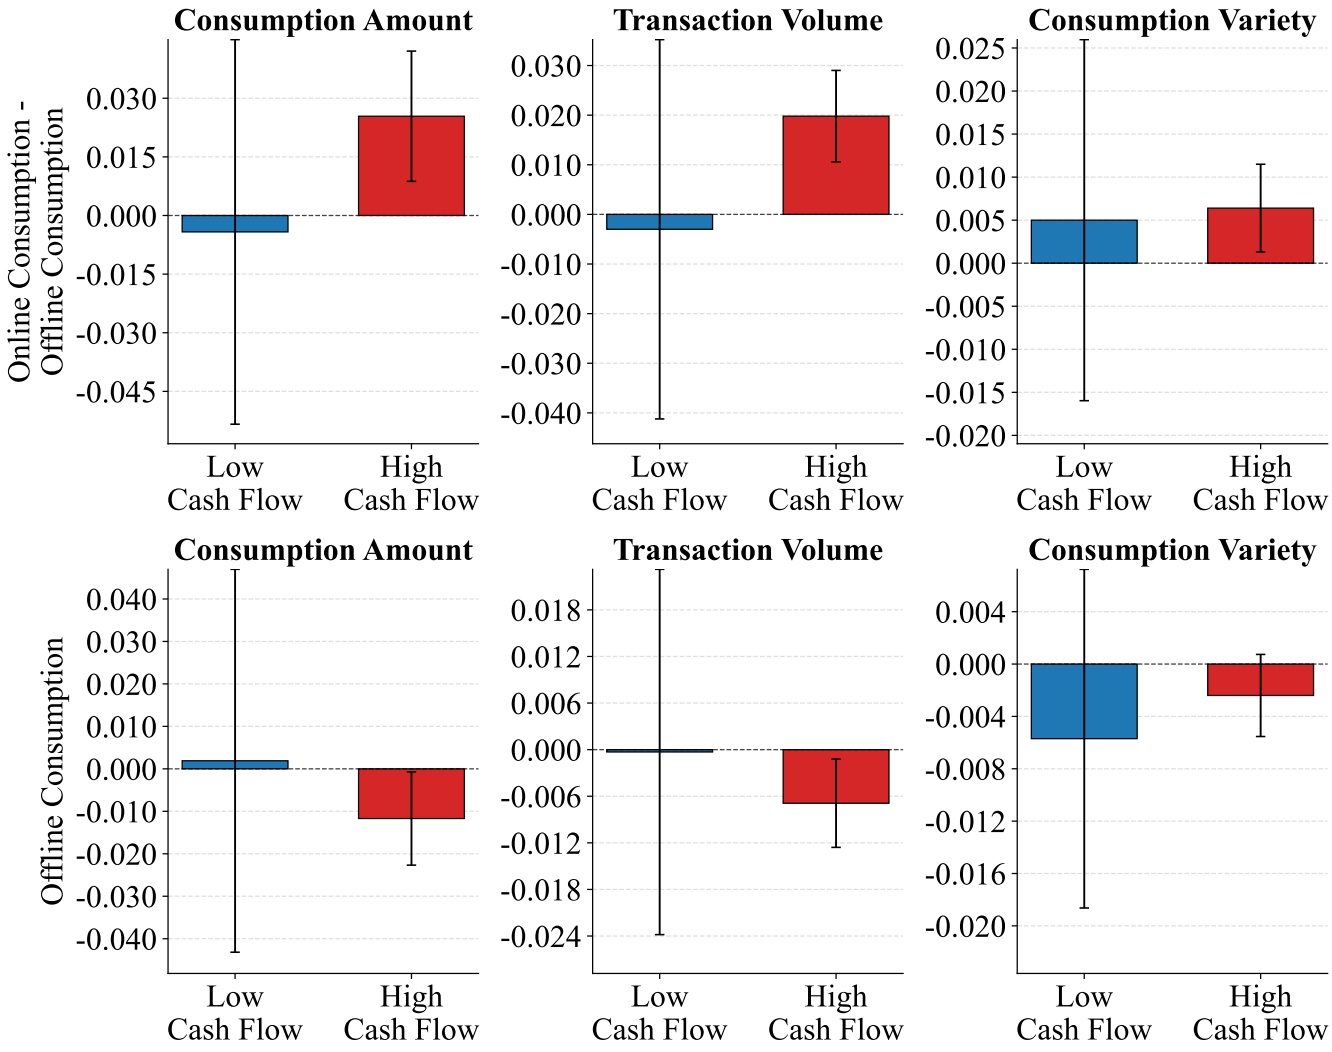


Figure 6: Consumer-Level Heterogeneity by Cash Flow

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across consumer cash flow levels, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left reports the effect of sentiment for low-cash-flow consumers, while the red bar on the right reports the effect of sentiment for high-cash-flow consumers. Cash flow captures the aggregate funds transferred into Alipay accounts by consumers throughout 2018 (the year preceding the sample period). Error bars represent 95% confidence intervals.

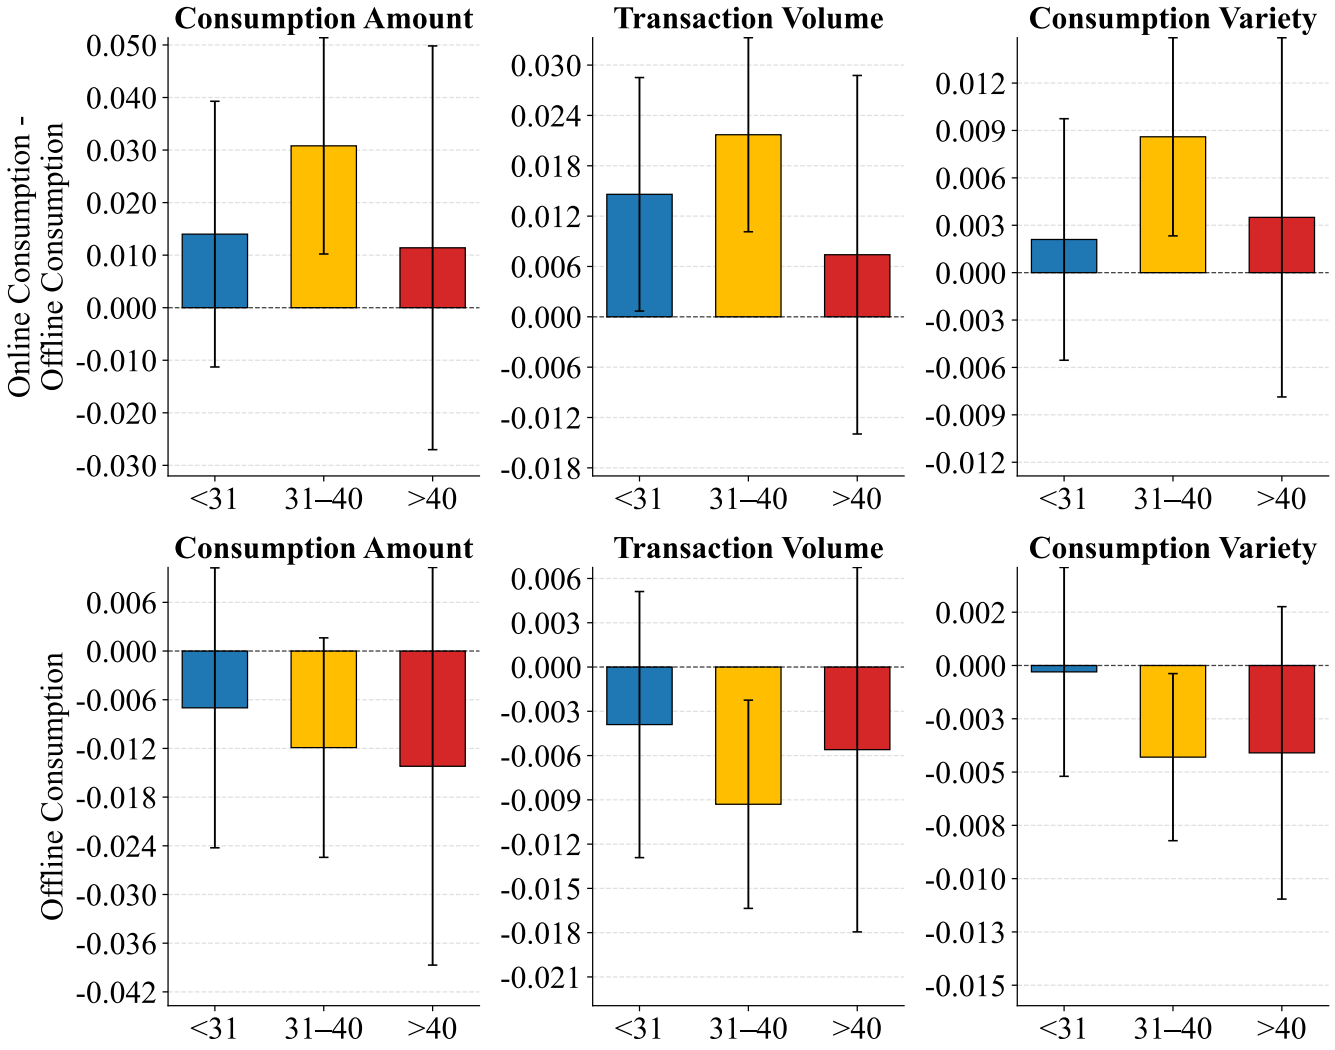


Figure 7: Consumer-Level Heterogeneity by Age

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across consumer age, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left reports the effect of sentiment for consumers under 31, the yellow bar in the middle reports the effect of sentiment for consumers between 31 and 40, and the red bar on the right reports the effect of sentiment for consumers over 40. Error bars represent 95% confidence intervals.

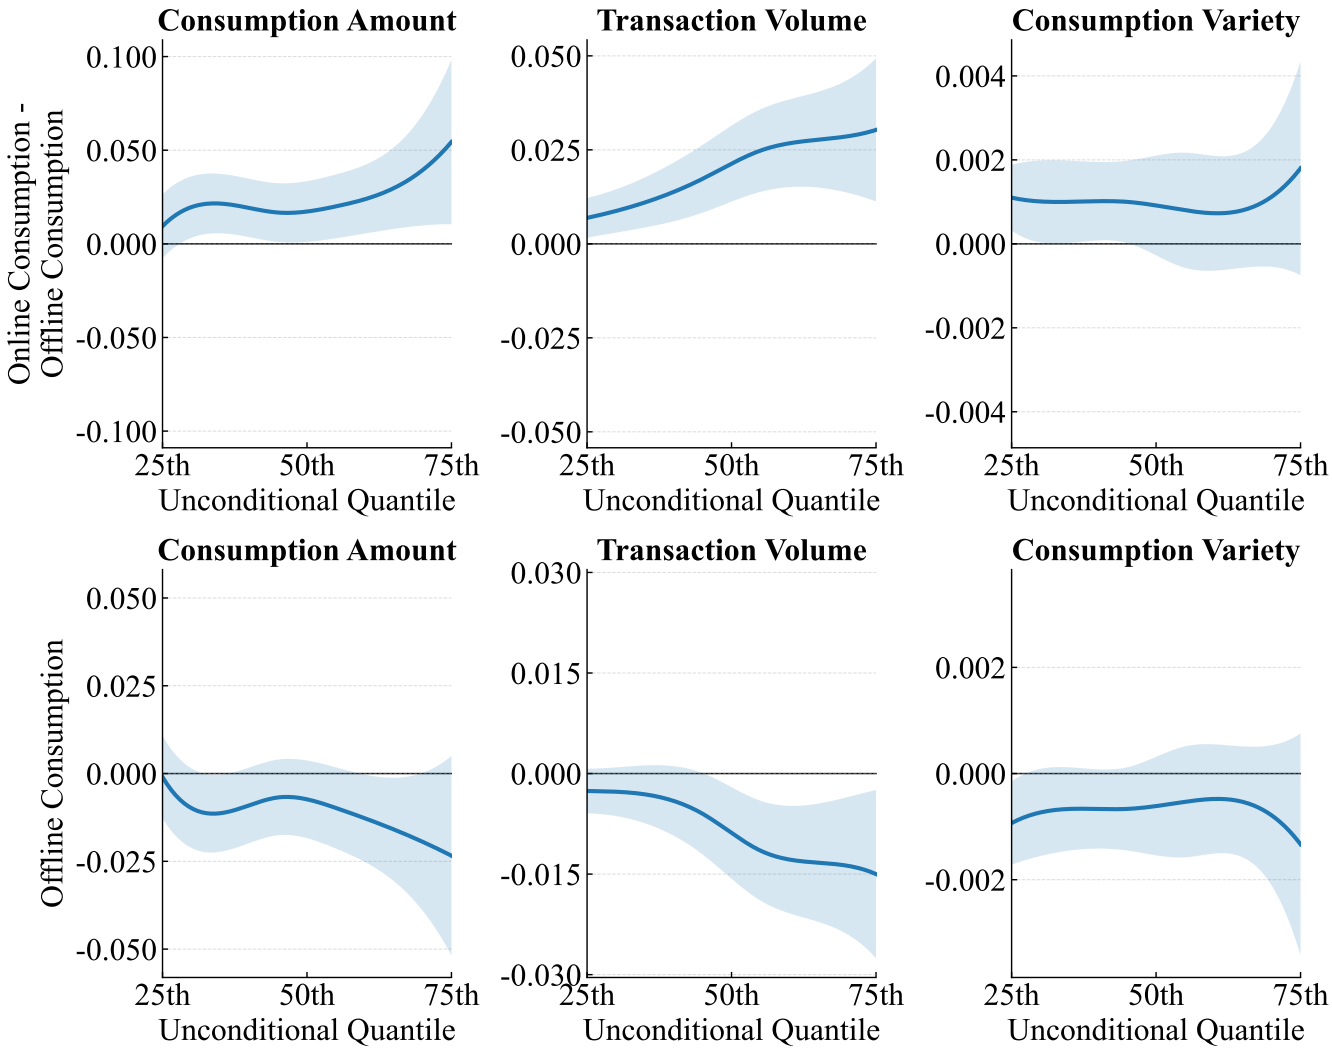


Figure 8: Sentiment Effects Across the Consumption Distribution: Unconditional Quantile Regression Estimates

Notes: This figure plots the effects of sentiment along the unconditional distribution of consumption, estimated using an unconditional quantile regression model following the methodology of Firpo, Fortin, and Lemieux (2009) and Borgen (2016). Across the three panels in each row, the dependent variables are the recentered influence function evaluated at different unconditional quantiles of the natural logarithm of consumption amount, transaction volume, and consumption variety, respectively. The top panels show the impact of sentiment on the difference between offline and online consumption, while the bottom panels show the impact of sentiment on offline consumption. In each panel, point estimates represent the marginal effect of sentiment on consumption at its 25th, 35th, 45th, 55th, 65th, and 75th unconditional quantiles, along with their corresponding 95% confidence intervals.

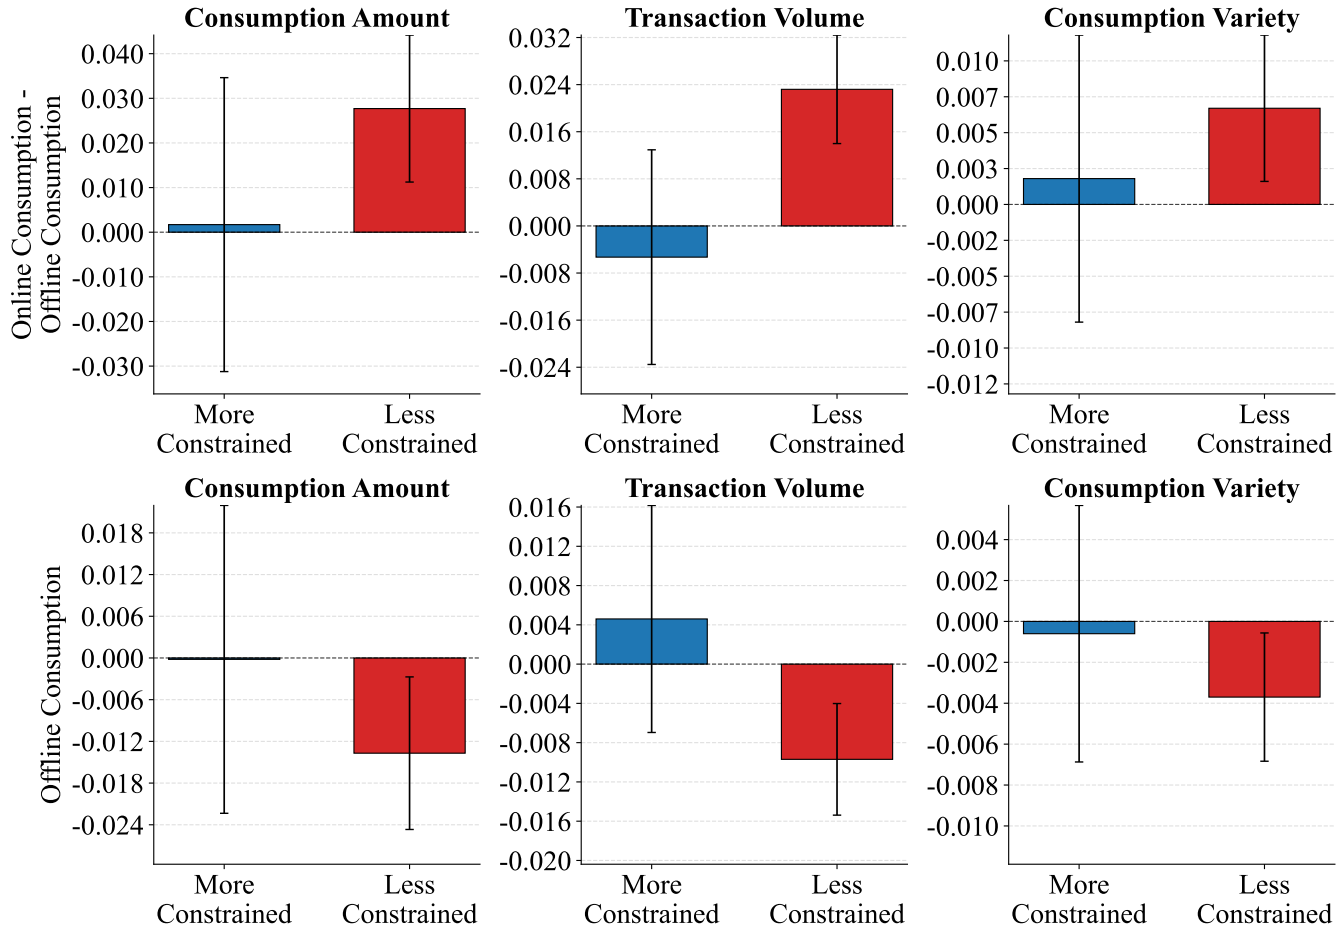


Figure 9: City-Level Heterogeneity by Traditional Credit Constraints

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across credit constraints consumers face when borrowing from traditional financial institutions, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment among consumers facing higher credit constraints from traditional financial institutions, and the red bar on the right shows the effect of sentiment among those facing lower credit constraints. Credit constraints are proxied by the number of bank branches in the consumer's city of residence, with the sample partitioned at the median. Error bars represent 95% confidence intervals.

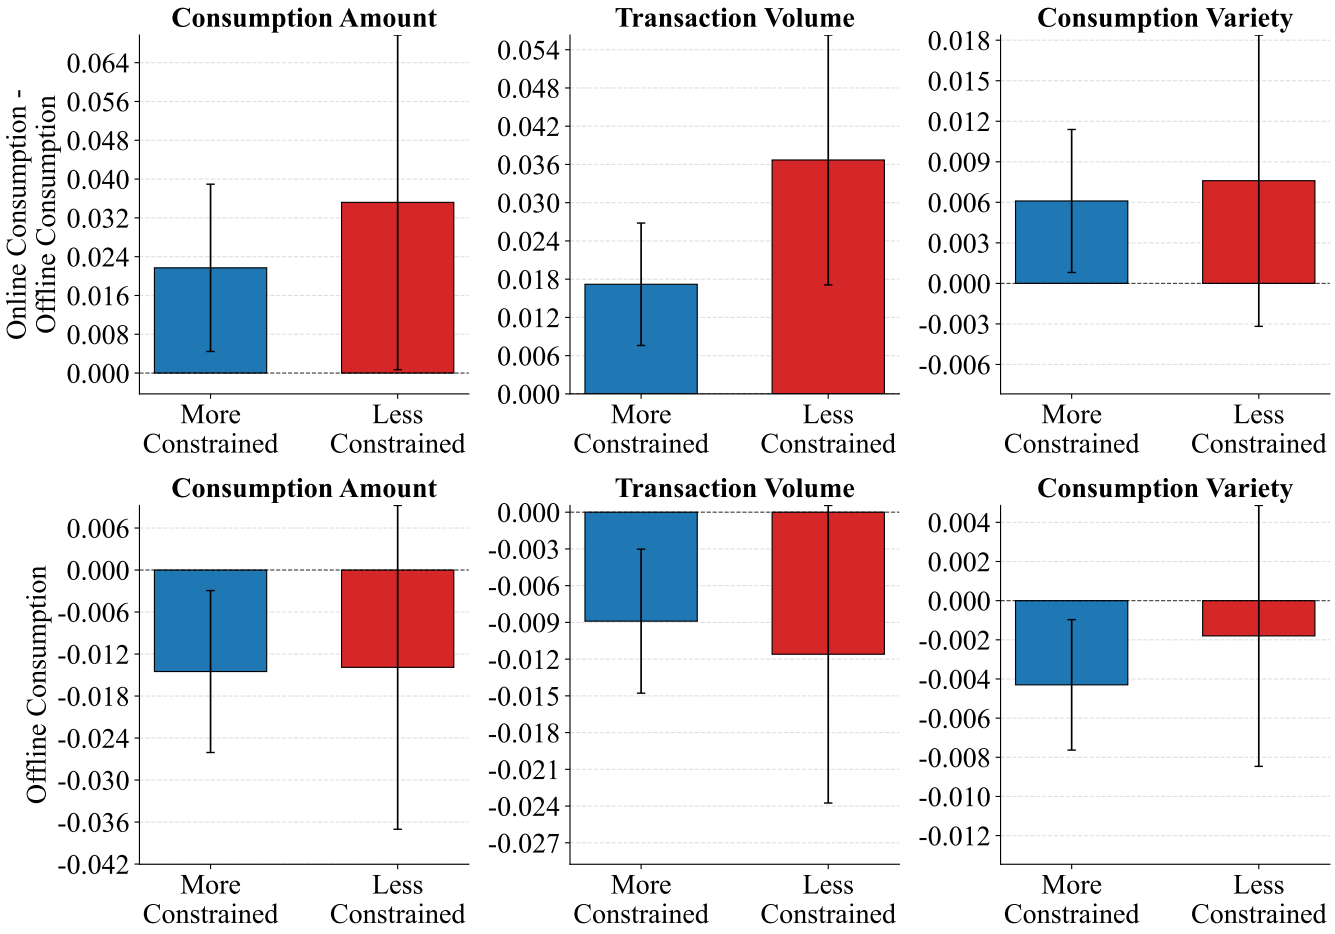


Figure 10: City-Level Heterogeneity by Digital Credit Constraints

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across credit constraints consumers face when borrowing from digital financial institutions, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment among consumers facing higher credit constraints from digital financial institutions, and the red bar on the right shows the effect of sentiment among those facing lower credit constraints. Credit constraints are proxied by a sub-index related to the credit of the digital inclusive finance index developed by Guo et al. (2020), with the sample partitioned at the median. Error bars represent 95% confidence intervals.

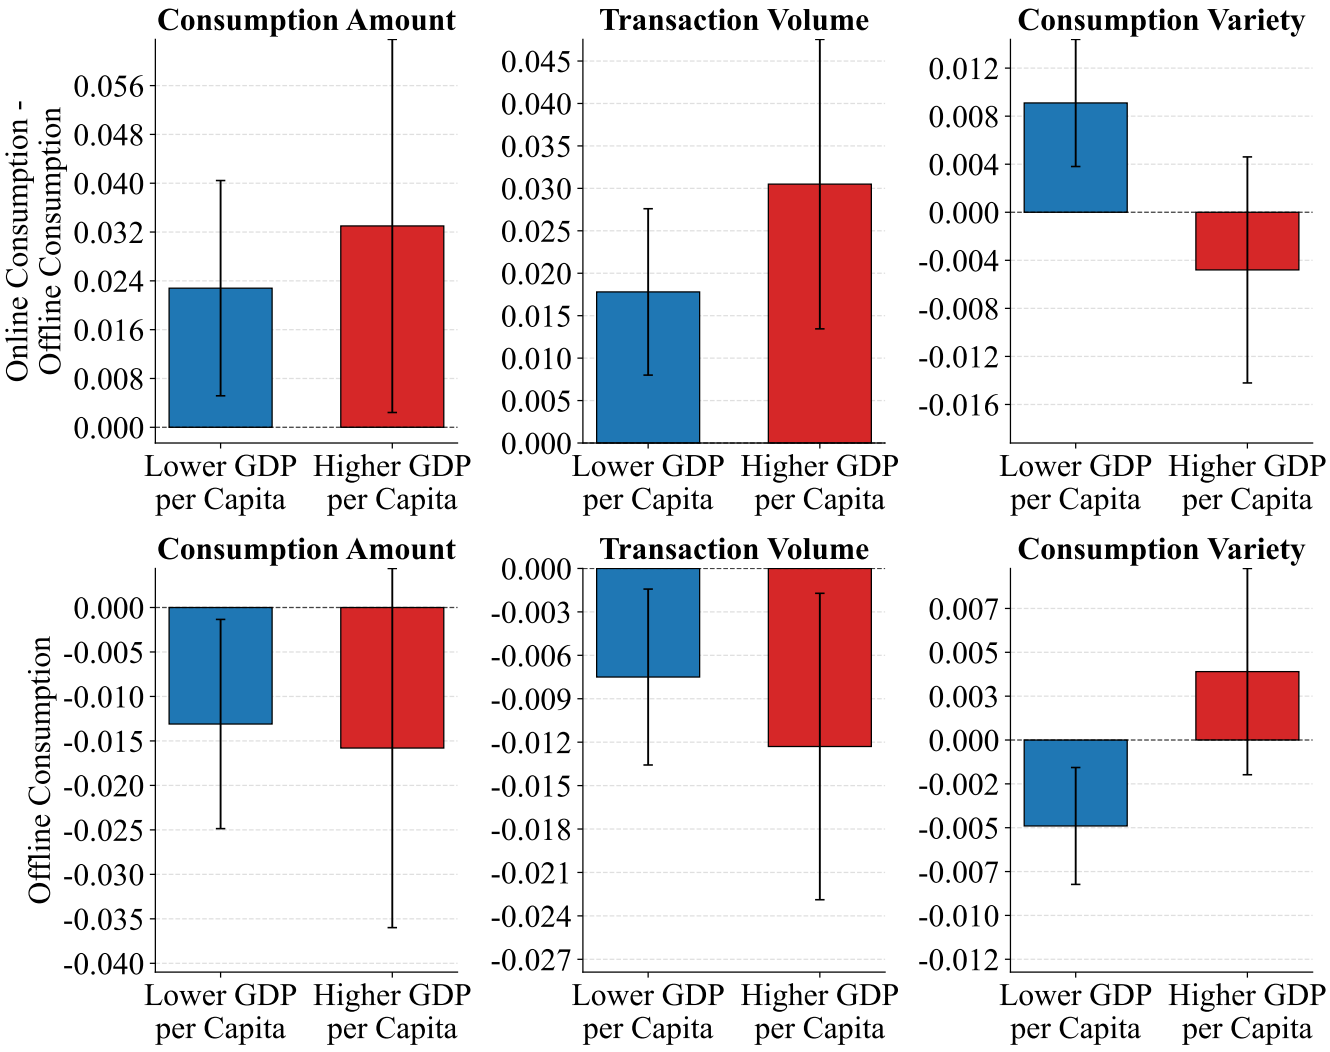


Figure 11: City-Level Heterogeneity by Local Economic Development

Notes: This figure illustrates the heterogeneous effects of sentiment on consumption across local economic development, estimated using ordinary least squares regression (Equation 1). Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment among consumers living in cities with lower local economic development, and the red bar on the right shows the effect of sentiment among those living in cities with higher local economic development. Local economic development is proxied by city-level GDP per capita, with the sample partitioned at the median. Error bars represent 95% confidence intervals.

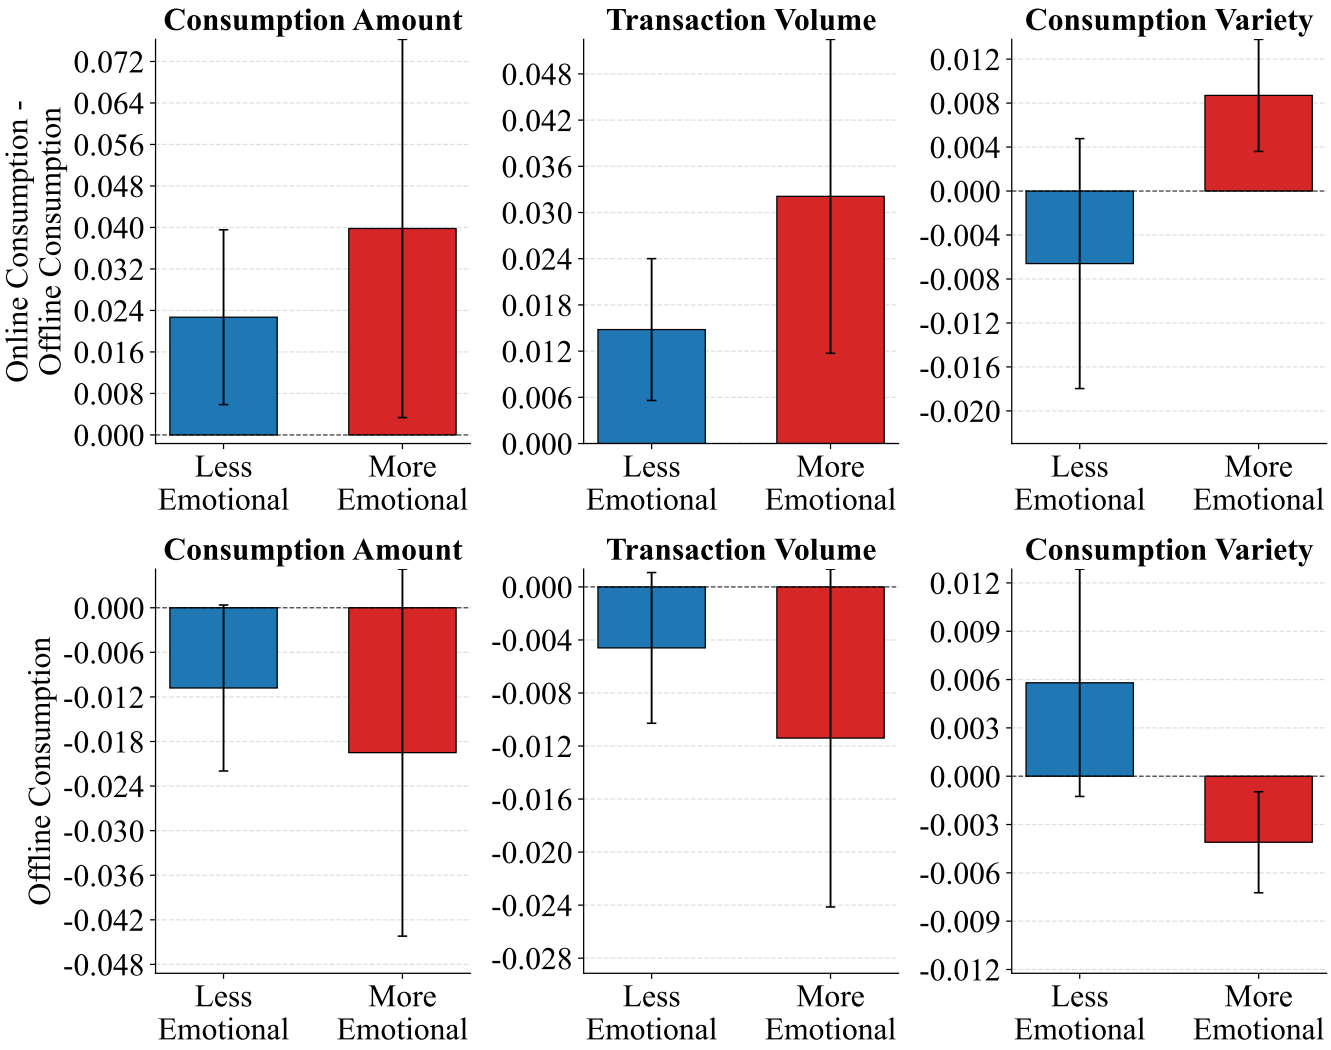


Figure 12: Heterogeneity by Emotionality: Public Attention to a Patriotism Campaign

Notes: This figure presents heterogeneity in the effects of sentiment on consumption by emotionality, estimated using ordinary least squares regressions (Equation 1). Emotionality is proxied by local public attention to a patriotic campaign, measured by the weekly Baidu Search Index for the term "1.4 Billion Flag Guardians" in the consumer's city of residence. Higher public attention to this campaign corresponds to periods of heightened consumer emotionality. Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment among less emotional consumers, and the red bar on the right shows the effect of sentiment among more emotional consumers. Error bars denote 95% confidence intervals.

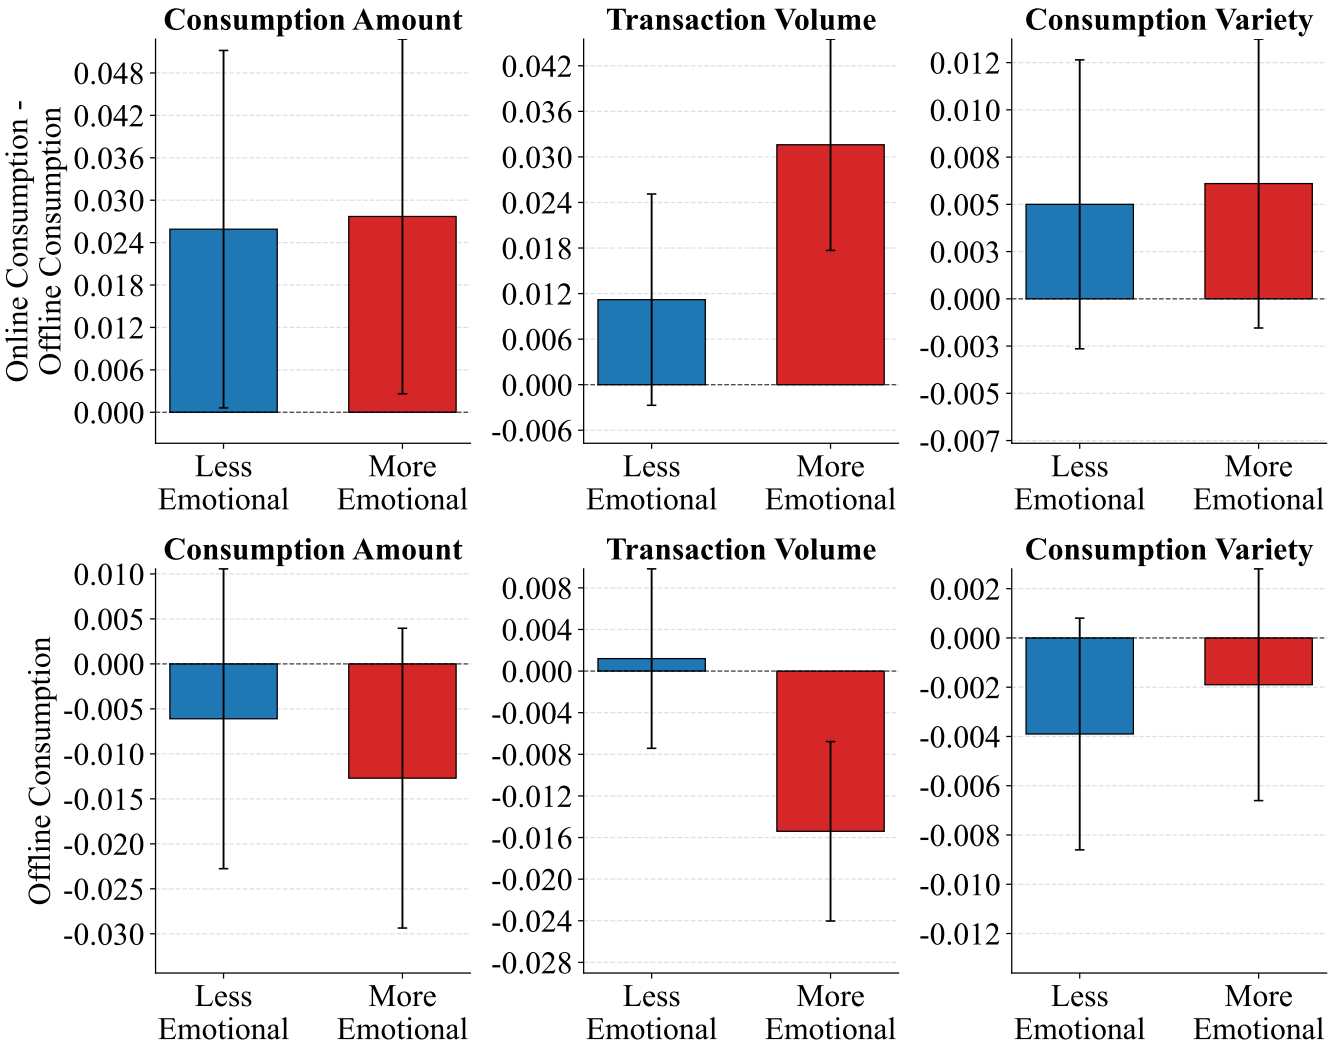


Figure 13: Heterogeneity by Emotionality: Public Attention to a Large Consumption Fraud

Notes: This figure presents heterogeneity in the effects of sentiment on consumption by emotionality, estimated using ordinary least squares regressions (Equation 1). Emotionality is proxied by local public attention to a major consumer fraud scandal, measured by the weekly Baidu Search Index for the term "Live Streaming False Propaganda" in the consumer's city of residence. Elevated public attention to this event corresponds to periods of dampened consumer emotionality. Across the three panels in each row, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety, respectively. The bars in the top panels show the impact of sentiment on the difference between offline and online consumption, while the bars in the bottom panels show the impact of sentiment on offline consumption. Within each panel, the blue bar on the left shows the effect of sentiment among less emotional consumers, and the red bar on the right shows the effect of sentiment among more emotional consumers. Error bars denote 95% confidence intervals.

Table 1: Level Effects of Sentiment on Consumption

Panel A: Overall Consumption			
	$Log(Amount_{i,c,t})$	$Log(Transactions_{i,c,t})$	$Log(Variety_{i,c,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0013	0.0006	-0.0006
	(0.0034)	(0.0017)	(0.0008)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.6246	0.7196	0.4678
N	3681478	3681478	3681478
Panel B: Online and Offline Consumption			
	$Log(Amount_{i,c,f,t})$	$Log(Transactions_{i,c,f,t})$	$Log(Variety_{i,c,f,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0106**	-0.0070***	-0.0030**
	(0.0050)	(0.0026)	(0.0014)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0222***	0.0171***	0.0056**
	(0.0075)	(0.0042)	(0.0023)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.4986	0.4856	0.4635
N	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables in columns (1)–(3) are the natural logarithms of total consumption amount, transaction volume, and consumption variety in week t, respectively. In Panel B, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table 2: Impacts of Sentiment Changes on Consumption Growth

Panel A: Overall Consumption			
	$\Delta Log(Amount_{i,c,t})$	$\Delta Log(Transactions_{i,c,t})$	$\Delta Log(Variety_{i,c,t})$
	(1)	(2)	(3)
$\Delta Sentiment_{c,t-1}$	0.0006	0.0012	0.0005
	(0.0021)	(0.0010)	(0.0005)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.1475	0.1732	0.1110
N	3585708	3585708	3585708
Panel B: Online and Offline Consumption			
	$\Delta Log(Amount_{i,c,f,t})$	$\Delta Log(Transactions_{i,c,f,t})$	$\Delta Log(Variety_{i,c,f,t})$
	(1)	(2)	(3)
$\Delta Sentiment_{c,t-1}$	-0.0042	-0.0031***	-0.0018**
	(0.0026)	(0.0012)	(0.0007)
$\Delta Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0078**	0.0071***	0.0034***
	(0.0036)	(0.0017)	(0.0011)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.4346	0.4202	0.3959
N	6419900	6419900	6419900

Notes: This table reports the effects of sentiment changes on consumption growth. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables in columns (1)–(3) are the changes in the natural logarithms of total consumption amount, transaction volume, and consumption variety, respectively, from week $t-1$ to week t. In Panel B, the dependent variables are the changes in the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f from week $t-1$ to week t, where $f=0$ denotes the offline channel and $f=1$ denotes the online channel. The explanatory variable, $\Delta Sentiment_{c,t-1}$, represents the change in sentiment level in city c from week $t-2$ to week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table 3: Asymmetric Impacts of Sentiment on Consumption

	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	Low Sentiment (1)	High Sentiment (2)	Low Sentiment (3)	High Sentiment (4)	Low Sentiment (5)	High Sentiment (6)
$Sentiment_{c,t-1}$	-0.0142**	-0.0022	-0.0078***	-0.0028	-0.0047***	0.0030
	(0.0058)	(0.0097)	(0.0030)	(0.0050)	(0.0017)	(0.0028)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0261***	0.0106	0.0177***	0.0153*	0.0083***	-0.0024
	(0.0088)	(0.0145)	(0.0049)	(0.0081)	(0.0027)	(0.0045)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.5146	0.4813	0.4961	0.4744	0.4808	0.4450
N	3673668	3689288	3673668	3689288	3673668	3689288

Notes: This table reports estimates of Equation 1. The dependent variables are the natural logarithms of consumption amount (Columns 1–2), transaction volume (Columns 3–4), and consumption variety (Columns 5–6) by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. Columns (1), (3), and (5) report estimates for the subsample where $Sentiment_{c,t-1}$ is below the sample median, while Columns (2), (4), and (6) report estimates for the subsample where $Sentiment_{c,t-1}$ is above the sample median. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table 4: Impacts of Sentiment on Consumption Volatility

Panel A: Overall Consumption Volatility			
	$Amount\ SD_{i,c,t}$	$Transactions\ SD_{i,c,t}$	$Variety\ SD_{i,c,t}$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	0.0021*	0.0060*	-0.0003**
	(0.0011)	(0.0036)	(0.0001)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.8437	0.8903	0.8772
N	2959258	2959258	2959258
Panel B: Online and Offline Consumption Volatility			
	$Amount\ SD_{i,c,f,t}$	$Transactions\ SD_{i,c,f,t}$	$Variety\ SD_{i,c,f,t}$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0024*	-0.0076***	-0.0017***
	(0.0014)	(0.0024)	(0.0002)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0071***	0.0199***	0.0028***
	(0.0019)	(0.0031)	(0.0003)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.8443	0.9257	0.8670
N	6522014	6522014	6522014

Notes: This table reports the impact of sentiment on consumption volatility, estimated using ordinary least squares regression. In Panel A, the dependent variables in Columns (1)–(3) are the volatilities of total consumption amount, transaction volume, and consumption variety, respectively. In Panel B, the dependent variables in Columns (1)–(3) are the volatilities of consumption amount, transaction volume, and consumption variety, respectively, by channel f, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. Volatility is computed as the standard deviation of weekly consumption over the horizon from week t to $t+8$; for consumption amount and variety, the standard deviation is normalized by the simple average. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table 5: City-Level Heterogeneity by Sentiment Volatility

	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	Low Volatility (1)	High Volatility (2)	Low Volatility (3)	High Volatility (4)	Low Volatility (5)	High Volatility (6)
$Sentiment_{c,t-1}$	-0.0100	-0.0116*	-0.0078**	-0.0066*	-0.0009	-0.0048**
	(0.0074)	(0.0068)	(0.0038)	(0.0036)	(0.0021)	(0.0020)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0187*	0.0253**	0.0126**	0.0212***	0.0001	0.0106***
	(0.0110)	(0.0103)	(0.0061)	(0.0057)	(0.0033)	(0.0032)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.4913	0.5050	0.4854	0.4856	0.4565	0.4696
N	3578922	3784034	3578922	3784034	3578922	3784034

Notes: This table reports the heterogeneous effects of sentiment on consumption across cities with varying levels of sentiment volatility, based on estimates of Equation (1). The dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. Columns (1), (3), and (5) report estimates for the subsample of cities with sentiment volatility below the sample median; Columns (2), (4), and (6) report estimates for cities with sentiment volatility above the sample median. City-level sentiment volatility is calculated as the standard deviation of weekly sentiment levels within each city over the 52 weeks of 2019. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table 6: City-Level Heterogeneity by Offline Consumption Opportunities

Panel A: Less Offline Consumption Opportunities

	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	-0.0093	-0.0093	-0.0027	-0.0025	-0.0036**	-0.0034**
	(0.0057)	(0.0057)	(0.0029)	(0.0030)	(0.0016)	(0.0016)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0177**	0.0167*	0.0106**	0.0097**	0.0071***	0.0067**
	(0.0085)	(0.0086)	(0.0047)	(0.0048)	(0.0026)	(0.0026)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.5147	0.5156	0.4925	0.4926	0.4799	0.4808
N	3712946	3583072	3712946	3583072	3712946	3583072

Panel B: More Offline Consumption Opportunities

	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	-0.0251*	-0.0257**	-0.0269***	-0.0274***	0.0040	0.0013
	(0.0129)	(0.0123)	(0.0067)	(0.0064)	(0.0037)	(0.0036)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0635***	0.0651***	0.0683***	0.0681***	0.0016	0.0040
	(0.0195)	(0.0186)	(0.0110)	(0.0105)	(0.0060)	(0.0058)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.4737	0.4742	0.4739	0.4744	0.4385	0.4389
N	3381776	3511650	3381776	3511650	3381776	3511650

Notes: This table reports the heterogeneous effects of sentiment on consumption across regions with varying offline consumption opportunities, based on estimates of Equation (1). Panel A reports estimates for the subsample of consumers in regions with fewer offline consumption opportunities, while Panel B reports estimates for the subsample in regions with more offline consumption opportunities. In each panel, Columns (1), (3), and (5) are based on the density of consumption-related points of interest (POIs), while Columns (2), (4), and (6) are based on the density of shopping malls (both measured as of 2019). The dependent variables are the natural logarithms of consumption amount (Columns 1–2), transaction volume (Columns 3–4), and consumption variety (Columns 5–6) by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

References

Acemoglu, D. and Scott, A. (1994). Consumer confidence and rational expectations: Are agents' beliefs consistent with the theory? *The Economic Journal*, 104(422):1–19.

Angeletos, G., Collard, F., and Dellas, H. (2018). Quantifying confidence. *Econometrica*, 86(5):1689–1726.

Angeletos, G. and La'O, J. (2013). Sentiments. *Econometrica*, 81(2):739–779.

Angeletos, G.-M. and La'O, J. (2010). Noisy business cycles. *NBER Macroeconomics Annual*, 24:319–378.

Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review*, 102(4):1343–1377.

Benhabib, J., Liu, X., and Wang, P. (2016). Endogenous information acquisition and countercyclical uncertainty. *Journal of Economic Theory*, 165:601–642.

Benhabib, J. and Spiegel, M. M. (2018). Sentiments and economic activity: Evidence from us states. *The Economic Journal*, 129(618):715–733.

Benhabib, J., Wang, P., and Wen, Y. (2015). Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(Jan):993–1022.

Borgen, N. T. (2016). Fixed effects in unconditional quantile regression. *The Stata Journal*, 16(2):403–415.

Bounie, D., Camara, Y., and Galbraith, J. W. (2025). Consumer mobility, offline and online regional trade: Evidence from high-frequency transaction data. *International Economic Review*.

Brodeur, A. (2018). The effect of terrorism on employment and consumer sentiment: Evidence from successful and failed terror attacks. *American Economic Journal: Applied Economics*, 10(4):246–282.

Carroll, C. D., Fuhrer, J. C., and Wilcox, D. W. (1994). Does consumer sentiment forecast household spending? if so, why? *The American Economic Review*, 84(5):1397–1408.

Cavallo, A. (2017). Are online and offline prices similar? evidence from large multi-channel retailers. *American Economic Review*, 107(1):283–303.

Cavallo, A. (2018). More amazon effects: Online competition and pricing behaviors. Working Paper 25138, National Bureau of Economic Research.

Chang, A., Dong, X., Martin, X., and Zhou, C. (2023). Ai democratization, return predictability, and trading inequality. Working Paper 4543999, SSRN.

Chen, H., Qian, W., and Wen, Q. (2021). The impact of the covid-19 pandemic on consumption: Learning from high-frequency transaction data. *AEA Papers and Proceedings*, 111:307–11.

Chen, H., Qian, W., and Wen, Q. (2023). Singles Day and Consumer Spending Offline. Working Paper 4644663, SSRN.

Chhaochharia, V., Kim, D., Korniotis, G. M., and Kumar, A. (2019). Mood, firm behavior, and aggregate economic outcomes. *Journal of Financial Economics*, 132(2):427–450.

Cortés, K., Duchin, R., and Sosyura, D. (2016). Clouded judgment: The role of sentiment in credit origination. *Journal of Financial Economics*, 121(2):392–413.

Couture, V., Faber, B., Gu, Y., and Liu, L. (2021). Connecting the Countryside via E-Commerce: Evidence from China. *American Economic Review: Insights*, 3(1):35–50.

Dautović, E., Hau, H., and Huang, Y. (2024). Consumption response to minimum wages: Evidence from chinese households. *The Review of Economics and Statistics*, pages 1–47.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859.

Dolfen, P., Einav, L., Klenow, P. J., Klopack, B., Levin, J. D., Levin, L., and Best, W. (2023). Assessing the Gains from E-Commerce. *American Economic Journal: Macroeconomics*, 15(1):342–70.

Edmans, A., Garcia, D., and Norli, (2007). Sports sentiment and stock returns. *The Journal of Finance*, 62(4):1967–1998.

Firpo, S., Fortin, N. M., and Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3):953–973.

Fisman, R., Hamao, Y., and Wang, Y. (2014). Nationalism and economic exchange: Evidence from shocks to sino-japanese relations. *The Review of Financial Studies*, 27(9):2626–2660.

Forman, C., Ghose, A., and Goldfarb, A. (2009). Competition between local and electronic markets: How the benefit of buying online depends on where you live. *Management Science*, 55(1):47–57.

Ganong, P., Greig, F., Noel, P., Sullivan, D. M., and Vavra, J. (2024). Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data. *American Economic Review*, 114(9):2898–2939.

Ganong, P., Jones, D., Noel, P., Farrell, D., Greig, F., and Wheat, C. (2025). Liquid wealth and consumption smoothing of typical labor income shocks.

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.

Gao, Z., Xiong, W., and Yuan, J. (2026). Structured Beliefs and Fund Performance: An LLM-Based Approach.

Garmaise, M. J., Levi, Y., and Lustig, H. (2024). Spending less after (seemingly) bad news. *The Journal of Finance*, 79(4):2429–2471.

Gelman, M., Gorodnichenko, Y., Kariv, S., Koustas, D., Shapiro, M. D., Silverman, D., and Tadelis, S. (2023). The response of consumer spending to changes in gasoline prices. *American Economic Journal: Macroeconomics*, 15(2):129–160.

Gillitzer, C. and Prasad, N. (2018). The effect of consumer sentiment on consumption: Cross-sectional evidence from elections. *American Economic Journal: Macroeconomics*, 10(4):234–269.

Goolsbee, A. (2000). In a world without borders: The impact of taxes on internet commerce*. *The Quarterly Journal of Economics*, 115(2):561–576.

Gu, C. and Wang, Y. (2022). Consumer online search with partially revealed information. *Management Science*, 68(6):4215–4235.

Guo, F., Wang, J., Wang, F., Kong, T., Zhang, X., and Cheng, Z. (2020). Measuring China's Digital Financial Inclusion: Index Compilation and Spatial Characteristics. *China Economic Quartely*, 19(4):1401–1418.

Heese, J., Pérez-Cavazos, G., and Pérez-Silva, A. (2024). Human bias in the oversight of firms: evidence from workplace safety violations. *Review of Accounting Studies*, 29(4):3413–3448.

Houde, J.-F., Newberry, P., and Seim, K. (2023). Nexus tax laws and economies of density in e-commerce: A study of amazon's fulfillment center network. *Econometrica*, 91(1):147–190.

Huo, Z. and Takayama, N. (2024). Rational expectations models with higher-order beliefs. *The Review of Economic Studies*, 92(5):3138–3173.

Jappelli, T. and Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, 2(Volume 2, 2010):479–506.

Jha, M., Liu, H., and Manela, A. (2025). Does finance benefit society? a language embedding approach. *The Review of Financial Studies*, page hhaf012.

Jo, Y., Matsumura, M., and Weinstein, D. E. (2024). The impact of retail e-commerce on relative prices and consumer welfare. *The Review of Economics and Statistics*, 106(6):1675–1689.

Jordá, , Schularick, M., and Taylor, A. M. (2016). Sovereigns versus banks: Credit, crises, and consequences. *Journal of the European Economic Association*, 14(1):45–79.

Kirtac, K. and Germano, G. (2025). Large language models in finance : what is financial sentiment?

Lagerborg, A., Pappa, E., and Ravn, M. O. (2022). Sentimental business cycles. *The Review of Economic Studies*, 90(3):1358–1393.

Li, H., Yin, P., and Zhou, Y. (2025). Public welfare-oriented media promotion and corporate donations: Evidence from "touching china". *Journal of Economic Behavior Organization*, 238:107195.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of Finance*, 66(1):35–65.

Ludvigson, S. C. (2004). Consumer confidence and consumer spending. *Journal of Economic Perspectives*, 18(2):29–50.

Malmendier, U. and Shen, L. S. (2024). Scarred consumption. *American Economic Journal: Macroeconomics*, 16(1):322–55.

Maxted, P. (2023). A macro-finance model with sentiment. *The Review of Economic Studies*, 91(1):438–475.

Mojon, B., Qiu, H., Wang, F., and Weber, M. (2025). Housing wealth effects in China. Working Paper 1319, Bank for International Settlements.

Moraga-González, J. L., Sándor, Z., and Wildenbeest, M. R. (2022). Consumer search and prices in the automobile market. *The Review of Economic Studies*, 90(3):1394–1440.

Paiella, M. and Pistaferri, L. (2017). Decomposing the wealth effect on consumption. *Review of Economics and Statistics*, 99(4):710–721.

Shapiro, A. H., Sudhof, M., and Wilson, D. J. (2022). Measuring news sentiment. *Journal of Econometrics*, 228(2):221–243.

Sun, S. T. and Yannelis, C. (2016). Credit constraints and demand for higher education: Evidence from financial deregulation. *The Review of Economics and Statistics*, 98(1):12–24.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, 97(2):305–346.

Òscar Jordà (2005). Estimation and inference of impulse responses by local projections. *The American Economic Review*, 95(1):161–182.

A Data Appendix

A.1 Collecting Consumption-Related Posts

We provide further details of the two-step topic-modeling procedure used to collect posts explicitly conveying consumption-specific sentiment.

First, to mitigate seasonal confounding driven by the Spring Festival (typically occurring in late January to mid-February) and academic vacations in China, we assemble a seed corpus of 100,000 Weibo posts published between March 1 and March 31, 2019. These posts contain at least one of six initial consumption-related keywords: "buy", "purchase", "shopping", "consumption", "cross-border shopping", and "procurement".

We subject this seed corpus to a rigorous natural language preprocessing pipeline. This involves Chinese word segmentation and stop-word removal utilizing established open-source lexicons developed by authoritative academic institutions (Tsinghua University, Harbin Institute of Technology, and Sichuan University) and leading internet companies (Tencent and Baidu). Following preprocessing, we apply a Latent Dirichlet Allocation (LDA) model to the corpus and, in parallel, apply a BERT-based topic model using a pre-trained Chinese encoder (Blei, Ng, and Jordan (2003); Devlin et al. (2019)). We extract the top 50 keywords generated by each topic model. After a manual review of this pooled vocabulary, we retain a refined dictionary of 58 keywords strictly related to consumption activities (detailed in Table A1).

Second, leveraging this refined keyword dictionary, we retrieve all geotagged Weibo posts published between January 1 and December 31, 2019, that contain at least one of the 58 target keywords. A primary empirical concern at this stage is the prevalence of commercial advertisements within the raw data. Because advertisements systematically convey positive tones, their inclusion would introduce a severe upward bias into our constructed consumer sentiment measure. To ensure data integrity, we implement a robust two-step filtering protocol to eliminate commercial noise: (i) we compile another dictionary of common promotional phrases to screen out algorithmic and advertising content, and (ii) we exclude high-frequency accounts, strictly defined as users averaging more than one post per week over the sample period. This rigorous selection procedure yields a final sample of approximately 100,000 validated, geographically tagged posts related to consumption.

To verify that the collected posts indeed pertain to attitudes toward consumption, we estimate a two-topic LDA model on the full sample of retained Weibo posts. Figures A1 and A2 present word clouds generated from the 50 terms with the highest posterior probabilities within each topic, rendered in English and Chinese, respectively. The pronounced concentration of consumption-specific terminology across both topics confirms that the data are strictly anchored to the attitude towards consumption. This establishes robust content validity for our empirical design, ensuring that the subsequent sentiment classification cleanly isolates true individual attitudes toward consumption.

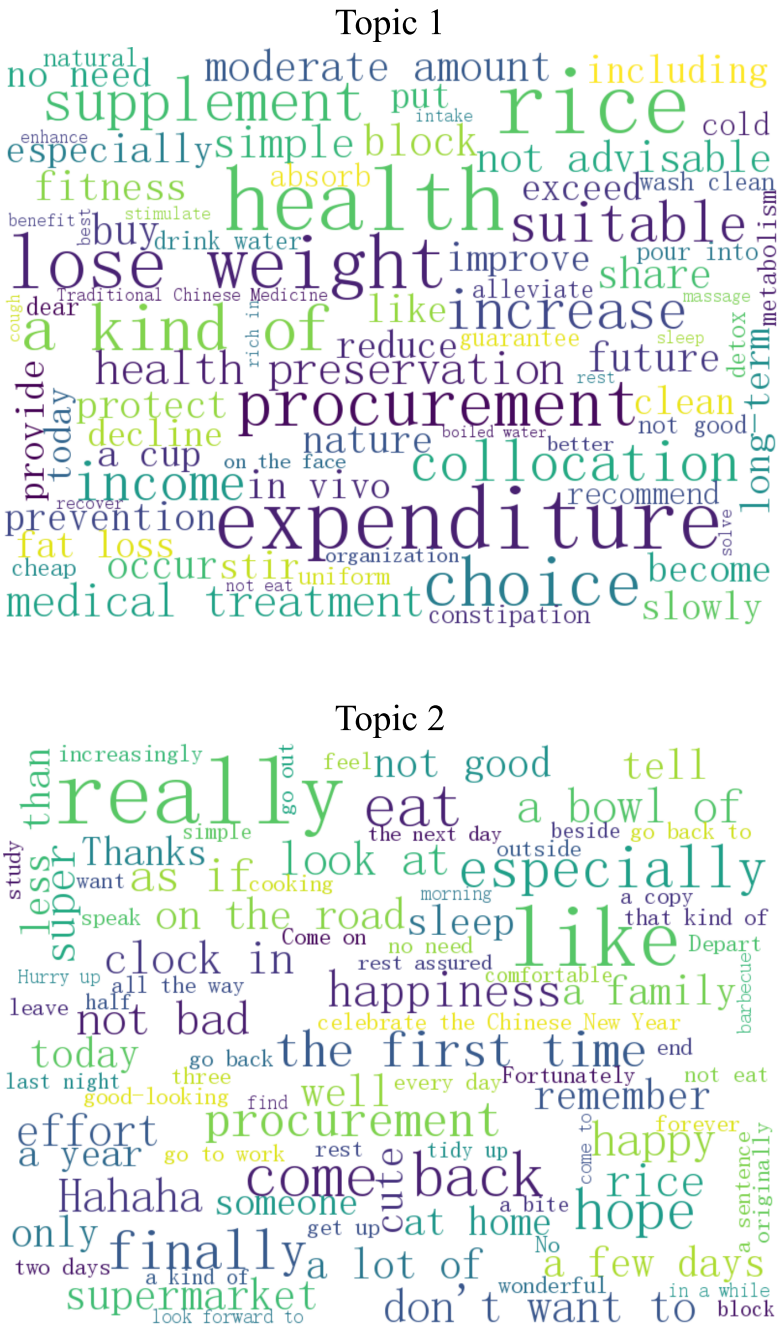


Figure A1: Word Cloud

Notes: This figure illustrates the lexical distribution extracted from the full sample of consumer-generated posts on Weibo, one of China's largest social media platforms. The underlying corpus consists of posts published in 2019 with city-level geotags that have been rigorously screened to exclude commercial advertisements. The visualization is generated via a natural language processing pipeline: (1) corpus preprocessing, including tokenization and the removal of stopwords and common words to standardize the text; (2) topic modeling, which estimates a two-topic Latent Dirichlet Allocation model to cluster the resultant vocabulary.

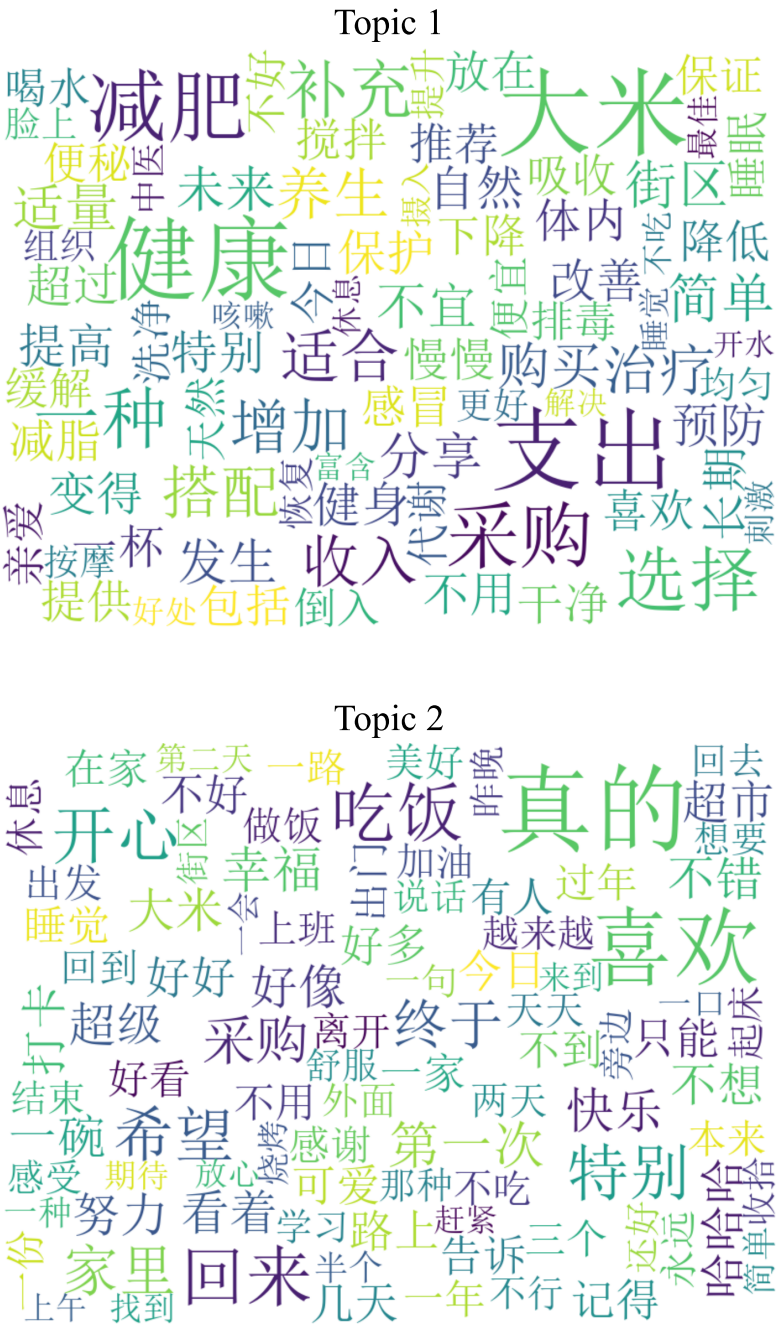


Figure A2: Word Cloud in Chinese

Notes: This figure illustrates the lexical distribution extracted from the full sample of consumer-generated posts on Weibo, one of China's largest social media platforms. The underlying corpus consists of posts published in 2019 with city-level geotags that have been rigorously screened to exclude commercial advertisements. The visualization is generated via a natural language processing pipeline: (1) corpus preprocessing, including tokenization and the removal of stopwords and common words to standardize the text; (2) topic modeling, which estimates a two-topic Latent Dirichlet Allocation model to cluster the resultant vocabulary. All terms are presented in Chinese.

A.2 Labeling the Sentiment of Posts

Traditional sentiment-classification methods typically rely on rigid lexicons or supervised models trained on narrow, domain-specific corpora. In contrast, large language models (LLMs) are pre-trained on vast, open-domain text. This extensive pre-training enables them to parse complex contextual relationships and nuanced emotional valences, a capability that is crucial for accurately identifying subtle, context-dependent shifts in sentiment polarity.

To classify the sentiment of each Weibo post, we utilize OpenAI's GPT-4o-mini API. Because the length of all posts in our sample falls well within the model's context window, we process the full text without truncation. We input the raw text alongside a carefully engineered zero-shot prompt, instructing the model to assign a binary sentiment score (+1 for positive sentiment, -1 for negative sentiment). Crucially, to isolate pure sentiment from fundamental product quality, our prompt simultaneously directs the model to flag posts that discuss product quality. We strictly exclude these posts related to product quality prior to constructing our sentiment measure, ensuring our final measure is strictly purged of underlying quality shocks. Appendix A.6 provides the prompt in detail.

A.3 Measurement of the Consumer Sentiment

Let $N_{c,t}$ denote the number of consumer-generated posts in city c during week t. We construct the city–week measurement of consumer sentiment as

$$Sentiment_{c,t} = \frac{\sum_{k=1}^{N_{c,t}} \{sentiment_{k,c,t}\}}{N_{c,t}}$$

where $sentiment_{k,c,t}$ denotes the sentiment of post k published by a consumer residing in city c in week t, taking the value 1 for positive sentiment and -1 for negative sentiment. As described in section A.2, each post's sentiment is labelled by a large language model.

A.4 Validation of the Sentiment Measurement

We provide several pieces of evidence that our measures are valid for capturing sentiment toward consumption, including comparisons with established macroeconomic sentiment indices and well-known consumption-related measures.

First, we benchmark our consumer sentiment measure against a well-established exogenous proxy for local mood: cloud thickness (Chhaochharia et al. (2019); Heese, Pérez-Cavazos, and Pérez-Silva (2024)). Prior literature documents a robust negative relationship between cloud thickness and individual sentiment. Consequently, a valid sentiment measure should be negatively covary with cloud thickness. To test this, we partition our city-week panel into quintiles based on local cloud thickness and construct corresponding indicator variables. We then regress our consumer sentiment

measure on these five indicators—omitting the constant term to directly recover the group-specific conditional means—while controlling for the identical suite of weather and air quality index used in Equation 1.

The estimated coefficients, plotted in Figure A3, reflect the conditional mean of consumer sentiment within each cloud-thickness quintile. These point estimates strictly decline as cloud thickness increases, confirming the expected negative correlation and validating that our sentiment measure captures genuine sentiment-driven variation. Crucially, the modest R^2 of this regression indicates that cloud thickness explains only a fraction of the total variance in our sentiment measure. This confirms that our text-based sentiment measure encapsulates substantial, independent variation well beyond standard, coarse environmental proxies.

Second, we validate our text-based consumer sentiment measure against two official macroeconomic benchmarks published by the National Bureau of Statistics (NBS) of China: the Consumer Confidence Index and the Consumer Expectation Index. To facilitate this comparison, we aggregate our high-frequency city-week consumer sentiment measure into a national monthly series, calculated as the weighted average of city-level sentiment, with city-level GDP per capita as weights. Figure A4 plots our aggregate sentiment index alongside the official Consumer Confidence Index. The two time series exhibit a pronounced positive comovement, confirming that our text-based measure effectively captures aggregate macroeconomic consumption dynamics. An analogous comparison with the NBS Consumer Expectation Index, presented in Figure A5, reveals a similarly strong positive association. Taken together, this robust macroeconomic alignment provides compelling external validation for our constructed sentiment measure.

Additionally, we leverage our LLM classification framework to categorize each post as pertaining to either online or offline activities, allowing us to compute the city-level share of offline-related posts. Figure A6 plots the cross-city distribution of this share, demonstrating that the overwhelming majority of the collected Weibo posts reflect offline rather than online consumption. This empirical distribution mitigates the concern that our baseline reallocation effects are driven by a compositional artifact—specifically, the threat that the sentiment measure is mechanically dominated by online discourse that disproportionately drives online expenditure.

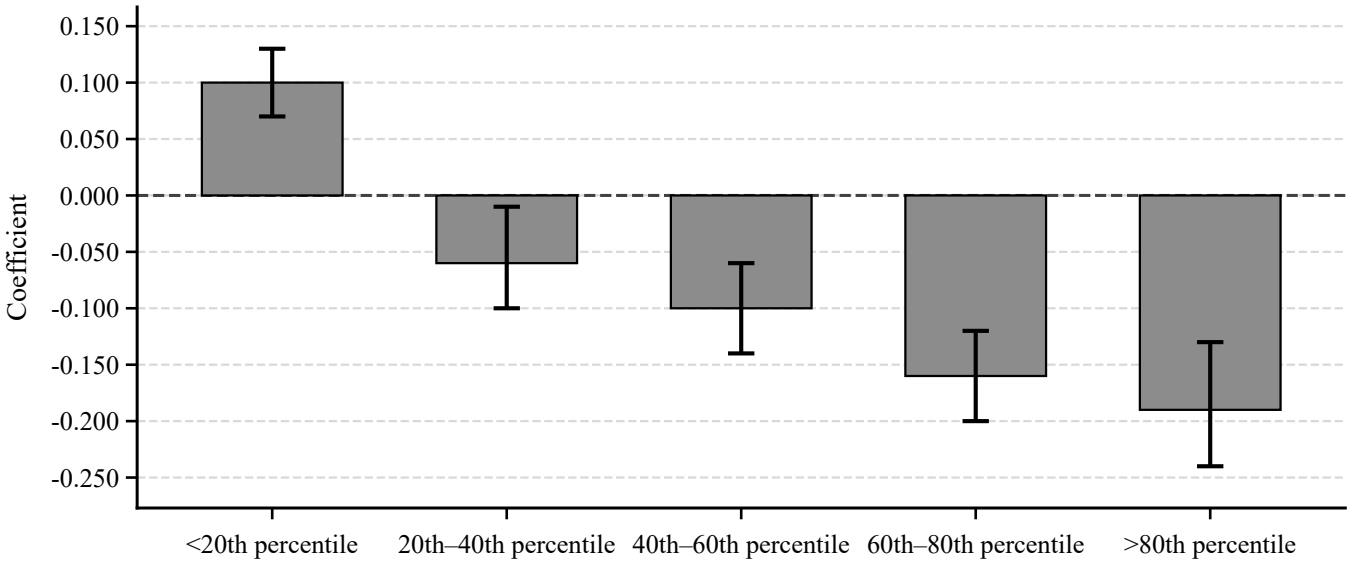


Figure A3: Variation in Consumer Sentiment within Each Quintile Interval of the Cloud Thickness

Notes: This figure plots the estimated coefficients and 95% confidence intervals from a regression of the consumer sentiment measure on five indicator variables corresponding to quintiles of local cloud thickness. To directly recover the group-specific conditional means, the regression omits the constant term. The quintile indicators are defined sequentially from the lowest (below the 20th percentile) to the highest (above the 80th percentile) cloud thickness. The regression model controls for temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level, and includes week-fixed effects to account for temporal variations. The x-axis denotes the cloud thickness quintiles, and the y-axis displays the corresponding point estimates.

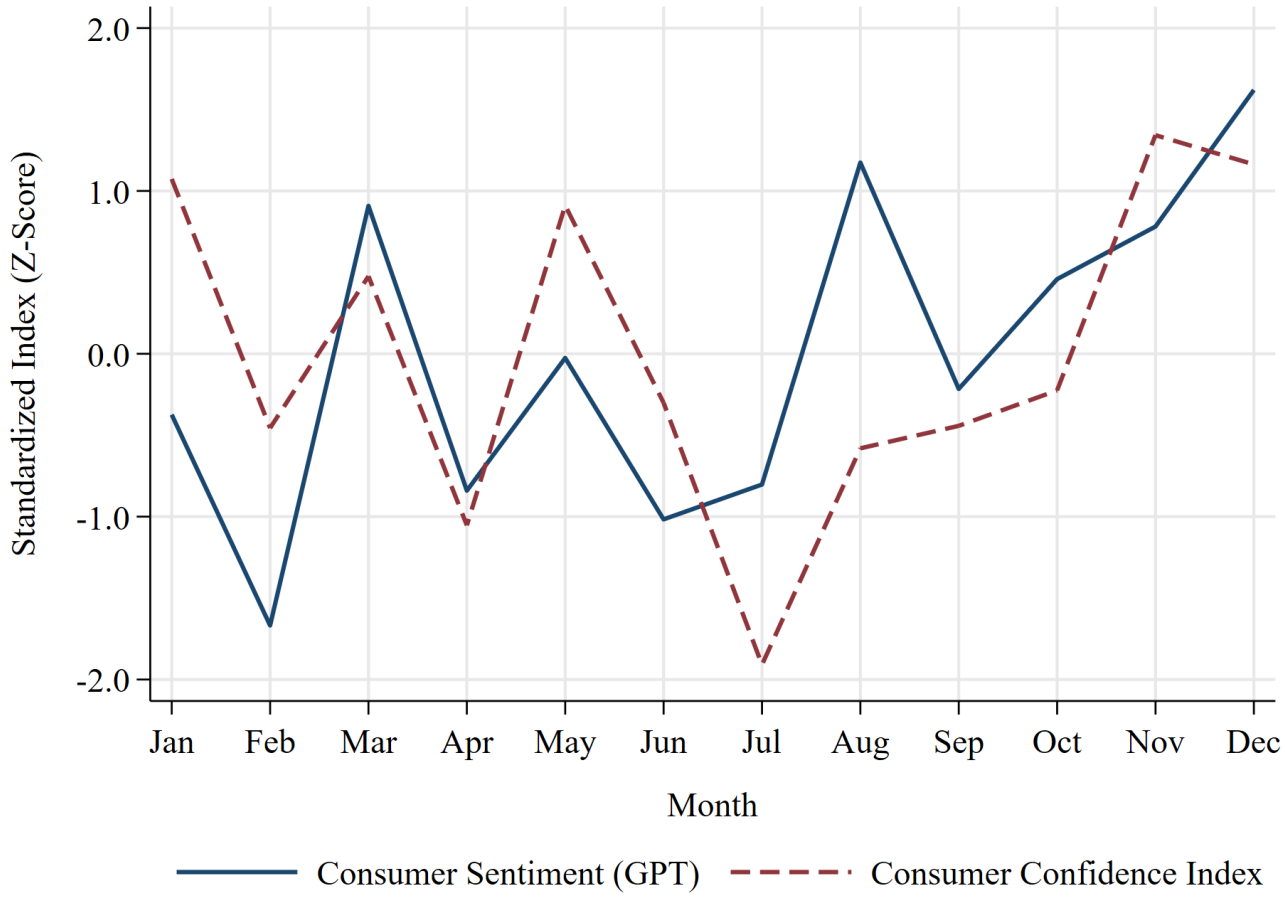


Figure A4: Correlation Between Consumer Sentiment and the Consumer Confidence Index

Notes: This figure plots our aggregate monthly national consumer sentiment index alongside the detrended Consumer Confidence Index published by the National Bureau of Statistics of China. The blue line depicts our constructed sentiment index for the current month, calculated as the weighted average of cities' consumer sentiment using cities' real GDP as weights. The red line plots the official, detrended Consumer Confidence Index for the subsequent month. To facilitate direct visual comparison, both indices are standardized to mean zero and unit variance. The horizontal axis denotes the calendar month, and the vertical axis displays the standardized index values. The correlation coefficient between the two time series is 0.52.

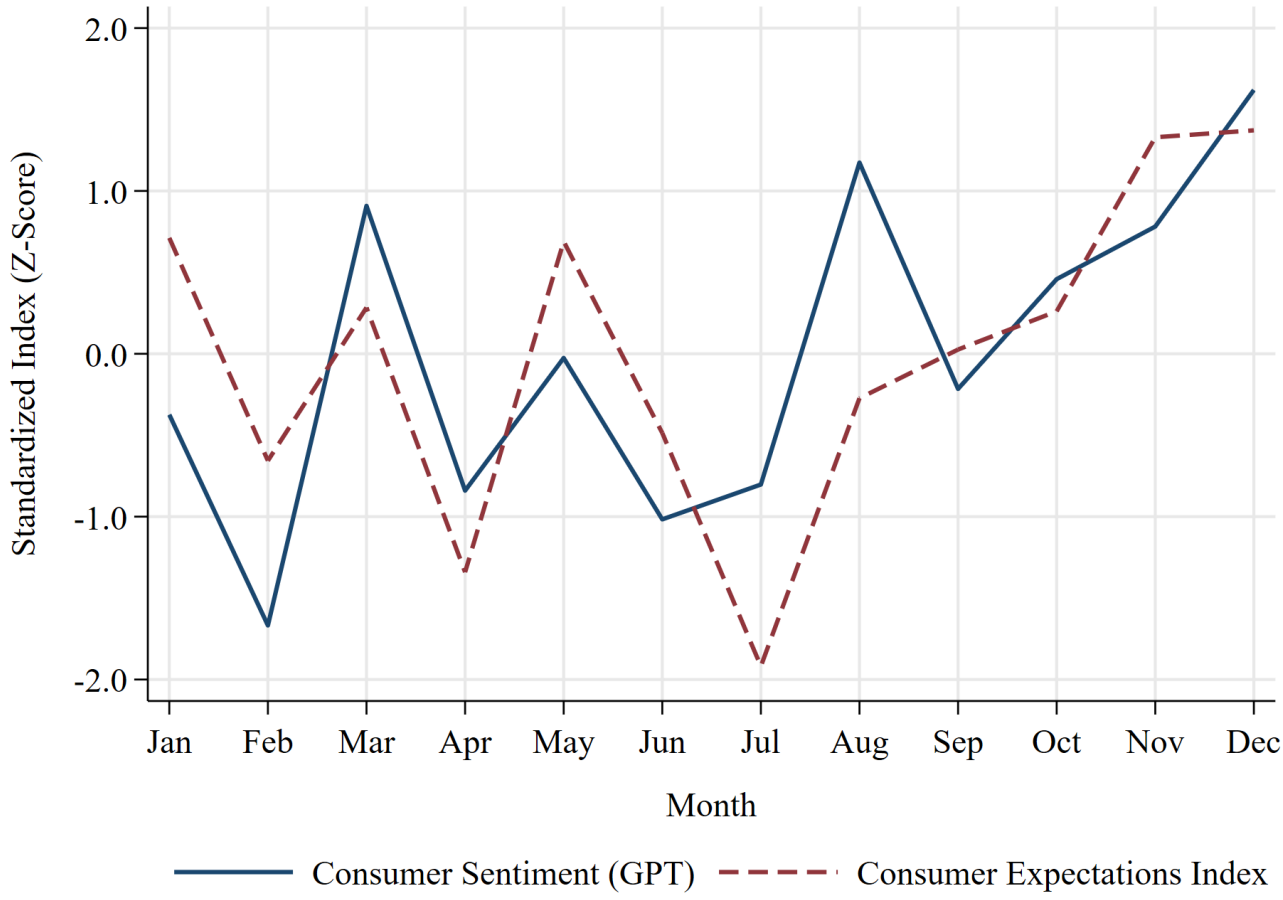


Figure A5: Correlation Between Consumer Sentiment and the Consumer Expectations Index

Notes: This figure plots our aggregate monthly national consumer sentiment index alongside the detrended Consumer Expectations Index published by the National Bureau of Statistics of China. The blue line depicts our constructed sentiment index for the current month, calculated as the weighted average of cities' consumer sentiment using cities' real GDP as weights. The red line plots the official, detrended Consumer Expectations Index for the subsequent month. To facilitate direct visual comparison, both indices are standardized to mean zero and unit variance. The horizontal axis denotes the calendar month, and the vertical axis displays the standardized index values. The correlation coefficient between the two time series is 0.66.

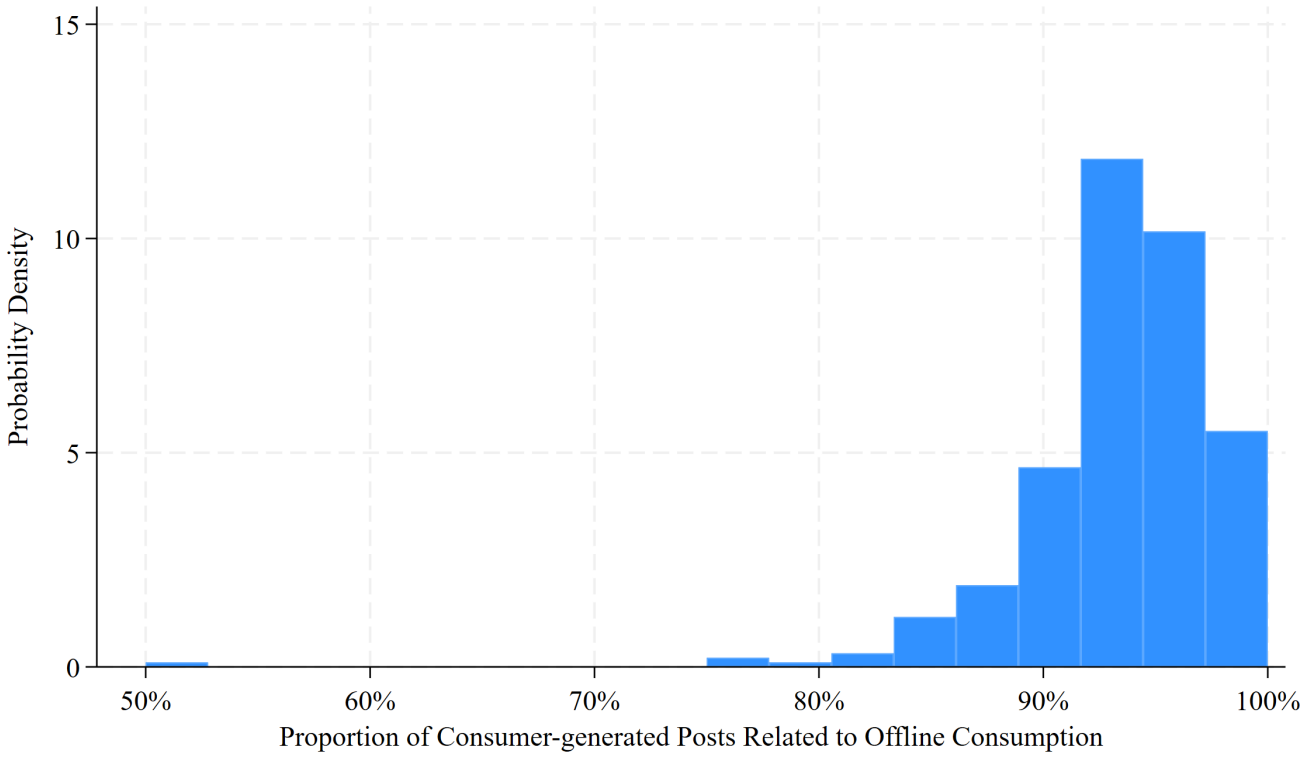


Figure A6: Distribution of Offline Posts Proportion

Notes: This figure plots the cross-city distribution of the share of posts related to offline activities. The horizontal axis denotes the city-level share of offline posts, and the vertical axis indicates the frequency of cities.

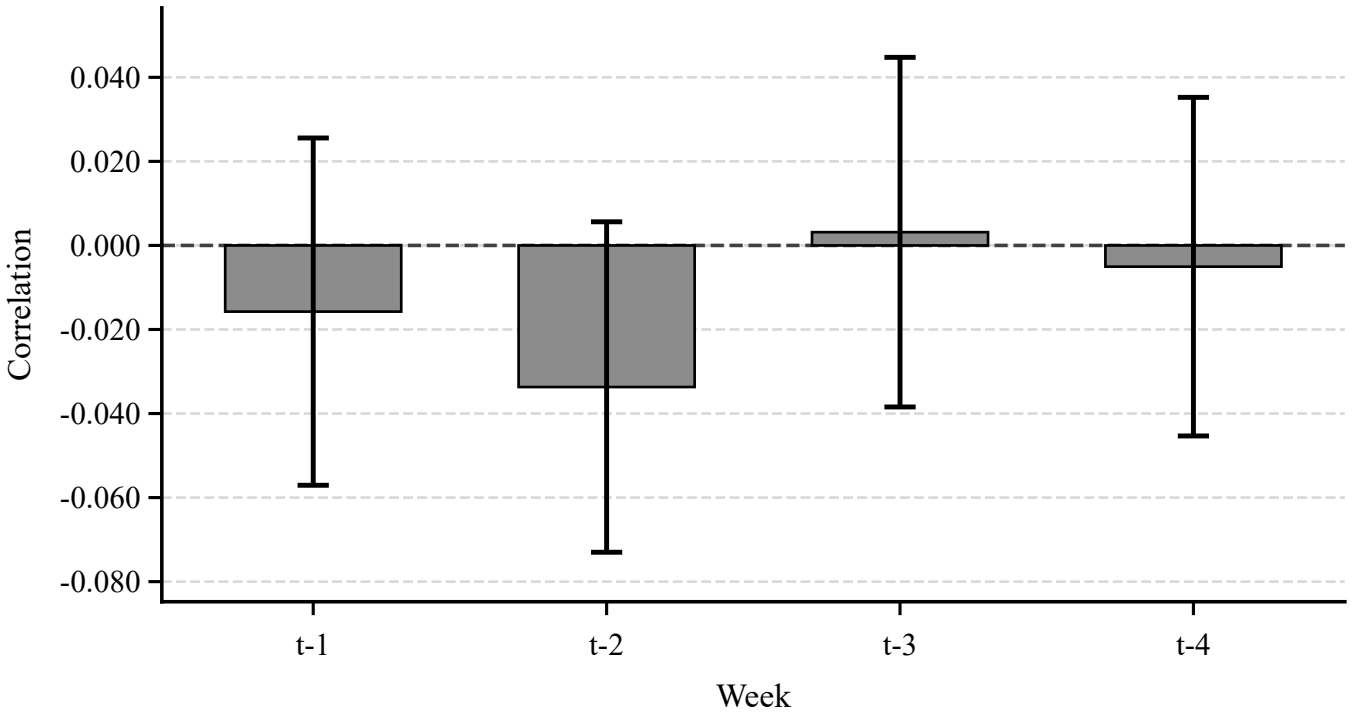


Figure A7: Correlation Between Current Sentiment and Sentiment over the Prior Four Weeks

Notes: This figure reports the estimated coefficients from regressions of current-period sentiment on its lagged values over the prior four weeks. All estimates are reported with 95% confidence intervals.

A.5 Other Sentiment Measurements for Robustness Checks

We construct other sentiment measures for robustness checks from three alternative perspectives.

First, we construct a city-week measure of cloud thickness to serve as an alternative proxy for local consumer sentiment. Specifically, we source global raster data on cloud cover from NASA Earth Observations[15]. We aggregate these pixel-level rasters within the administrative boundaries of Chinese cities to compute a daily cloud thickness for each city. We then calculate the simple average of daily cloud thickness within a week for each city, which generates our city-week measure of cloud thickness. As documented in prior literature, elevated local cloud thickness is posited to exogenously dampen consumer sentiment during the same period.

Second, we deploy DeepSeek as a substitute for OpenAI's GPT-4o-mini to classify the sentiment of the Weibo posts. This alternative classification protocol strictly mirrors the methodology detailed in Section 2.2, using prompts adapted to the Chinese linguistic context. As a complementary exercise, we also reclassify the sentiment of the Weibo posts using OpenAI's GPT-4o-mini with alternative prompts designed following Gao, Xiong, and Yuan (2026). Appendix A.7 provides details about the alternative prompts.

Third, we classify the sentiment of each Weibo post using a traditional dictionary-based approach (e.g., Loughran and McDonald (2011)). To account for the specific linguistic nuances of Chinese social media text, we execute this classification using two established Chinese sentiment dictionaries developed by Tsinghua University and the Dalian University of Technology.

A.6 Prompt that Removes Product Quality-related Posts and Performs Sentiment Scoring

You are an expert in Natural Language Processing and Sentiment Analysis. You will be provided with a social media post in Chinese. Your tasks are to (1) classify the sentiment of the given post as Positive or Negative based on its content and (2) identify whether the given post is related to product quality.

1. You need to think step by step, following the chains below: (1) Read the entire post carefully to fully understand the context. Focus on the overall message, tone, and the author's intent. (2) Break down the post into key phrases or expressions that reflect sentiment. Pay attention to words that carry emotional weight. (3) Analyze how each key phrase and expression contributes to the overall sentiment under its context. It is important to consider cultural nuances, idiomatic expressions, and linguistic features unique to Chinese. (4) Detect any sarcasm, irony, or contradictions that may alter the apparent sentiment, especially for those unique to Chinese. When encountering those complex emotions, carefully analyze the sentiment within its full context. (5) Based on all the analyses above, classify the sentiment as Positive or Negative. (6) After finishing the classification work, re-evaluate each analysis step contributing to your final classification. Consider how to

[15] https://neo.gsfc.nasa.gov

improve your analysis and refine your classification based on those considerations. (7) After completing the sentiment classification, identify whether the given post is related to product quality. Look for mentions of defects, performance, durability, or other quality-related attributes. (8) Iterate the refinements in steps (6) - (7) a total of 10 times. In each iteration, consider how to improve the sentiment classification based on the previous iteration. (9) Keep track of the iteration number after each refinement. Your final output should include all sentiment labels and scores from the 10 iterations, and an indicator of whether the post is related to product quality.

2. Important reminds

(1) Consider specific cultural characteristics, idiomatic expressions, linguistic features, and slang unique to the Chinese language.

(2) Try your best to avoid measurement errors and account for any biases towards positive or negative sentiment.

3. Do not carry any inference beyond the contents of the given post. Stick strictly to the content provided. Do not introduce external information or assumptions.

The output should be a JSON array containing a single object for each iteration. For each of the 10 iterations, provide the output in the following JSON structure: "Class": "[Positive or Negative]", "Score": "[a number between 0 and 1]", "Product Quality": "[Yes or No]". The Class field should be either "Positive" or "Negative". The Score field should be a numerical value between 0 and 1, where 0 represents the lowest sentiment intensity, and 1 represents the highest sentiment intensity. The Product Quality field should be either "Yes" or "No"., where "Yes" indicates relevance to product quality and "No" indicates irrelevance to product quality. Do not include any additional text or explanation outside of this JSON structure.

A.7 Additional Prompt for Robustness Checks

Following Gao, Xiong, and Yuan (2026), we use another two prompts as the input of OpenAI's GPT Model. The first prompt is a zero-shot prompt:

Forget all your previous instructions. Assume the role of a Chinese expert specializing in sentiment analysis. Your task is to assess the sentiment of a social media post in Chinese. Respond with 'Positive', 'Weakly Positive', 'Neutral', 'Weakly Negative', 'Negative', or 'Not Mentioned' (not applicable for the overall assessment). Do not carry any inference beyond the contents of the given post. Stick strictly to the content provided. Do not introduce external information or assumptions.

The second prompt is a few-shot prompt:

Forget all your previous instructions. Assume the role of a Chinese expert specializing in sentiment analysis.
Your task is to assess the sentiment of a social media post in Chinese.
Respond with 'Positive', 'Weakly Positive', 'Neutral', 'Weakly Negative', 'Negative', or 'Not Mentioned' (not applicable for the overall assessment).
Do not carry any inference beyond the contents of the given post. Stick strictly to the content provided. Do not introduce external information or assumptions.
Some examples:
Example 1:
Weibo post: '新年第一发前天一家人去逛花市，马上都要关门时我们非常冲动的在根本不知道尺寸情况下买了一个大花架，又匆忙买了几盆小花，回家组装后发现装不进去然而感谢我有一个动手能力强的理工男老公，在我们智慧的结合下，家里多了一个非常浪漫、美丽、梦幻的小角落，让这个跨年的夜晚炫丽出彩，幸福感满满花花还没有采购齐全，明天继续花市采购，属于我们的小角落啊，美起来！2019??用心装点每一天.'
Output: 'Positive'.
Example 2:
Weibo post: 在这儿吃了三年啦！真的是越来越不如以前了啊！买了一根猪蹄一份鸡爪、看到牌子上写的第二份半价就问：'你好、帅哥我想问一下猪蹄是第二份半价吗？'旁边一个帅哥语气特别不好的回了一句：'美女麻烦你看清楚了再说嘛！真的是醉了！啥子态度哦！'我说：'对啊、因为不知道所以问一下啊。'他有怼了一句：'是啊、叫你看清楚得嘛！'哎！真的是成为因为态度问题！就算再好吃也不会去的店啦！黑名单加入.'
Output: 'Negative'.

When constructing our main sentiment measure, we excluded posts related to product quality. In the prompt described below, we further instructed the large language model to identify whether a post discussed the quality, price, or quality–price tradeoff of a product or service, and to exclude posts referring to these fundamental product or service characteristics:

You are an expert in natural language processing, consumer behavior, and sentiment analysis. You will be given a Chinese social media post. Your task is to classify the post along three dimensions, using ONLY the permitted labels defined below.
1. CLASSIFICATION DIMENSIONS AND PERMITTED LABELS
Dimension 1: Sentiment (assign exactly one): - "Positive" - "Weakly Positive" - "Neutral" - "Weakly Negative" - "Negative"
Dimension 2: Sentiment_Object (assign one or more labels): - "Product/Service Advertising" - "Product/Service Quality" - "Product/Service Price" - "Value for Money" - "None"
Dimension 3: Channel (assign exactly one): - "Online" - "Offline" - "Both" - "Other"

2. STEP-BY-STEP CLASSIFICATION PROCEDURE

Follow the reasoning chain below. All reasoning should be conducted silently; only the final JSON object is returned.

Step 1: Comprehend the post. Read the post in full and identify: (a) the overall message; (b) the tone; (c) the poster's intent.

Step 2: Identify sentiment-related components. Decompose the post into the phrases or expressions that convey sentiment.

Step 3: Contextual sentiment analysis. Evaluate how each sentiment-related component contributes to the overall sentiment within its surrounding context.

Account for Chinese linguistic and cultural conventions, including idioms, slang, sarcasm, rhetorical questions, and context-dependent expressions.

Step 4: Detect complex sentiment. Identify any sarcasm, irony, or contradiction—particularly those rooted in Chinese-specific cultural or linguistic conventions—that may invert or attenuate the surface-level sentiment. Re-evaluate affected phrases within the full context of the post before proceeding.

Step 5: Assign the Sentiment label. On the basis of Steps 1–4, assign exactly one label from Dimension 1.

Step 6: Assign the Sentiment_Object label(s).

Determine whether the sentiment-related components of the post articulate an evaluative stance toward product- or service-related advertising or toward specific attributes of a product or service, namely quality, price level, or value for money.

Apply the following rules:

(a) Restrict the analysis to text segments that convey sentiment. Within those segments, identify expressions that explicitly or implicitly evaluate quality, price level, value for money, or advertising of a product or service. Disregard expressions that merely state a monetary amount, mention a product, service, or brand, or describe a product or service feature without conveying an evaluative judgment.

(b) Assign "Product/Service Quality" when the sentiment-related text evaluates performance, functionality, reliability, durability, design, user experience, or other quality-related attributes of a product or service. For services, this includes evaluations of service attitude, friendliness, professionalism, responsiveness, or wait time, as well as praise or criticism directed at staff, owners, or other service providers when the trait being evaluated is service-relevant.

(c) Assign "Product/Service Price" when the sentiment-bearing text evaluates the price level itself, such as whether the price is high, low, fair, cheap, expensive, discounted, or excessive. A mere statement of the amount paid does not qualify.

(d) Assign "Value for Money" when the sentiment-related text evaluates whether a product's or a service's quality, utility, performance, or satisfaction justifies its price, such as whether a product or a service is worth the price, cost-effective, overpriced relative to quality, or poor value for money.
(e) Assign "Product/Service Advertising" when the sentiment-related text either constitutes promotional content for a product or service, or evaluates product- or service-related advertising, sponsored promotion, or marketing communication.
(f) Assign "None" when the sentiment-related components of the post contain no evaluation of advertising, quality, price, or value for money of a product or service.
(g) Multiple labels among "Product/Service Quality", "Product/Service Price", and "Value for Money" may be co-assigned when the post addresses several attributes with comparable emphasis, provided that each label is supported by a distinct evaluative judgment.
(h) "Product/Service Advertising" and "None" are each mutually exclusive with all other labels; when either is assigned, it must be the sole element of the array.
(i) Content irrelevant to sentiment must be disregarded entirely when assigning this dimension.
Step 7: Assign the Channel label. Determine whether the activities described in the post occurred online, offline, or in both settings, and assign the corresponding label. If the post does not provide sufficient information to determine whether the activity occurred online or offline, assign "Other".
Step 8: Refinement. Re-examine the labels assigned in Steps 5, 6, and 7. Revise any label whose justification is weak or inconsistent with the content of the post.
Step 9: Internal verification. Before returning the JSON object, verify that: (a) "Sentiment" is exactly one of the five labels listed in Dimension 1. (b) "Sentiment_Object" is a non-empty array drawn only from the vocabulary listed in Dimension 2. (c) If "Sentiment_Object" contains "Product/Service Advertising" or "None", that label is the sole element of the array. (d) "Channel" is exactly one of the four labels listed in Dimension 3. (e) The classification relies solely on the content of the post.
3. IMPORTANT REMINDERS (1) Avoid measurement error and systematic bias toward either positive or negative sentiment. (2) Do not draw inferences beyond the content of the post. (3) Do not introduce external information or assumptions.

Table A1: List of Keywords

purchase	overseas shopping	price	supplier	rice	brand
free shipping	shopping guide	direct mail	after discount	member	travel
code	unit price	product	medicine	price reduction	shop owner
kilogram	jewelry	expenditure	eggs	honey	eat
impulse buying	purchasing agent	discount	bargain hunting	buy in bulk	can buy with confidence
buy more	shop with confidence	purchasing agents	buy replicas	buy	better to buy
select products	reliable purchasing agent	reliable shopping agent	buy	consumption	bought
cushion	purchase	shop for luxury brands	set	second-hand items	this product
fruits	pork	rice	kilogram	potatoes	meat
vegetables	oil	dining	seasoning		

A.8 Addition Results

A.8.1 Summary Statistics

Table A2: Summary Statistics

Variable	**Obs.**	**Mean**	**SD**	**Min**	**Max**
$Log(Amount_{i,c,f,t})$	7362956	9.8897	10.6728	5.2983	13.9989
$Log(Transactions_{i,c,f,t})$	7362956	1.9286	1.2112	0.0000	3.6109
$Log(Variety_{i,c,f,t})$	7362956	0.8892	0.6580	0.0000	1.7918
$Online_{i,c,f,t}$	7362956	0.5000	0.0000	0.0000	1.0000
$Sentiment_{c,t-1}$	7362956	0.4773	0.3277	-1.0000	1.0000
$Humidity_{c,t-1}$	7362956	0.0489	0.0642	0.0000	1.0000
$Sunshine_{c,t-1}$	7362956	0.4028	0.1865	0.0000	1.0000
$Temperature_{c,t-1}$	7362956	0.7179	0.1713	0.0000	1.0000
$Wind\ Speed_{c,t-1}$	7362956	0.1494	0.0602	0.0000	1.0000
$AQI_{c,t-1}$	7362956	0.1629	0.0661	0.0000	1.0000

Notes: This table reports the descriptive statistics for the key variables employed in our empirical analysis.

A.8.2 Quarterly Effects of Sentiment on Consumption

Table A3: Quarterly Effects of Sentiment on Consumption

	$Log(OverallAmount_{i,c,q})$	$Log(OverallTransactions_{i,c,q})$
	(1)	(2)
$Sentiment_{c,q-1}$	0.0086**	0.0220**
	(0.0038)	(0.0101)
Controls	Y	Y
Quarter FE	Y	Y
Consumer FE	Y	Y
R^2	0.8954	0.8688
N	290331	290331

Notes: This table presents the quarterly effects of sentiment on consumption, estimated by ordinary least squares regression. The dependent variables in columns (1) and (2) are the natural logarithms of total consumption amount and transaction volume in quarter q, respectively. The explanatory variable, $Sentiment_{c,q-1}$, represents the sentiment level in city c at the end of quarter $q-1$. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-quarter level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.3 Non-linear Relationship Between Sentiment and Consumption

Although the level effects of sentiment on overall consumption are insignificant on average, we document a nonlinear, J-shaped relationship between consumer sentiment and overall consumption. Empirically, we augment Equation 2 with a quadratic sentiment term, $Sentiment^2_{c,t-1}$, allowing the marginal effect of sentiment to vary with its baseline level. Panel A of Table A4 reports the corresponding estimates. The coefficient on the linear term ($Sentiment_{c,t-1}$) is negative and statistically indistinguishable from zero, whereas the coefficient on the quadratic term ($Sentiment^2_{c,t-1}$) is positive and significant. This combination formally illustrates a J-shaped curve, indicating that the stimulative effect of sentiment on aggregate consumption accelerates at higher sentiment tiers. Furthermore, Panel B of Table A4 reports the estimates from an augmented version of Equation (1), which incorporates the quadratic sentiment term, $Sentiment^2_{c,t-1}$, alongside its interaction with the online channel indicator. In stark contrast to the case of overall consumption, the marginal effects of sentiment on online and offline expenditure do not exhibit the convexity observed in Panel A.

Table A4: Non-linear Relationship Between Sentiment and Consumption

Panel A: Overall Consumption			
	$Log(Amount_{i,c,t})$	$Log(Transactions_{i,c,t})$	$Log(Variety_{i,c,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0031	-0.0005	-0.0009
	(0.0035)	(0.0018)	(0.0008)
$Sentiment^2_{c,t-1}$	0.0117**	0.0074***	0.0017
	(0.0057)	(0.0028)	(0.0013)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.6246	0.7196	0.04678
N	3681478	3681478	3681478
Panel B: Online and Offline Consumption			
	$Log(Amount_{i,c,f,t})$	$Log(Transactions_{i,c,f,t})$	$Log(Variety_{i,c,f,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0119**	-0.0073***	-0.0037**
	(0.0052)	(0.0027)	(0.0015)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0240***	0.0172***	0.0071***
	(0.0078)	(0.0043)	(0.0024)
$Sentiment^2_{c,t-1}$	0.0082	0.0022	0.0049**
	(0.0083)	(0.0043)	(0.0024)
$Sentiment^2_{c,t-1} \times Online_{i,c,f,t}$	-0.0118	-0.0001	-0.0093**
	(0.0124)	(0.0069)	(0.0038)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.4986	0.4856	0.4635
N	7362956	7362956	7362956

Notes: This table presents the non-linear association between sentiment and consumption, estimated by ordinary least squares regression. In Panel A, the dependent variables in columns (1)–(3) are the natural logarithms of total consumption amount, transaction volume, and consumption variety in week t, respectively. In Panel B, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. $Sentiment^2_{c,t-1}$ is the squared term of $Sentiment_{c,t-1}$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.4 Consumption Price

Table A5: Level Effects of Sentiment on Payment Amount per Transaction

	$Log(Overall\ Price_{i,c,t})$ (1)	$Log(Price_{i,c,f,t})$ (2)
$Sentiment_{c,t-1}$	-0.0019	-0.0069*
	(0.0028)	(0.0042)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$		0.0123**
		(0.0060)
Controls	Y	Y
Week FE	Y	Y
Consumer FE	Y	Y
Consumer-Month FE	Y	Y
Online-Consumer FE	N	Y
Online-Consumer-Month FE	N	Y
R^2	0.6246	0.5842
N	3681478	7362956

Notes: This table presents the impact of sentiment on consumption amount per transaction, estimated by ordinary least squares regression. In column (1), the dependent variable is the natural logarithms of overall payment amounts per transaction in week t. In column (2), the dependent variable is the natural logarithms of payment amounts per transaction by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.5 Product-Level Heterogeneity

In this section, we conduct a more granular, category-level heterogeneity analysis. Specifically, we disaggregate each consumer's transaction volume across the 17 distinct product categories, separating transaction volumes by the online and offline channels. We then estimate the impact of sentiment independently for each category. The results, reported in Figure A9, demonstrate that positive sentiment induces a significant reallocation from offline to online expenditure predominantly concentrated within five specific categories: daily necessities, clothing, electronic goods, automotive services, and education.

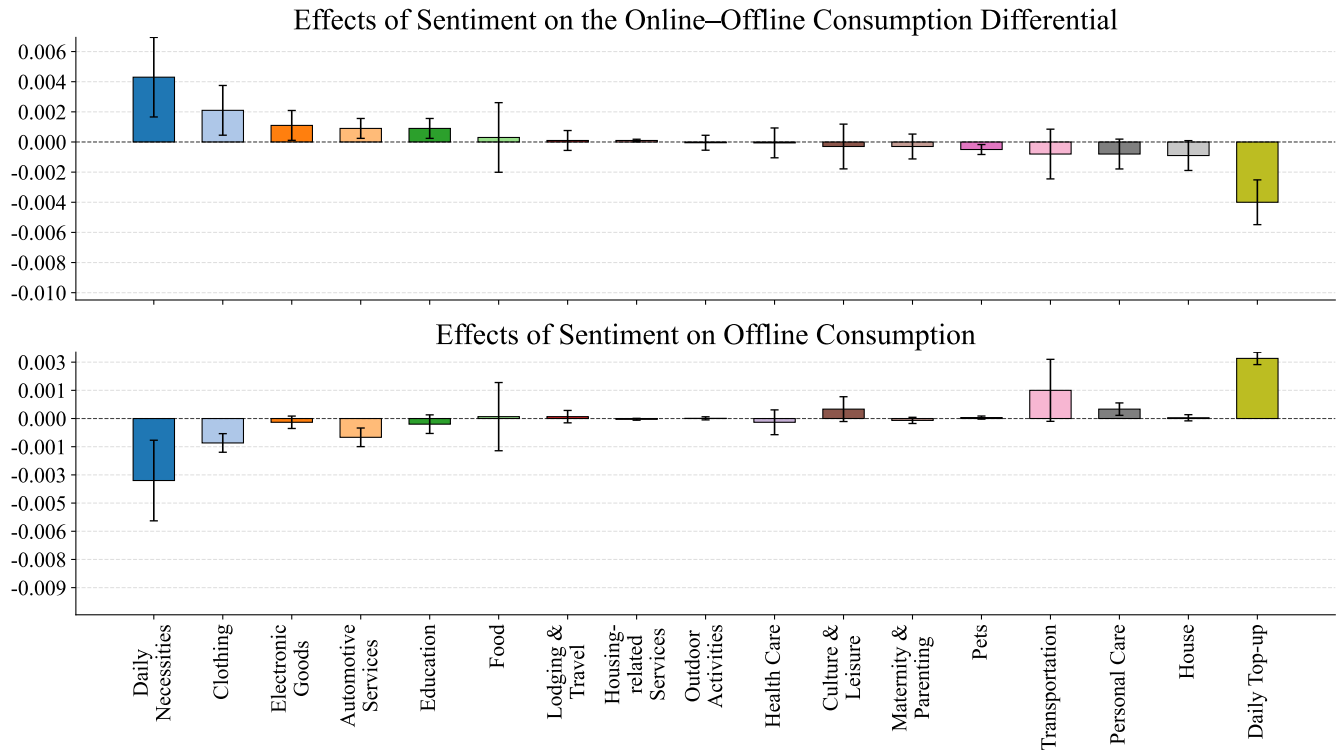


Figure A8: Impacts of Sentiment on Consumption by Product Category

Notes: This figure presents the heterogeneous effects of sentiment on transaction volume across different categories of consumption goods, estimated by ordinary least squares regression. All estimates are reported with 95% confidence intervals.

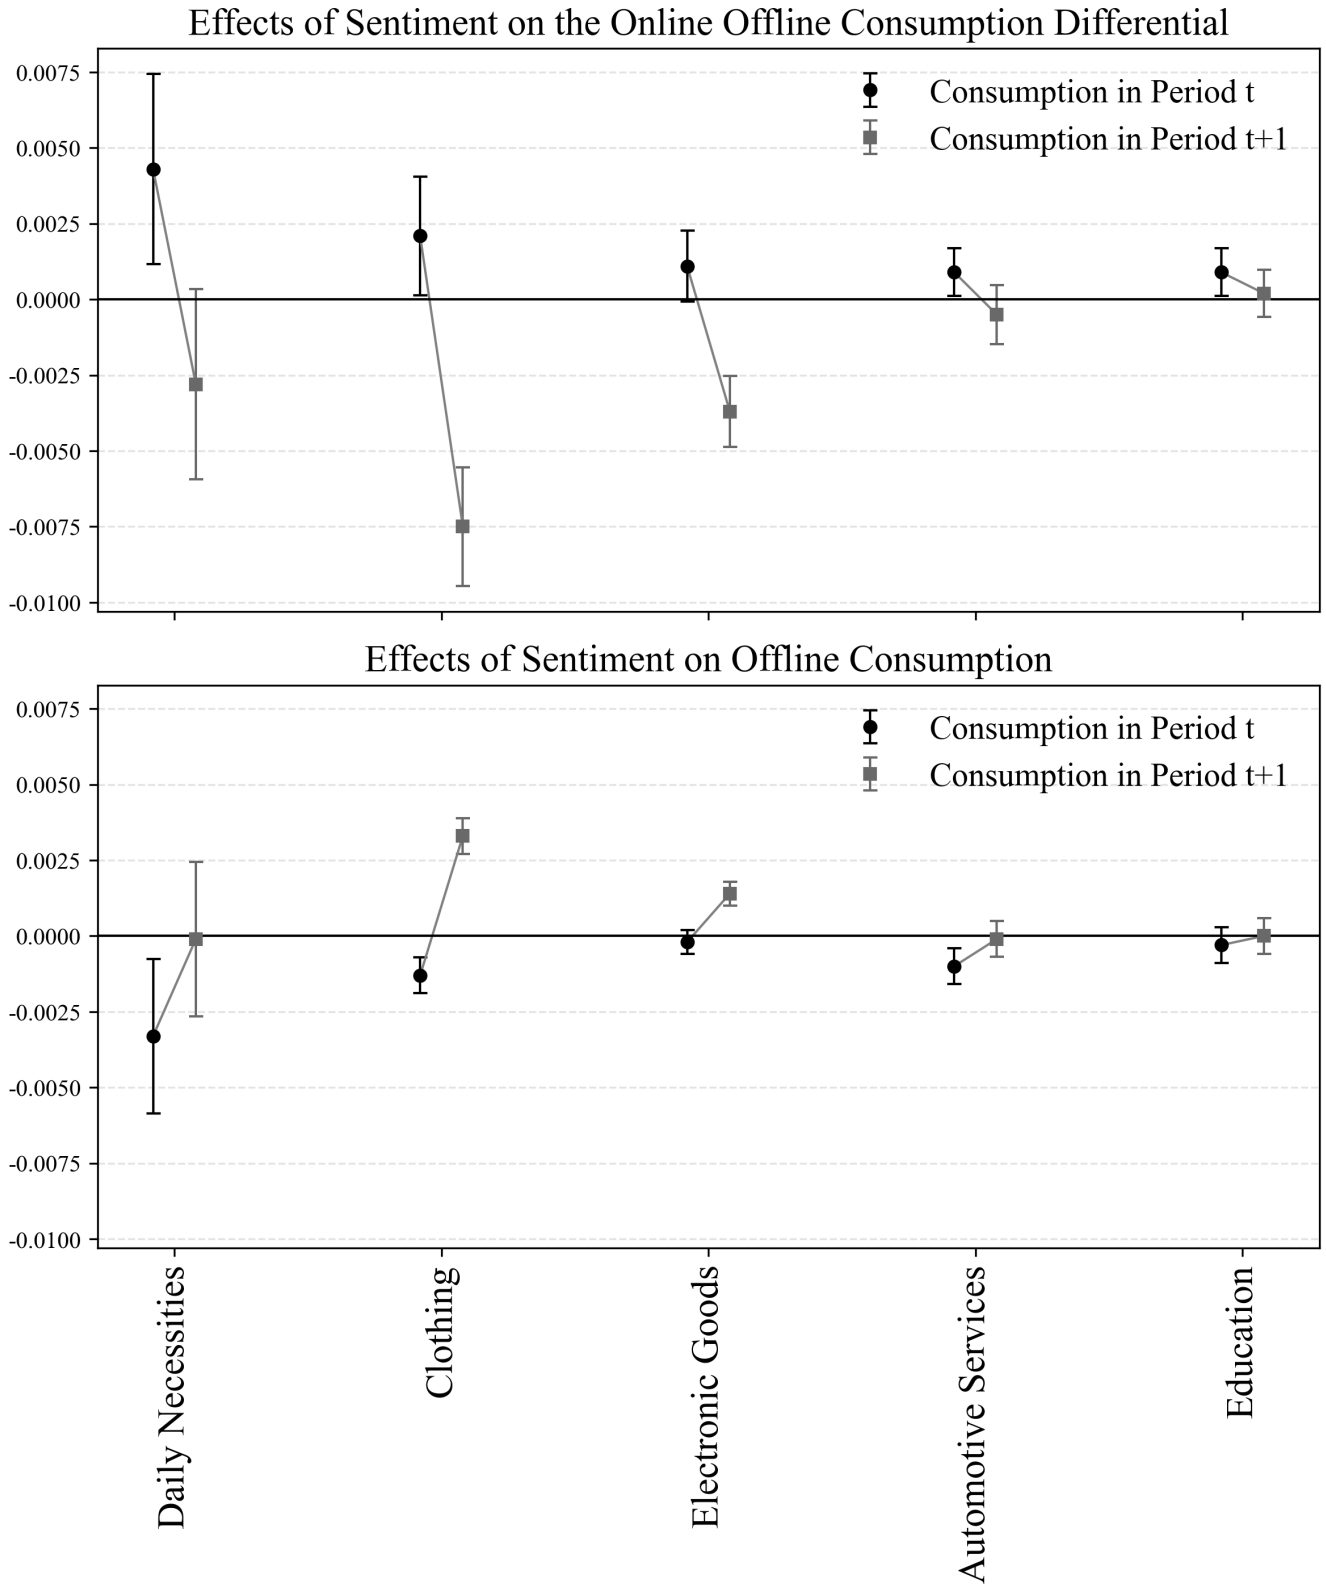


Figure A9: Dynamic Effects of Sentiment on Consumption by Product Category

Notes: This figure presents the heterogeneous effects of city-level sentiment in week $t-1$ on transaction volume in weeks t and $t+1$ across different categories of consumption goods, estimated by ordinary least squares regression. All estimates are reported with 95% confidence intervals.

A.8.6 Alternative sentiment measure

As described in Section 2.2, Weibo posts related to product quality were excluded to avoid interference from fundamental information about product quality. To test the impact of this sample exclusion on the results, we retain consumers' Weibo posts related to product quality and use the full sample of Weibo posts to construct an alternative measure of city-week consumer sentiment. Table A6 reports the results based on this sentiment measure, which bolsters the robustness of our main results.

Second, we measure consumer sentiment by city-week cloud thickness, as detailed in Section A.5. Table A7 reports the regression results based on the cloud-based sentiment measure, which support the robustness of our main results.

Furthermore, we employ DeepSeek to identify the sentiment of Weibo posts and also use a traditional dictionary-based approach to identify the sentiment of Weibo posts, as detailed in Section A.5. In Table A8, the odd columns present the results based on the sentiment measure constructed using DeepSeek, while the even columns show the results based on the sentiment measure constructed using a dictionary-based approach. The estimates of results are consistent with our main results.

Finally, we design alternative prompts following Gao, Xiong, and Yuan (2026) and use them to classify the sentiment of Weibo posts, as described in Section A.5. Based on these classifications, we construct alternative city–week sentiment measures. Table A9 reports regression results using these sentiment measures, which corroborate our main findings.

Table A6: Robustness Checks: Sentiment Constructed before Deleting Quality-related Posts

Panel A: Overall Consumption			
	$Log(Amount_{i,c,t})$	$Log(Transactions_{i,c,t})$	$Log(Variety_{i,c,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0044	-0.0014	-0.0010
	(0.0033)	(0.0016)	(0.0008)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.6246	0.7196	0.4678
N	3681478	3681478	3681478
Panel B: Online and Offline Consumption			
	$Log(Amount_{i,c,f,t})$	$Log(Transactions_{i,c,f,t})$	$Log(Variety_{i,c,f,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0136***	-0.0134***	-0.0049***
	(0.0048)	(0.0025)	(0.0014)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0257***	0.0271***	0.0080***
	(0.0071)	(0.0040)	(0.0022)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.4986	0.4856	0.4635
N	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption, using an alternative sentiment measure constructed before excluding posts discussing fundamental product quality. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables in columns (1)–(3) are the natural logarithms of total consumption amount, transaction volume, and consumption variety in week t, respectively. In Panel B, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table A7: Robustness Checks: Proxy Sentiment by Cloud Thickness

Panel A: Overall Consumption

	$Log(Amount_{i,c,t})$			$Log(Transactions_{i,c,t})$			$Log(Variety_{i,c,t})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Cloud_{c,t-1}$	-0.0088	-0.0121	-0.0096	-0.0054	-0.0025	-0.0060	-0.0026	-0.0044	-0.0029
	(0.0077)	(0.0132)	(0.0085)	(0.0040)	(0.0069)	(0.0045)	(0.0020)	(0.0034)	(0.0022)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.6386	0.6386	0.6386	0.7326	0.7326	0.7326	0.4933	0.4933	0.4933
N	3112199	3112199	3112199	3112199	3112199	3112199	3112199	3112199	3112199

Panel B: Online and Offline Consumption

	$Log(Amount_{i,c,f,t})$			$Log(Transactions_{i,c,f,t})$			$Log(Variety_{i,c,f,t})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Cloud_{c,t-1}$	0.0775***	0.1191***	0.0878***	0.0602***	0.0934***	0.0664***	0.0241***	0.0338***	0.0277***
	(0.0103)	(0.0180)	(0.0114)	(0.0054)	(0.0094)	(0.0059)	(0.0030)	(0.0052)	(0.0033)
$Cloud_{c,t-1} \times Online_{i,c,f,t}$	-0.1590***	-0.2311***	-0.1798***	-0.1180***	-0.1699***	-0.1300***	-0.0487***	-0.0618***	-0.0560***
	(0.0133)	(0.0237)	(0.0147)	(0.0074)	(0.0131)	(0.0081)	(0.0041)	(0.0073)	(0.0045)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.5232	0.5232	0.5232	0.5116	0.5115	0.5116	0.4900	0.4900	0.4900
N	6224398	6224398	6224398	6224398	6224398	6224398	6224398	6224398	6224398

Notes: This table reports the level effects of sentiment on consumption, using local cloud thickness to construct alternative proxies for sentiment. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are the natural logarithms of total consumption amount in week t from columns (1) to (3), transaction volume in week t from columns (4) to (6), and consumption variety in week t from columns (7) to (9), respectively. In Panel B, the dependent variables include the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The key explanatory variable, $Cloud_{c,t-1}$, captures the local sentiment, in three distinct ways: columns (1), (4), and (7) employ the local cloud thickness in city c during week $t-1$; columns (2), (5), and (8) use the difference between local cloud thickness in week $t-1$ and the average local cloud thickness for the same week over the past five years; and columns (3), (6), and (9) use the detrended cloud thickness in city c during week $t-1$. Higher values of $Cloud_{c,t-1}$ indicate more negative local sentiment. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table A8: Robustness Checks: Alternative Sentiment Measures

Panel A: Overall Consumption	$Log(Amount_{i,c,t})$		$Log(Transactions_{i,c,t})$		$Log(Variety_{i,c,t})$	
	DeepSeek	Dictionary	DeepSeek	Dictionary	DeepSeek	Dictionary
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	-0.0022	0.0026	-0.0015	-0.0033*	-0.0019**	-0.0003
	(0.0038)	(0.0037)	(0.0019)	(0.0019)	(0.0009)	(0.0009)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.6246	0.6246	0.7196	0.7196	0.4678	0.4678
N	3681478	3681478	3681478	3681478	3681478	3681478
Panel B: Online and Offline Consumption	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	DeepSeek	Dictionary	DeepSeek	Dictionary	DeepSeek	Dictionary
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	-0.0115**	-0.0104**	-0.0121***	-0.0136***	-0.0036**	-0.0056***
	(0.0054)	(0.0053)	(0.0028)	(0.0028)	(0.0016)	(0.0015)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0261***	0.0253***	0.0269***	0.0236***	0.0066***	0.0090***
	(0.0081)	(0.0081)	(0.0045)	(0.0045)	(0.0025)	(0.0024)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.4986	0.4986	0.4856	0.4856	0.4635	0.4635
N	7362956	7362956	7362956	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption, using other textual analysis approaches to construct alternative proxies for sentiment. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are the natural logarithms of total consumption amount in week t from columns (1) to (2), transaction volume in week t from columns (3) to (4), and consumption variety in week t from columns (5) to (6), respectively. In Panel B, the dependent variables include the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. In columns (1), (3), and (5), the explanatory variable $Sentiment_{c,t-1}$ is constructed by DeepSeek. In columns (2), (4), and (6), the explanatory variable $Sentiment_{c,t-1}$ is constructed by a dictionary-based method. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table A9: Robustness Checks: Alternative Prompt for Measuring Sentiment

Panel A: Overall Consumption	$Log(Amount_{i,c,t})$		$Log(Transactions_{i,c,t})$		$Log(Variety_{i,c,t})$	
	Zero Shot	Few Shot	Zero Shot	Few Shot	Zero Shot	Few Shot
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	0.0066*	0.0005	-0.0011	-0.0018	-0.0010	-0.0015
	(0.0038)	(0.0040)	(0.0019)	(0.0020)	(0.0009)	(0.0009)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.6246	0.6246	0.7196	0.7196	0.4678	0.4678
N	3681478	3681478	3681478	3681478	3681478	3681478
Panel B: Online and Offline Consumption	$Log(Amount_{i,c,f,t})$		$Log(Transactions_{i,c,f,t})$		$Log(Variety_{i,c,f,t})$	
	Zero Shot	Few Shot	Zero Shot	Few Shot	Zero Shot	Few Shot
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{c,t-1}$	-0.0174***	-0.0222***	-0.0221***	-0.0231***	-0.0037**	-0.0049***
	(0.0054)	(0.0058)	(0.0028)	(0.0030)	(0.0016)	(0.0016)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0451***	0.0478***	0.0458***	0.0470***	0.0091***	0.0098***
	(0.0082)	(0.0086)	(0.0046)	(0.0048)	(0.0025)	(0.0026)
Controls	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y
R^2	0.4986	0.4986	0.4856	0.4856	0.4635	0.4635
N	7362956	7362956	7362956	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption, using alternative prompt architectures for the large language model to perform sentiment classification. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are the natural logarithms of total consumption amount in week t from columns (1) to (2), transaction volume in week t from columns (3) to (4), and consumption variety in week t from columns (5) to (6), respectively. In Panel B, the dependent variables include the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. In columns (1), (3), and (5), the explanatory variable $Sentiment_{c,t-1}$ is constructed using GPT-4o-mini with a zero-shot prompt designed following Gao, Xiong, and Yuan (2026). In columns (2), (4), and (6), the explanatory variable $Sentiment_{c,t-1}$ is constructed using GPT-4o-mini with a few-shot prompt designed following Gao, Xiong, and Yuan (2026). The detailed contents of the two prompts are shown in Appendix A.7. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table A10: Robustness Checks: Excluding Price Components and Price–Quality Tradeoffs

Panel A: Overall Consumption			
	$Log(Amount_{i,c,t})$	$Log(Transactions_{i,c,t})$	$Log(Variety_{i,c,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	0.0059	-0.0034	-0.0018
	(0.0061)	(0.0025)	(0.0020)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
R^2	0.6323	0.7364	0.6326
N	3681478	3681478	3681478
Panel B: Online and Offline Consumption			
	$Log(Amount_{i,c,f,t})$	$Log(Transactions_{i,c,f,t})$	$Log(Variety_{i,c,f,t})$
	(1)	(2)	(3)
$Sentiment_{c,t-1}$	-0.0787***	-0.0207***	-0.0124***
	(0.0130)	(0.0030)	(0.0017)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.1536***	0.0361***	0.0219***
	(0.0205)	(0.0041)	(0.0025)
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Online-Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y
R^2	0.6529	0.7632	0.6448
N	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption, using alternative prompt architectures for the large language model to perform sentiment classification and exclude posts related to product/service price, quality, and price-quality tradeoff. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are the natural logarithms of total consumption amount in week t from columns (1), transaction volume in week t from columns (2), and consumption variety in week t from columns (3), respectively. In Panel B, the dependent variables include the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable $Sentiment_{c,t-1}$ is constructed using GPT-4o-mini with a zero-shot prompt shown in Appendix A.7. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.7 Alternative consumption measure

In our baseline specifications (Equations 1 and 2), the dependent variable is the natural logarithm of consumption. Table A11 assesses the robustness of our results to several alternative functional forms for dependent variables. Columns (1), (4), and (7) employ the first difference of the inverse hyperbolic sine (IHS) transformation of raw consumption. Columns (2), (5), and (8) use indicator variables equal to one if the consumption exceeds the sample mean, and zero otherwise. The remaining columns specify the dependent variables as the one-week-ahead consumption growth rate, defined as $\frac{c_{i,c,f,t}-c_{i,c,f,t-1}}{c_{i,c,f,t-1}}$. Across all specifications, the estimation results consistently indicate that elevated consumer sentiment depresses offline expenditure while simultaneously inducing a reallocation toward online channels. These results confirm that our main findings are robust to the functional form of the consumption variables.

Table A11: Robustness Checks: Alternative Functional Forms for Consumption Variables

Panel A: Overall Consumption

	$Amount_{i,c,t}$			$Transactions_{i,c,t}$			$Variety_{i,c,t}$		
	IHS	Dummy	Growth	IHS	Dummy	Growth	IHS	Dummy	Growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Sentiment_{c,t-1}$	-0.0066	-0.0008	-0.0046	-0.0021	0.0001	0.0016	-0.0009	-0.0013	-0.0010
	(0.0046)	(0.0007)	(0.0028)	(0.0021)	(0.0008)	(0.0019)	(0.0009)	(0.0010)	(0.0010)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.1497	0.5113	0.1492	0.1731	0.6358	0.1721	0.1110	0.4733	0.1110
N	3681478	3681478	3681478	3681478	3681478	3681478	3681478	3681478	3681478

Panel B: Online and Offline Consumption

	$Amount_{i,c,f,t}$			$Transactions_{i,c,f,t}$			$Variety_{i,c,f,t}$		
	IHS	Dummy	Growth	IHS	Dummy	Growth	IHS	Dummy	Growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Sentiment_{c,t-1}$	-0.0239**	-0.0014**	-0.0154***	-0.0062**	-0.0018**	-0.0079***	-0.0060***	-0.0015*	-0.0086***
	(0.0105)	(0.0007)	(0.0037)	(0.0025)	(0.0007)	(0.0029)	(0.0008)	(0.0025)	(0.0025)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0291*	0.0018*	0.0299***	0.0087**	0.0039***	0.0156***	0.0083***	0.0021*	0.0162***
	(0.0169)	(0.0009)	(0.0054)	(0.0036)	(0.0011)	(0.0045)	(0.0012)	(0.0041)	(0.0041)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Online-Consumer-Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.1340	0.5130	0.1674	0.1445	0.6295	0.1809	0.1309	0.6039	0.1755
N	7362956	7362956	7362956	7362956	7362956	7362956	7362956	7362956	7362956

Notes: This table reports the level effects of sentiment on consumption, using alternative functional forms for the consumption variables. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are total consumption amount in week t from columns (1) to (2), transaction volume in week t from columns (3) to (4), and consumption variety in week t from columns (5) to (6), respectively. In Panel B, the dependent variables include the consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. Across these underlying measures, the dependent variable is transformed in three distinct ways: columns (1), (4), and (7) employ the inverse hyperbolic sine (IHS) transformation; columns (2), (5), and (8) construct an indicator variable equal to one if the underlying measure exceeds its sample mean, and zero otherwise; and columns (3), (6), and (9) use the one-week-ahead growth rate. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.8 Alternative sample

Table A12 re-estimates the baseline specifications (Equations 1 and 2) using alternative subsamples. Columns (1), (4), and (6) exclude weeks associated with major shopping festivals in China—Women's Day, the "618" mid-year festival, and "Double 11". The latter two are the most prominent online shopping events in China's e-commerce calendar, occurring in the first and second halves of the year, respectively. Columns (2), (5), and (7) exclude observations from the five cities with the highest GDP per capita in 2019—Beijing, Shanghai, Shenzhen, Guangzhou, and Chongqing. The remaining two columns aggregate individual-level expenditure to the city-week level to estimate the impact of sentiment on aggregate regional expenditure. Across these specifications, the estimation results consistently corroborate our baseline findings, demonstrating that our main conclusions are robust to the exclusion of macroeconomic outliers and alternative sample aggregation schemes.

Table A12: Robustness Checks: Alternative Sample

Panel A: Overall Consumption

	$Log(Amount_{i,c,t})$			$Log(Transactions_{i,c,t})$			$Log(Variety_{i,c,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Sentiment_{c,t-1}$	-0.0027	-0.0009	0.0161	0.0004	0.0010	0.0081	-0.0010	-0.0008
	(0.0035)	(0.0035)	(0.0128)	(0.0018)	(0.0018)	(0.0063)	(0.0008)	(0.0008)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	N	Y	Y	N	Y	Y
Consumer-Month FE	Y	Y	N	Y	Y	N	Y	Y
City FE	N	N	Y	N	N	Y	N	N
City-Month FE	N	N	Y	N	N	Y	N	N
R^2	0.6300	0.6247	0.9702	0.7236	0.7155	0.9913	0.4762	0.4729
N	3518203	2889717	10340	3518203	2889717	10340	3518203	2889717

Panel B: Online and Offline Consumption

	$Log(Amount_{i,c,f,t})$			$Log(Transactions_{i,c,f,t})$			$Log(Variety_{i,c,f,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Sentiment_{c,t-1}$	-0.0106**	-0.0106**	0.0019	-0.0066**	-0.0071***	-0.0042	-0.0030**	-0.0034**
	(0.0051)	(0.0051)	(0.0145)	(0.0027)	(0.0027)	(0.0061)	(0.0015)	(0.0015)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	0.0215***	0.0209***	0.0375*	0.0169***	0.0166***	0.0350***	0.0054**	0.0054**
	(0.0077)	(0.0077)	(0.0212)	(0.0043)	(0.0043)	(0.0116)	(0.0024)	(0.0024)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	N	Y	Y	N	Y	Y
Online-Consumer FE	Y	Y	N	Y	Y	N	Y	Y
Consumer-Month FE	Y	Y	N	Y	Y	N	Y	Y
Online-Consumer-Month FE	Y	Y	N	Y	Y	N	Y	Y
City FE	N	N	Y	N	N	Y	N	N
Online-City FE	N	N	Y	N	N	Y	N	N
City-Month FE	N	N	Y	N	N	Y	N	N
Online-City-Month FE	N	N	Y	N	N	Y	N	N
R^2	0.5059	0.5041	0.9550	0.4929	0.4877	0.9827	0.4717	0.4688
N	7036406	5779434	20680	7036406	5779434	20680	7036406	5779434

Notes: This table reports the level effects of sentiment on consumption, using alternative samples. Panel A presents estimates of Equation (2), and Panel B presents estimates of Equation (1). In Panel A, the dependent variables are the nature logarithm of total consumption amount in week t from columns (1) to (3), the nature logarithm of transaction volume in week t from columns (4) to (6), and the nature logarithm of consumption variety in week t from columns (7) to (8). In Panel B, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The regression samples are modified as follows: columns (1), (4), and (7) exclude major Chinese shopping festivals; columns (2), (5), and (8) exclude the five cities with the highest per capita GDP; and columns (3) and (6) utilize panel data aggregated to the city-week level. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

A.8.9 Instrument variable

A potential threat to our identification strategy is that the text-based sentiment measure might be correlated with unobservable time-varying shocks in the error term, thereby biasing the estimated relationship between sentiment and consumption. To address this endogeneity and isolate the causal effect of sentiment on consumption, we implement an instrumental variable (IV) strategy. Building on the precedent of Benhabib, Wang, and Wen (2015) and Brodeur (2018), who use localized social events as instruments for sentiment, we exploit exogenous variation from three distinct sources.

First, we construct our primary instrumental variable, $Exemplary\ People_{c,t-1}$, an indicator that equals one if city c had at least one local resident featured on the "List of Exemplary People"[16] during the two months preceding week $t-1$, and zero otherwise. Similar to the well-known "Touching China" broadcast, which has been documented to significantly influence public sentiment (Li, Yin, and Zhou (2025)), the "List of Exemplary People" is a prominent state-sponsored civic campaign known to elevate sentiment. A distinct empirical advantage of this event is its higher temporal frequency relative to other national media promotions. Published monthly by China Civilization Online—the official portal of the Central Propaganda Department—the "List of Exemplary People" highlights extraordinary altruistic acts across the country, such as donating blood hundreds of times over decades or risking one's life to rescue a drowning individual. Because the recognition of these incidents is strictly intended to promote civic virtue, the timing and location of these awards are entirely orthogonal to local economic policy adjustments and unobserved consumption shocks, thereby cleanly satisfying the exclusion restriction. Simultaneously, these highly publicized local honors are expected to exogenously induce shifts in the public sentiment of the honoree's home city, thereby fulfilling the instrument relevance criterion.

Second, we construct the second instrumental variable, $Patriotic\ Shock_{c,t-1}$, which captures local public attention to a prominent patriotic campaign in city c during week $t-1$. The theoretical motivation for this instrument relies on the established premise that exogenous geopolitical events or nationalistic shocks can induce widespread variations in social sentiment (Fisman, Hamao, and Wang (2014)). Specifically, we exploit the 2019 nationwide patriotic campaign, "1.4 Billion Flag Guardians." As detailed in Section 2.4, heightened public attention to this event exogenously stimulates local sentiment by evoking strong nationalistic responses, thereby satisfying the instrument relevance condition. Crucially, the genesis of this macro-level political event is strictly orthogonal to local, unobserved economic fundamentals or consumption shocks, satisfying the exclusion restriction. Empirically, we construct $Patriotic\ Shock_{c,t-1}$ as the week-over-week change (from week $t-2$ to week $t-1$) in the Baidu Search Index for the exact keyword "1.4 Billion Flag Guardians" within city c.

Third, following Edmans, Garcia, and Norli (2007), we use soccer results to construct an instrumental variable that generates plausibly exogenous variation in local sentiment. Specifically,

[16] http://www.wenming.cn/

we collect soccer results from several major national soccer competitions in China organized by official authorities, including the Chinese FA Cup, the Chinese Super League, China League One, and China League Two.[17] We define the instrument, $Soccer_{c,t-1}$, as the number of soccer matches in which teams from city c win or draw in week $t-1$, divided by the total number of soccer matches involving teams from city c in that week.

We first present the first-stage estimates from our instrumental variable regressions in Table A13. Across all specifications, $Exemplary\ People_{c,t-1}$, $Patriotic\ Shock_{c,t-1}$, and $Soccer_{c,t-1}$ have significant coefficients and exhibit first-stage F-statistics well above conventional critical values, strictly fulfilling the relevance condition and confirming the strength of our instruments.

Table A14 then reports the corresponding second-stage estimates. The instrument variable is $Exemplary\ People_{c,t-1}$ in Columns (1), (4), and (7), $Patriotic\ Shock_{c,t-1}$ in Columns (2), (5), and (8), and $Soccer_{c,t-1}$ in Columns (3), (6), and (9). Across both identification strategies, the causal estimates strongly corroborate our baseline findings: exogenously elevated sentiment induces a significant reallocation of expenditure, strictly dampening offline transactions while simultaneously stimulating the online channel.

[17] https://www.thecfa.cn/

Table A13: First-Stage Estimates from the Instrumental Variables Regressions

Panel A			
	$Sentiment_{c,t-1}$	$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	$Sentiment_{c,t-1}$
	(1)	(2)	(3)
$Exemplary\ People_{c,t-1}$	-0.0095***	-0.0018***	-0.0095***
	(0.0005)	(0.0001)	(0.0005)
$Exemplary\ People_{c,t-1} \times Online_{i,c,f,t}$	-0.0000	-0.0059***	
	(0.0007)	(0.0005)	
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Quarter FE	Y	Y	Y
Online-Consumer FE	Y	Y	N
Online-Consumer-Quarter FE	Y	Y	N
F-statistic	574.8236	313.4993	287.4118
R^2	0.2953	0.3949	0.2953
N	7222490	7222490	3611245
Panel B			
	$Sentiment_{c,t-1}$	$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	$Sentiment_{c,t-1}$
	(4)	(5)	(6)
$Patriotic\ Shock_{c,t-1}$	-0.0017***	0.0002***	-0.0017***
	(0.0001)	(0.0000)	(0.0002)
$Patriotic\ Shock_{c,t-1} \times Online_{i,c,f,t}$	-0.000	-0.0021***	
	(0.0002)	(0.0001)	
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Quarter FE	Y	Y	Y
Online-Consumer FE	Y	Y	N
Online-Consumer-Quarter FE	Y	Y	N
F-statistic	215.0959	320.8096	107.5479
R^2	0.2953	0.3949	0.2953
N	7222490	7222490	3611245

Panel C

	$Sentiment_{c,t-1}$	$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	$Sentiment_{c,t-1}$
	(1)	(2)	(3)
$Soccer_{c,t-1}$	0.0092***	-0.0016***	0.0159***
	(0.0007)	(0.0001)	(0.0006)
$Soccer_{c,t-1} \times Online_{i,c,f,t}$	0.0000	0.0123***	
	(0.0010)	(0.0007)	
Controls	Y	Y	Y
Week FE	Y	Y	Y
Consumer FE	Y	Y	Y
Consumer-Month FE	Y	Y	Y
Online-Consumer FE	Y	Y	N
Online-Consumer-Month FE	Y	Y	N
F-statistic	121.0466	799.5387	605.2329
R^2	0.2953	0.3949	0.5771
N	7222490	7222490	3611245

Notes: This table reports the first-stage estimates for the instrumental variable regressions. Panel A employs *Exemplary People*$_{c,t-1}$, Panel B uses *Patriotic Shock*$_{c,t-1}$, and Panel C adopts $Soccer_{c,t-1}$ as the instrument. In both panels, columns (1) and (2) report the two first-stage equations required when instrumenting for both $Sentiment_{c,t-1}$ and its interaction with the dummy variable $Online_{i,c,f,t}$, using the respective instrument and its corresponding interaction as excluded instruments. Column (3) reports the single first-stage equation when instrumenting solely for $Sentiment_{c,t-1}$. *Exemplary People*$_{c,t-1}$ is an indicator equal to one if city c featured a resident on the "List of Exemplary People" in the two months prior to week $t-1$. *Patriotic Shock*$_{c,t-1}$ is the week-over-week change in the Baidu Search Index for "1.4 Billion Flag Guardians" in city c. $Soccer_{c,t-1}$ is soccer results for city c in week $t-1$. $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$. The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. The F-statistic corresponds to the F-statistic of the excluded instrument in the respective first-stage regression. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

Table A14: Robustness Checks: Instrumental Variable Regression Results

Panel A: Overall Consumption									
	$Log(Amount_{i,c,t})$			$Log(Transactions_{i,c,t})$			$Log(Variety_{i,c,t})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Sentiment_{c,t-1}$	-0.0210	0.2721	-0.4218	-0.2138	-0.3377	-0.0291	0.0374	0.1781	-0.0011
	(0.3195)	(0.4575)	(0.2883)	(0.1799)	(0.2210)	(0.1473)	(0.0726)	(0.1109)	(0.0739)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Quarter FE	Y	Y	N	Y	Y	N	Y	Y	N
Consumer-Month FE	N	N	Y	N	N	Y	N	N	Y
R^2	0.4959	0.4959	0.6246	0.5961	0.5961	0.7196	0.3120	0.3120	0.4678
N	3,611,245	3,611,245	3,611,245	3,611,245	3,611,245	3,611,245	3,611,245	3,611,245	3,611,245

Panel B: Online and Offline Consumption

	$Log(Amount_{i,c,f,t})$			$Log(Transactions_{i,c,f,t})$			$Log(Variety_{i,c,f,t})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Sentiment_{c,t-1}$	-10.338***	-3.2557***	-1.9937***	-7.8248***	-2.9916***	-0.5824***	-1.4522***	-1.2201***	-0.5924***
	(0.5900)	(0.5760)	(0.4372)	(0.3139)	(0.2984)	(0.2248)	(0.1700)	(0.1664)	(0.1248)
$Sentiment_{c,t-1} \times Online_{i,c,f,t}$	19.418***	7.0433***	3.7104***	14.5350***	5.8782***	1.2286***	2.6573***	2.8847***	1.1060***
	(1.0209)	(0.7320)	(0.6714)	(0.5689)	(0.4058)	(0.3714)	(0.3100)	(0.2239)	(0.2056)
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Online-Consumer FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Consumer-Quarter FE	Y	Y	N	Y	Y	N	Y	Y	N
Online-Consumer-Quarter FE	Y	Y	N	Y	Y	N	Y	Y	N
Consumer-Month FE	N	N	Y	N	N	Y	N	N	Y
Online-Consumer-Month FE	N	N	Y	N	N	Y	N	N	Y
R^2	0.3392	0.3391	0.4986	0.3214	0.3214	0.4856	0.2954	0.2954	0.4635
N	7,222,490	7,222,490	7,222,490	7,222,490	7,222,490	7,222,490	7,222,490	7,222,490	7,222,490

Notes: This table reports the second-stage instrumental variable estimates for the effect of sentiment on consumption. In Panel A, the dependent variables are the natural logarithm of total consumption amount in week t from columns (1) to (3), the natural logarithm of transaction volume in week t from columns (4) to (6), and the natural logarithm of consumption variety in week t from columns (7) to (9). In Panel B, the dependent variables are the natural logarithms of consumption amount, transaction volume, and consumption variety by channel f in week t, where $f = 0$ denotes the offline channel and $f = 1$ denotes the online channel. The explanatory variable, $Sentiment_{c,t-1}$, represents the sentiment level in city c during week $t-1$, and is instrumented by $Exemplary\ People_{c,t-1}$ in columns (1), (4), and (7), $Patriotic\ Shock_{c,t-1}$ in columns (2), (5), and (8), and $Soccer_{c,t-1}$ in columns (3), (6), and (9). The indicator variable $Online_{i,c,f,t}$ equals one if the dependent variable corresponds to online consumption and zero otherwise. Controls include temperature, humidity, wind speed, sunshine, and the air quality index, measured at the city-week level. Standard errors clustered at the consumer level are reported in parentheses. Coefficients marked with *, **, and *** are significant at 10%, 5%, and 1%, respectively.

B Appendix for the Model

This appendix provides detailed proofs for propositions in Section 3.

B.1 Proof of Proposition 1

We solve the consumer's problem in two stages: an intratemporal allocation of total expenditure across goods, followed by an intertemporal consumption-saving decision.

Intratemporal optimization. For a given realization of the matching values γ_{jt} and channel choices $\mathbf{1}_{jt}$, the consumption bundle is given by (6). For a given total expenditure $X_t \equiv \int_0^1 P_{jt}C_{jt}dj$, minimizing expenditure to achieve a given level of C_t yields the standard CES demand functions:

$$C_{jt} = \left[\mathbf{1}_{jt} \cdot \gamma_{jt} + (1 - \mathbf{1}_{jt})\right]^{\theta} \left(\frac{P_{jt}}{P_t}\right)^{-\theta} C_t, \tag{19}$$

where the aggregate price index is

$$P_t^{1-\theta} = \int_0^1 \left[\mathbf{1}_{jt} \cdot \gamma_{jt} + (1 - \mathbf{1}_{jt})\right]^{\theta} P_{jt}^{1-\theta} dj. \tag{20}$$

It follows that $\int_0^1 P_{jt}C_{jt}dj = P_tC_t$. Using similar logic, we have $\int_0^1 P_{jt+1}C_{jt+1}dj = P_{t+1}C_{t+1}$.

Intertemporal optimization. Using this result, the budget constraints (8) and (9) become

$$P_tC_t + S_t = Y, \tag{21}$$

$$P_{t+1}C_{t+1} = RS_t. \tag{22}$$

Substituting these into the utility function (5), the consumer's problem reduces to

$$\max_{S_t} \log\left(\frac{Y - S_t}{P_t}\right) + \beta \log\left(\frac{RS_t}{P_{t+1}}\right) + u_t. \tag{23}$$

where u_t will later be shown to be independent of C_t.

The first-order condition with respect to S_t yields $1/(Y - S_t) = \beta/S_t$, which implies $S_t = \frac{\beta}{1+\beta}Y$. Substituting back, we obtain total expenditure in period 't':

$$P_tC_t = \frac{1}{1+\beta}Y. \tag{24}$$

Therefore, the intertemporal multiplier for the budget constraint (8) is $\Lambda_t = 1/(P_tC_t)$, which satisfies the Euler equation $1/(C_tP_t) = \beta R/(C_{t+1}P_{t+1})$. Thus we have

$$P_{t+1}C_{t+1} = \frac{\beta R}{1+\beta}Y. \tag{25}$$

From the above equation, total expenditure depends only on the household's income Y and is independent of consumer sentiment.

Substituting (24) and (25) into (11), (12), (13) and (14), we have

$$
\begin{aligned}
u^o_{jt} &= \frac{1}{1+\beta}\big(v^o_j \cdot \eta^\alpha_t - z^o_j\big)Y, && (26)\\
u^f_{jt} &= \frac{1}{1+\beta}\big(v^f_j \cdot \eta^\alpha_t - z^f_j\big)Y, && (27)\\
u^o_{jt+1} &= \frac{\beta R}{1+\beta}\big(v^o_j - z^o_j\big)Y, && (28)\\
u^f_{jt+1} &= \frac{\beta R}{1+\beta}\big(v^f_j - z^f_j\big)Y. && (29)
\end{aligned}
$$

Substituting all the above expressions into u_t confirms that u_t is independent of C_t.

B.2 Proof of Proposition 2

From the CES consumption aggregator (6), the demand for type j within each period is given by the standard CES form:

$$
\begin{aligned}
C_{jt} &= [\mathbf{1}_{jt}\gamma_{jt} + (1-\mathbf{1}_{jt})]^\theta \left(\frac{P_{jt}}{P_t}\right)^{-\theta} C_t, && (30)\\
C_{jt+1} &= [\mathbf{1}_{jt+1}\gamma_{jt+1} + (1-\mathbf{1}_{jt+1})]^\theta \left(\frac{P_{jt+1}}{P_{t+1}}\right)^{-\theta} C_{t+1}. && (31)
\end{aligned}
$$

The corresponding price indices are:

$$
\begin{aligned}
P_t &= \left[\int_0^1 \big[\mathbf{1}_{jt}\gamma^\theta_{jt} + (1-\mathbf{1}_{jt})\big] P^{1-\theta}_{jt} dj\right]^{\frac{1}{1-\theta}}, && (32)\\
P_{t+1} &= \left[\int_0^1 \big[\mathbf{1}_{jt+1}\gamma^\theta_{jt+1} + (1-\mathbf{1}_{jt+1})\big] P^{1-\theta}_{jt+1} dj\right]^{\frac{1}{1-\theta}}, && (33)
\end{aligned}
$$

which imply that the price index faced by an individual consumer could be affected by her shopping channel choices.

Using (24) and (25), the consumer's utility $U_t = \log(C_t) + \beta\log(C_{t+1}) + u_t$ can be expressed as:

$$
\begin{aligned}
U_t &= \text{constant} - \log(P_t) - \beta\log(P_{t+1}) + u_t\\
&= \frac{\log\left(\int_0^1 \big[\mathbf{1}_{jt}\gamma^\theta_{jt} + (1-\mathbf{1}_{jt})\big] P^{1-\theta}_{jt} dj\right)}{\theta - 1} + \frac{\beta\log\left(\int_0^1 \big[\mathbf{1}_{jt+1}\gamma^\theta_{jt+1} + (1-\mathbf{1}_{jt+1})\big] P^{1-\theta}_{jt+1} dj\right)}{\theta - 1}\\
&\quad + \int_0^1 \left[\left(\mathbf{1}_{jt}u^o_{jt} + \mathbf{1}_{jt+1}u^o_{jt+1}\right)^\omega + \left((1-\mathbf{1}_{jt})u^f_{jt} + (1-\mathbf{1}_{jt+1})u^f_{jt+1}\right)^\omega\right] dj + \text{constant}. \qquad (34)
\end{aligned}
$$

Then we solve the shopping channel choice for type j good backward. In period $t+1$, the

consumer prefers to purchase type j good online if and only if

$$\begin{aligned}\Delta V_{jt+1} &= \frac{\beta}{\theta-1}\frac{P_{jt+1}^{1-\theta}}{P_t^{1-\theta}}(\gamma_{jt+1}^{\theta}-1)P_{jt+1}^{1-\theta} \\ &\quad +\Big[(\mathbf{1}_{jt}u_{jt}^{o}+u_{jt+1}^{o})^{\omega}+((1-\mathbf{1}_{jt})u_{jt}^{f})^{\omega}-(\mathbf{1}_{jt}u_{jt}^{o})^{\omega}-((1-\mathbf{1}_{jt})u_{jt}^{f}+u_{jt+1}^{f})^{\omega}\Big] \quad (35)\\ &\geq 0. \quad (36)\end{aligned}$$

This condition implies that there is a threshold of γ_{jt+1}^{*}, above which the consumer purchase the good online and offline otherwise. This threshold γ_{jt+1}^{*} depends on $\mathbf{1}_{jt}$. We denote γ_{jt+1}^{o*} for $\mathbf{1}_{jt}=1$ and γ_{jt+1}^{f*} for $\mathbf{1}_{jt}=0$.

Then we consider the consumer in period t. Considering that there is a probability $F_j(\gamma_{jt+1}^{*})$ that the consumer purchase the good offline and the probability $1-F_j(\gamma_{jt+1}^{*})$ online in period $t+1$, the net benefit of purchasing the good online is then given by

$$\begin{aligned}\Delta V_{jt} &= \frac{1}{\theta-1}\frac{P_{jt}^{1-\theta}}{P_t^{1-\theta}}(\gamma_{jt}^{\theta}-1) \\ &\quad +\left[1-F_j\left(\gamma_{jt+1}^{o*}\right)\right]\left(u_{jt}^{o}+u_{jt+1}^{o}\right)^{\omega}+F_j\left(\gamma_{jt+1}^{o*}\right)\left(u_{jt}^{o}\right)^{\omega}+F_j\left(\gamma_{jt+1}^{o*}\right)\left(u_{jt+1}^{f}\right)^{\omega} \\ &\quad -\left[1-F_j\left(\gamma_{jt+1}^{f*}\right)\right]\left(u_{jt+1}^{o}\right)^{\omega}-\left[1-F_j\left(\gamma_{jt+1}^{f*}\right)\right]\left(u_{jt}^{f}\right)^{\omega}-F_j\left(\gamma_{jt+1}^{f*}\right)\left(u_{jt}^{f}+u_{jt+1}^{f}\right)^{\omega}, \quad (37)\end{aligned}$$

which defines a threshold γ_{jt}^{*} for γ_{jt}, above which the consumer purchase the good online and offline otherwise.

The optimal channel choices are therefore:

$$\mathbf{1}_{jt}=1 \text{ iff } \gamma_{jt}\geq\gamma_{jt}^{*}; \mathbf{1}_{jt+1}=1 \text{ iff } \gamma_{jt+1}\geq\gamma_{jt+1}^{*}, \quad (38)$$

which is the result stated in Proposition 2. The thresholds γ_{jt}^{*} and γ_{jt+1}^{*} depend on the parameters regarding the shopping utilities $\{v_j^{o}, v_j^{f}, z_j^{o}, z_j^{f}\}$ as well as the consumer sentiment η_t.

To do 1: (i) Setting $\Delta V_{jt+1}=0$ where the expression of ΔV_{jt+1} is given by (35). Then this equation is an implicit function of γ_{jt+1}^{*}. Derive the partial derivative $\frac{\partial\gamma_{jt+1}^{*}}{\partial\eta_t}$ when $\mathbf{1}_{jt}=1$. Derive the partial derivative $\frac{\partial\gamma_{jt+1}^{*}}{\partial\eta_t}$ when $\mathbf{1}_{jt}=0$.

(ii) Check the sign of $\frac{\partial\gamma_{jt+1}^{*}}{\partial\eta_t}$ when $\mathbf{1}_{jt}=1$ and $\mathbf{1}_{jt}=0$.

Solution to (i) and (ii):

The threshold γ_{jt+1}^{*} is implicitly defined by $\Delta V_{jt+1}=0$ in (35). Let

$$K\equiv\frac{\beta}{\theta-1}\frac{(P_{jt+1}^{1-\theta})^2}{P_t^{1-\theta}}>0. \quad (39)$$

Case $\mathbf{1}_{jt} = 1$ (online today). Denote the threshold by γ^{o*}_{jt+1}. Setting $\mathbf{1}_{jt} = 1$ in (35) and $\Delta V_{jt+1} = 0$ gives:

$$K\left[(\gamma^{o*}_{jt+1})^{\theta} - 1\right] + (u^{o}_{jt} + u^{o}_{jt+1})^{\omega} + (0)^{\omega} - (u^{o}_{jt})^{\omega} - (u^{f}_{jt+1})^{\omega} = 0,$$

which simplifies to:

$$K\left[(\gamma^{o*}_{jt+1})^{\theta} - 1\right] = (u^{o}_{jt})^{\omega} + (u^{f}_{jt+1})^{\omega} - (u^{o}_{jt} + u^{o}_{jt+1})^{\omega}. \tag{40}$$

By the normalization $z^{o}_{j} = 0$ from the main text, we have $u^{o}_{jt} = \frac{1}{1+\beta} v^{o}_{j} \eta^{\alpha}_{t} Y$, $u^{o}_{jt+1} = \frac{\beta R}{1+\beta} v^{o}_{j} Y$, and $u^{f}_{jt+1} = \frac{\beta R}{1+\beta}(v^{f}_{j} - z^{f}_{j})Y$. The right-hand side of (40) depends on η_t only through u^{o}_{jt}.

Define the function $G^{o}(\eta_t) \equiv (u^{o}_{jt})^{\omega} + (u^{f}_{jt+1})^{\omega} - (u^{o}_{jt} + u^{o}_{jt+1})^{\omega}$. Differentiating with respect to η_t:

$$\frac{\partial G^{o}}{\partial \eta_t} = \omega \frac{\partial u^{o}_{jt}}{\partial \eta_t}\left[(u^{o}_{jt})^{\omega-1} - (u^{o}_{jt} + u^{o}_{jt+1})^{\omega-1}\right], \tag{41}$$

where $\frac{\partial u^{o}_{jt}}{\partial \eta_t} = \frac{1}{1+\beta} v^{o}_{j} \alpha \eta^{\alpha-1}_{t} Y > 0$. Since $\omega \in (0,1)$, the function $z \mapsto z^{\omega-1}$ is strictly decreasing. Because $u^{o}_{jt} + u^{o}_{jt+1} > u^{o}_{jt}$, we have $(u^{o}_{jt})^{\omega-1} > (u^{o}_{jt} + u^{o}_{jt+1})^{\omega-1}$. Hence the bracket is positive, so $\partial G^{o}/\partial \eta_t > 0$.

Applying the implicit function theorem to (40):

$$\frac{\partial \gamma^{o*}_{jt+1}}{\partial \eta_t} = \frac{\partial G^{o}/\partial \eta_t}{K\theta(\gamma^{o*}_{jt+1})^{\theta-1}} > 0. \tag{42}$$

Thus $\partial \gamma^{o*}_{jt+1}/\partial \eta_t > 0$: when the consumer shops online today, a positive sentiment shock raises the future threshold, making online shopping in $t+1$ less attractive.

Case $\mathbf{1}_{jt} = 0$ (offline today). Denote the threshold by γ^{f*}_{jt+1}. Setting $\mathbf{1}_{jt} = 0$ in (35) and $\Delta V_{jt+1} = 0$ gives:

$$K\left[(\gamma^{f*}_{jt+1})^{\theta} - 1\right] + (u^{o}_{jt+1})^{\omega} + (u^{f}_{jt})^{\omega} - (0)^{\omega} - (u^{f}_{jt} + u^{f}_{jt+1})^{\omega} = 0,$$

which simplifies to:

$$K\left[(\gamma^{f*}_{jt+1})^{\theta} - 1\right] = (u^{f}_{jt} + u^{f}_{jt+1})^{\omega} - (u^{o}_{jt+1})^{\omega} - (u^{f}_{jt})^{\omega}. \tag{43}$$

Now u^{f}_{jt} depends on η_t (see (27)), whereas u^{o}_{jt+1}, u^{f}_{jt+1} are independent of η_t. Define the right-

hand side of (43) as $G^f(\eta_t) \equiv (u^f_{jt} + u^f_{jt+1})^\omega - (u^o_{jt+1})^\omega - (u^f_{jt})^\omega$. Differentiating:

$$\begin{aligned}\frac{\partial G^f}{\partial \eta_t} &= \omega(u^f_{jt} + u^f_{jt+1})^{\omega-1}\frac{\partial u^f_{jt}}{\partial \eta_t} - \omega(u^f_{jt})^{\omega-1}\frac{\partial u^f_{jt}}{\partial \eta_t} \\ &= \omega\frac{\partial u^f_{jt}}{\partial \eta_t}\Big[(u^f_{jt} + u^f_{jt+1})^{\omega-1} - (u^f_{jt})^{\omega-1}\Big], \end{aligned} \tag{44}$$

where $\frac{\partial u^f_{jt}}{\partial \eta_t} = \frac{Y}{1+\beta} v^f_j \alpha \eta_t^{\alpha-1} > 0$ (since $v^f_j > 0$).

Since $\omega \in (0,1)$, $\omega - 1 < 0$, so the function $x \mapsto x^{\omega-1}$ is strictly decreasing for $x > 0$. The condition $v^f_j > z^f_j$ (positive shopping utility) implies $u^f_{jt+1} > 0$ (see (29)), hence $u^f_{jt} + u^f_{jt+1} > u^f_{jt}$. Consequently,

$$(u^f_{jt} + u^f_{jt+1})^{\omega-1} < (u^f_{jt})^{\omega-1}, \tag{45}$$

so the bracketed term in (44) is negative. Thus $\partial G^f / \partial \eta_t < 0$.

Applying the implicit function theorem to (43):

$$\frac{\partial \gamma^{f*}_{jt+1}}{\partial \eta_t} = \frac{\partial G^f / \partial \eta_t}{K\theta(\gamma^{f*}_{jt+1})^{\theta-1}} < 0. \tag{46}$$

When the consumer shops offline today, a positive sentiment shock *lowers* the future threshold γ^{f*}_{jt+1}, making online shopping in $t+1$ more attractive. Intuitively, higher sentiment raises the offline shopping utility u^f_{jt} in period t; due to the concavity of the shopping utility aggregator ($\omega < 1$), an additional offline purchase in $t+1$ yields diminishing marginal utility, so the consumer becomes more inclined to switch to the online channel in $t+1$.

To do 2: (i) Setting $\Delta V_{jt} = 0$ where the expression of ΔV_{jt} is given by (37). Then this equation is an implicit function of γ^*_{jt}. Derive the partial derivative $\frac{\partial \gamma^*_{jt}}{\partial \eta_t}$.

(ii) Check the sign of $\frac{\partial \gamma^*_{jt}}{\partial \eta_t}$.

Solution to (i) and (ii):

The threshold γ^*_{jt} is implicitly defined by $\Delta V_{jt} = 0$ in (37). Let

$$M \equiv \frac{1}{\theta - 1}\frac{P^{1-\theta}_{jt}}{P^{1-\theta}_t} > 0, \tag{47}$$

which does not depend on η_t (prices are independent of the individual consumer's sentiment). Write $\Delta V_{jt} = M\big[(\gamma_{jt})^\theta - 1\big] + \Psi(\eta_t)$, where $\Psi(\eta_t)$ collects the shopping-utility terms from (37).

By the implicit function theorem,

$$\frac{\partial \gamma^*_{jt}}{\partial \eta_t} = -\frac{\partial \Delta V_{jt}/\partial \eta_t}{\partial \Delta V_{jt}/\partial \gamma_{jt}}\bigg|_{\gamma_{jt}=\gamma^*_{jt}}. \tag{48}$$

The denominator is straightforward:

$$\frac{\partial \Delta V_{jt}}{\partial \gamma_{jt}} = M\theta(\gamma_{jt})^{\theta-1} > 0, \tag{49}$$

since $M > 0$, $\theta > 1$, and $\gamma_{jt} > 0$.

For the numerator, $\partial M/\partial \eta_t = 0$, so

$$\frac{\partial \Delta V_{jt}}{\partial \eta_t} = \frac{\partial \Psi}{\partial \eta_t}. \tag{50}$$

From (37), $\Psi(\eta_t)$ is given by:

$$\begin{aligned}\Psi(\eta_t) \equiv\ & \big[1 - F_j(\gamma^{o*}_{jt+1})\big](u^o_{jt} + u^o_{jt+1})^\omega + F_j(\gamma^{o*}_{jt+1})(u^o_{jt})^\omega + F_j(\gamma^{o*}_{jt+1})(u^f_{jt+1})^\omega \\ & - \big[1 - F_j(\gamma^{f*}_{jt+1})\big](u^o_{jt+1})^\omega - \big[1 - F_j(\gamma^{f*}_{jt+1})\big](u^f_{jt})^\omega - F_j(\gamma^{f*}_{jt+1})(u^f_{jt} + u^f_{jt+1})^\omega. \end{aligned} \tag{51}$$

The terms u^o_{jt+1} and u^f_{jt+1} are independent of η_t (see (28)–(29)). However, u^f_{jt} now depends on η_t through η_t^α (see (27)), and from (46) we have $\partial \gamma^{f*}_{jt+1}/\partial \eta_t < 0$. The dependence of Ψ on η_t therefore operates through four channels: (a) a *direct* channel via u^o_{jt}, (b) a *direct* channel via u^f_{jt}, (c) an *indirect* channel via γ^{o*}_{jt+1}, and (d) an *indirect* channel via γ^{f*}_{jt+1}.

Differentiating (51) with respect to η_t, and denoting $f^o_{t+1} \equiv f_j(\gamma^{o*}_{jt+1}) > 0$, $F^o_{t+1} \equiv F_j(\gamma^{o*}_{jt+1})$, $\Gamma'_o \equiv \partial \gamma^{o*}_{jt+1}/\partial \eta_t > 0$, and analogously $f^f_{t+1} \equiv f_j(\gamma^{f*}_{jt+1}) > 0$, $F^f_{t+1} \equiv F_j(\gamma^{f*}_{jt+1})$, $\Gamma'_f \equiv \partial \gamma^{f*}_{jt+1}/\partial \eta_t < 0$:

$$\begin{aligned}\frac{\partial \Psi}{\partial \eta_t} =\ & -f^o_{t+1}\Gamma'_o(u^o_{jt} + u^o_{jt+1})^\omega + (1 - F^o_{t+1})\,\omega(u^o_{jt} + u^o_{jt+1})^{\omega-1}\frac{\partial u^o_{jt}}{\partial \eta_t} \\ & + f^o_{t+1}\Gamma'_o(u^o_{jt})^\omega + F^o_{t+1}\,\omega(u^o_{jt})^{\omega-1}\frac{\partial u^o_{jt}}{\partial \eta_t} \\ & + f^o_{t+1}\Gamma'_o(u^f_{jt+1})^\omega + f^f_{t+1}\Gamma'_f(u^o_{jt+1})^\omega + f^f_{t+1}\Gamma'_f(u^f_{jt})^\omega - (1 - F^f_{t+1})\,\omega(u^f_{jt})^{\omega-1}\frac{\partial u^f_{jt}}{\partial \eta_t} \\ & - f^f_{t+1}\Gamma'_f(u^f_{jt} + u^f_{jt+1})^\omega - F^f_{t+1}\,\omega(u^f_{jt} + u^f_{jt+1})^{\omega-1}\frac{\partial u^f_{jt}}{\partial \eta_t}, \end{aligned} \tag{52}$$

Grouping the Γ'_o and Γ'_f terms and the $\partial u^o_{jt}/\partial \eta_t$ and $\partial u^f_{jt}/\partial \eta_t$ terms:

$$\begin{aligned}\frac{\partial \Psi}{\partial \eta_t} =\ & f^o_{t+1}\Gamma'_o\Big[(u^o_{jt})^\omega - (u^o_{jt} + u^o_{jt+1})^\omega + (u^f_{jt+1})^\omega\Big] \\ & + \omega\frac{\partial u^o_{jt}}{\partial \eta_t}\Big[(1 - F^o_{t+1})(u^o_{jt} + u^o_{jt+1})^{\omega-1} + F^o_{t+1}(u^o_{jt})^{\omega-1}\Big] \\ & + f^f_{t+1}\Gamma'_f\Big[(u^o_{jt+1})^\omega + (u^f_{jt})^\omega - (u^f_{jt} + u^f_{jt+1})^\omega\Big] \\ & - \omega\frac{\partial u^f_{jt}}{\partial \eta_t}\Big[(1 - F^f_{t+1})(u^f_{jt})^{\omega-1} + F^f_{t+1}(u^f_{jt} + u^f_{jt+1})^{\omega-1}\Big]. \end{aligned} \tag{53}$$

Now we evaluate the sign of $\partial \Delta V_{jt}/\partial \eta_t = \partial \Psi/\partial \eta_t$.

Direct effect via u^o_{jt}. All components in the term involving $\partial u^o_{jt}/\partial \eta_t$ are positive: $\omega > 0$, $\partial u^o_{jt}/\partial \eta_t > 0$, $F^o_{t+1} \in (0,1)$, and the power terms $(u^o_{jt} + u^o_{jt+1})^{\omega-1}$ and $(u^o_{jt})^{\omega-1}$ are positive. Hence this term is strictly positive.

Direct effect via u^f_{jt}. The term involving $\partial u^f_{jt}/\partial \eta_t$ is preceded by a negative sign. With $\partial u^f_{jt}/\partial \eta_t > 0$, $(u^f_{jt})^{\omega-1} > 0$, and $(u^f_{jt} + u^f_{jt+1})^{\omega-1} > 0$, this term is strictly negative. Hence the direct effect via u^f_{jt} (sentiment raising the offline shopping utility today) *lowers* ΔV_{jt}, making online shopping less attractive today.

Indirect effect through γ^{o*}_{jt+1}. The factor $f^o_{t+1}\Gamma'_o > 0$. The sign of the bracket $\left[(u^o_{jt})^\omega - (u^o_{jt} + u^o_{jt+1})^\omega + (u^f_{jt+1})^\omega\right]$ determines the sign.

Since $u^o_{jt+1} > 0$ (by $v^o_j > z^o_j = 0$), we have $u^o_{jt} + u^o_{jt+1} > u^o_{jt}$, and with $\omega \in (0,1)$ the function $x \mapsto x^\omega$ is strictly increasing, so $(u^o_{jt} + u^o_{jt+1})^\omega > (u^o_{jt})^\omega$. Hence $(u^o_{jt})^\omega - (u^o_{jt} + u^o_{jt+1})^\omega < 0$.

Without the normalization $v^f_j = 0$, we have $u^f_{jt+1} = \frac{\beta R}{1+\beta}(v^f_j - z^f_j)Y$. Under the assumption $v^f_j > z^f_j$ (positive shopping utility), $u^f_{jt+1} > 0$, and therefore $(u^f_{jt+1})^\omega > 0$. The bracket coincides with $G^o(\eta_t)$ defined in (40), which satisfies $G^o = K\left[(\gamma^{o*}_{jt+1})^\theta - 1\right]$. Hence the sign of the bracket is the sign of $(\gamma^{o*}_{jt+1})^\theta - 1$: it is positive if $\gamma^{o*}_{jt+1} > 1$ and negative if $\gamma^{o*}_{jt+1} < 1$. Consequently, the indirect effect through γ^{o*}_{jt+1} is positive when $\gamma^{o*}_{jt+1} > 1$ and negative when $\gamma^{o*}_{jt+1} < 1$.

Indirect effect through γ^{f*}_{jt+1}. The factor $f^f_{t+1}\Gamma'_f < 0$. The bracket $\left[(u^o_{jt+1})^\omega + (u^f_{jt})^\omega - (u^f_{jt} + u^f_{jt+1})^\omega\right]$ coincides with $-G^f(\eta_t)$. Since $G^f = K\left[(\gamma^{f*}_{jt+1})^\theta - 1\right]$ from (43), the bracket equals $K\left[1 - (\gamma^{f*}_{jt+1})^\theta\right]$, whose sign is the sign of $1 - \gamma^{f*}_{jt+1}$. With $f^f_{t+1}\Gamma'_f < 0$, the indirect effect through γ^{f*}_{jt+1} is therefore positive when $\gamma^{f*}_{jt+1} > 1$ and negative when $\gamma^{f*}_{jt+1} < 1$.

Overall sign. The numerator $\partial \Delta V_{jt}/\partial \eta_t$ is: (i) the direct effect via u^o_{jt} (> 0), (ii) the direct effect via u^f_{jt} (< 0), (iii) the indirect effect via γ^{o*}_{jt+1} (sign = sign of $(\gamma^{o*}_{jt+1})^\theta - 1$), and (iv) the indirect effect via γ^{f*}_{jt+1} (sign = sign of $1 - \gamma^{f*}_{jt+1}$). If $\gamma^{o*}_{jt+1} > 1$ and $\gamma^{f*}_{jt+1} > 1$, both indirect effects are positive, and the net sign depends on whether the sum of the two positive indirect effects and the positive direct effect via u^o_{jt} dominates the negative direct effect via u^f_{jt}. If either threshold is below 1, the corresponding indirect effect becomes negative, increasing the ambiguity.

Substituting into the implicit function theorem (48) and noting that the denominator $M\theta(\gamma^*_{jt})^{\theta-1} > 0$, we obtain:

$$\operatorname{sgn}\left(\frac{\partial \gamma^*_{jt}}{\partial \eta_t}\right) = -\operatorname{sgn}\left(\frac{\partial \Delta V_{jt}}{\partial \eta_t}\right). \tag{54}$$

The sign of $\partial \gamma^*_{jt}/\partial \eta_t$ depends on the relative magnitudes of the four channels. When both indirect effects are positive (i.e., $\gamma^{o*}_{jt+1} > 1$ and $\gamma^{f*}_{jt+1} > 1$) and the direct effects via u^o_{jt} and u^f_{jt} do not overturn them, $\partial \Delta V_{jt}/\partial \eta_t > 0$ and $\partial \gamma^*_{jt}/\partial \eta_t < 0$: a positive sentiment shock lowers the current-period threshold, making online purchase more likely today. When the direct effect via u^f_{jt} or negative indirect effects dominate, the sign may reverse.

By the normalization $z^o_j = 0$ and the explicit expressions for the shopping utilities and their

derivatives, we can write the numerator in closed form:

$$\begin{aligned}\frac{\partial \Delta V_{jt}}{\partial \eta_t} &= f^o_{t+1}\Gamma'_o\Big[\big(\tfrac{1}{1+\beta}v^o_j\eta^\alpha_t Y\big)^\omega - \big(\tfrac{1}{1+\beta}v^o_j\eta^\alpha_t Y + \tfrac{\beta R}{1+\beta}v^o_j Y\big)^\omega + \big(\tfrac{\beta R}{1+\beta}(v^f_j - z^f_j)Y\big)^\omega\Big] \\ &+\omega\tfrac{Y}{1+\beta}v^o_j\alpha\eta^{\alpha-1}_t\Big[(1-F^o_{t+1})\big(\tfrac{1}{1+\beta}v^o_j\eta^\alpha_t Y + \tfrac{\beta R}{1+\beta}v^o_j Y\big)^{\omega-1} + F^o_{t+1}\big(\tfrac{1}{1+\beta}v^o_j\eta^\alpha_t Y\big)^{\omega-1}\Big] \\ &+f^f_{t+1}\Gamma'_f\Big[\big(\tfrac{\beta R}{1+\beta}v^o_j Y\big)^\omega + \big(\tfrac{1}{1+\beta}(v^f_j\eta^\alpha_t - z^f_j)Y\big)^\omega - \big(\tfrac{1}{1+\beta}(v^f_j\eta^\alpha_t - z^f_j)Y + \tfrac{\beta R}{1+\beta}(v^f_j - z^f_j)Y\big)^\omega\Big] \\ &-\omega\tfrac{Y}{1+\beta}v^f_j\alpha\eta^{\alpha-1}_t\Big[(1-F^f_{t+1})\big(\tfrac{1}{1+\beta}(v^f_j\eta^\alpha_t - z^f_j)Y\big)^{\omega-1} + F^f_{t+1}\big(\tfrac{1}{1+\beta}(v^f_j\eta^\alpha_t - z^f_j)Y + \tfrac{\beta R}{1+\beta}(v^f_j - z^f_j)Y\big)^{\omega-1}\Big],\end{aligned} \tag{55}$$

where $\Gamma'_o = \partial G^o/\partial\eta_t \;/\; \big[K\theta(\gamma^{o*}_{jt+1})^{\theta-1}\big] > 0$ from (42), with $\partial G^o/\partial\eta_t = \omega\frac{Y}{1+\beta}v^o_j\alpha\eta^{\alpha-1}_t\big[(u^o_{jt})^{\omega-1} - (u^o_{jt}+u^o_{jt+1})^{\omega-1}\big] > 0$, and $\Gamma'_f = \partial G^f/\partial\eta_t \;/\; \big[K\theta(\gamma^{f*}_{jt+1})^{\theta-1}\big] < 0$ from (46), with $\partial G^f/\partial\eta_t = \omega\frac{Y}{1+\beta}v^f_j\alpha\eta^{\alpha-1}_t\big[(u^f_{jt}+u^f_{jt+1})^{\omega-1} - (u^f_{jt})^{\omega-1}\big] < 0$. In summary,

$$\frac{\partial\gamma^*_{jt}}{\partial\eta_t}\begin{cases} < 0, & \text{if the sum of the positive components outweighs the negative components,} \\ > 0, & \text{if negative components dominate.}\end{cases} \tag{56}$$

Resolving the sign requires a quantitative evaluation with calibrated parameters.

When both $\gamma^{o*}_{jt+1} > 1$ and $\gamma^{f*}_{jt+1} > 1$, we can strengthen this result and establish $\partial\gamma^*_{jt}/\partial\eta_t < 0$ under empirically plausible conditions. The reasoning proceeds as follows.

Sign of the brackets. From (40) and (43), the brackets in (53) satisfy:

$$B_o \equiv (u^o_{jt})^\omega - (u^o_{jt}+u^o_{jt+1})^\omega + (u^f_{jt+1})^\omega = K\big[(\gamma^{o*}_{jt+1})^\theta - 1\big], \tag{57}$$

$$B_f \equiv (u^o_{jt+1})^\omega + (u^f_{jt})^\omega - (u^f_{jt}+u^f_{jt+1})^\omega = K\big[1 - (\gamma^{f*}_{jt+1})^\theta\big]. \tag{58}$$

Thus $\gamma^{o*}_{jt+1} > 1$ implies $B_o > 0$, and $\gamma^{f*}_{jt+1} > 1$ implies $B_f < 0$.

Sign of the indirect effects. From (42) we have $\Gamma'_o > 0$, so $f^o_{t+1}\Gamma'_o B_o > 0$. From (46) we have $\Gamma'_f < 0$, so $f^f_{t+1}\Gamma'_f B_f = (\text{negative}) \times (\text{negative}) > 0$. Both indirect effects are therefore positive when the respective thresholds exceed 1.

Combining the f-channel terms. The sum of the γ^{f*}_{jt+1} indirect effect and the u^f_{jt} direct effect can be written as:

$$\begin{aligned}S_f \equiv \omega\frac{\partial u^f_{jt}}{\partial\eta_t}\Bigg[&-(u^f_{jt})^{\omega-1} \\ &+\Big(F^f_{t+1} + \frac{f^f_{t+1}}{\theta}\cdot\frac{(\gamma^{f*}_{jt+1})^\theta - 1}{(\gamma^{f*}_{jt+1})^{\theta-1}}\Big)\Big((u^f_{jt})^{\omega-1} - (u^f_{jt}+u^f_{jt+1})^{\omega-1}\Big)\Bigg].\end{aligned} \tag{59}$$

For $\gamma^{f*}_{jt+1} > 1$, the coefficient $C_f \equiv F^f_{t+1} + \frac{f^f_{t+1}}{\theta}\cdot\frac{(\gamma^{f*}_{jt+1})^\theta - 1}{(\gamma^{f*}_{jt+1})^{\theta-1}}$ exceeds F^f_{t+1} by a term that grows with

γ^{f*}_{jt+1}. When γ^{f*}_{jt+1} is sufficiently above 1, $C_f > 1$ and consequently $S_f > 0$. Together with $\gamma^{o*}_{jt+1} > 1$ (which guarantees $f^o_{t+1}\Gamma'_o B_o > 0$) and the always-positive u^o_{jt} direct effect, we have $\partial \Delta V_{jt}/\partial \eta_t > 0$ and hence $\partial \gamma^*_{jt}/\partial \eta_t < 0$.

Intuitively, when both future thresholds exceed 1, the online channel is already relatively attractive in period $t+1$. A positive sentiment shock amplifies this attractiveness through both future channels (raising γ^{o*}_{jt+1} and lowering γ^{f*}_{jt+1}), and these reinforcing indirect effects dominate the countervailing direct effect of higher current offline utility. The result is most cleanly obtained when $\gamma^{o*}_{jt+1} > 1$ and γ^{f*}_{jt+1} is sufficiently above 1, which corresponds to empirically relevant scenarios where online shopping offers a sufficiently large utility advantage.

B.3 Proof of Proposition 3

Step 1: Sign of $\partial \gamma^*_{jt}/\partial \eta_t$. From the threshold condition $\Delta V_{jt} = 0$ in (37), the implicit function theorem yields:

$$\frac{\partial \gamma^*_{jt}}{\partial \eta_t} = -\frac{\partial \Delta V_{jt}/\partial \eta_t}{\partial \Delta V_{jt}/\partial \gamma_{jt}}\bigg|_{\gamma_{jt}=\gamma^*_{jt}}. \tag{60}$$

The denominator is positive: $\partial \Delta V_{jt}/\partial \gamma_{jt} = M\theta(\gamma^*_{jt})^{\theta-1} > 0$. The sign of $\partial \gamma^*_{jt}/\partial \eta_t$ is therefore opposite to the sign of $\partial \Delta V_{jt}/\partial \eta_t$.

As established in the To do 2 solution of Appendix B.2, under the empirically relevant conditions $\gamma^{o*}_{jt+1} > 1$ and $\gamma^{f*}_{jt+1} > 1$, all four channels through which η_t affects ΔV_{jt} are jointly positive, giving $\partial \Delta V_{jt}/\partial \eta_t > 0$. Consequently:

$$\frac{\partial \gamma^*_{jt}}{\partial \eta_t} < 0. \tag{61}$$

Thus a positive sentiment shock lowers the current-period threshold, making an online purchase more likely.

Step 2: Derivative of the online share with respect to the threshold. From (78) and the definition of ω^o_{jt} in Step 3 below, the online expenditure share for type j goods is:

$$\omega^o_{jt} = \frac{G_j(\gamma^*_{jt})}{H_j(\gamma^*_{jt})}, \tag{62}$$

where $G_j(\gamma) \equiv \int_\gamma^\infty \gamma^\theta dF(\gamma)$ and $H_j(\gamma) \equiv G_j(\gamma) + F(\gamma)$. Differentiating with respect to γ^*_{jt}:

$$\begin{aligned}
\frac{d\omega^o_{jt}}{d\gamma^*_{jt}} &= \frac{G'_j(\gamma^*_{jt})H_j(\gamma^*_{jt}) - G_j(\gamma^*_{jt})H'_j(\gamma^*_{jt})}{H_j(\gamma^*_{jt})^2} \\
&= \frac{[-f_j(\gamma^*_{jt})(\gamma^*_{jt})^\theta]H_j - G_j\left[f_j(\gamma^*_{jt})(1-(\gamma^*_{jt})^\theta)\right]}{H_j^2} \\
&= -\frac{f_j(\gamma^*_{jt})\left[G_j(\gamma^*_{jt}) + (\gamma^*_{jt})^\theta F_j(\gamma^*_{jt})\right]}{H_j(\gamma^*_{jt})^2} < 0.
\end{aligned} \tag{63}$$

Thus a decrease in the threshold γ^*_{jt} (as occurs with a positive sentiment shock) raises the online expenditure share.

Step 3: Total effect of sentiment on ω^o_{jt}. Combining the results from Steps 1 and 2 via the chain rule:

$$\frac{\partial \omega^o_{jt}}{\partial \eta_t} = \frac{d\omega^o_{jt}}{d\gamma^*_{jt}} \cdot \frac{\partial \gamma^*_{jt}}{\partial \eta_t} = (\text{negative}) \times (\text{negative}) > 0. \tag{64}$$

Equivalently, substituting the explicit expressions from (63) and (60):

$$\frac{\partial \omega^o_{jt}}{\partial \eta_t} = \frac{f_j(\gamma^*_{jt})\left[G(\gamma^*_{jt}) + (\gamma^*_{jt})^\theta F_j(\gamma^*_{jt})\right]}{H_j(\gamma^*_{jt})^2} \cdot \frac{\partial \Delta V_{jt}/\partial \eta_t}{\partial \Delta V_{jt}/\partial \gamma_{jt}}. \tag{65}$$

Since $\partial \Delta V_{jt}/\partial \eta_t > 0$ (from To do 2 under $\gamma^{o*}_{jt+1} > 1$ and $\gamma^{f*}_{jt+1} > 1$) and $\partial \Delta V_{jt}/\partial \gamma_{jt} = M\theta(\gamma^*_{jt})^{\theta-1} > 0$, all terms are strictly positive. Therefore:

$$\frac{\partial \omega^o_{jt}}{\partial \eta_t} > 0. \tag{66}$$

Thus a positive sentiment shock raises the online expenditure share for each good type j.

Step 4: Aggregate online expenditure share. From (78), the expenditure share of type j goods satisfies:

$$s_{jt} \equiv \frac{P_{jt}C_{jt}}{P_t C_t} = P_{jt}^{1-\theta} P_t^{\theta-1} H_j(\gamma^*_{jt}), \tag{67}$$

where $H_j(\gamma) \equiv \int_\gamma^\infty \gamma^\theta dF(\gamma) + F(\gamma)$. The online expenditure share for type j goods is:

$$\omega^o_{jt} = \frac{\int_{\gamma \geq \gamma^*_{jt}} \gamma^\theta dF(\gamma)}{H_j(\gamma^*_{jt})} = \frac{G_j(\gamma^*_{jt})}{H_j(\gamma^*_{jt})}, \tag{68}$$

where $G_j(\gamma) \equiv \int_\gamma^\infty \gamma^\theta dF(\gamma)$. Combining these two expressions, the product simplifies to:

$$s_{jt}\omega^o_{jt} = P_{jt}^{1-\theta} P_t^{\theta-1} G_j(\gamma^*_{jt}). \tag{69}$$

The aggregate online expenditure share is therefore:

$$\omega^o_t = \int_0^1 s_{jt}\omega^o_{jt} dj = P_t^{\theta-1} \int_0^1 P_{jt}^{1-\theta} G_j(\gamma^*_{jt}) dj. \tag{70}$$

Using the price index $P_t^{1-\theta} = \int_0^1 P_{jt}^{1-\theta} H_j(\gamma_{jt}^*) dj$, we rewrite this as:

$$\omega_t^o = \frac{N}{D}, \quad N \equiv \int_0^1 P_{jt}^{1-\theta} G_j(\gamma_{jt}^*) dj, \quad D \equiv \int_0^1 P_{jt}^{1-\theta} H_j(\gamma_{jt}^*) dj. \tag{71}$$

Differentiating with respect to η_t:

$$\frac{\partial \omega_t^o}{\partial \eta_t} = \frac{N'D - ND'}{D^2}, \tag{72}$$

where:

$$N' \equiv \int_0^1 P_{jt}^{1-\theta} G_j'(\gamma_{jt}^*) \frac{\partial \gamma_{jt}^*}{\partial \eta_t} dj = \int_0^1 P_{jt}^{1-\theta} \left[-f_j(\gamma_{jt}^*)(\gamma_{jt}^*)^\theta \right] \frac{\partial \gamma_{jt}^*}{\partial \eta_t} dj, \tag{73}$$

$$D' \equiv \int_0^1 P_{jt}^{1-\theta} H_j'(\gamma_{jt}^*) \frac{\partial \gamma_{jt}^*}{\partial \eta_t} dj = \int_0^1 P_{jt}^{1-\theta} f_j(\gamma_{jt}^*) \left[1 - (\gamma_{jt}^*)^\theta \right] \frac{\partial \gamma_{jt}^*}{\partial \eta_t} dj. \tag{74}$$

From Step 2, $\partial \gamma_{jt}^* / \partial \eta_t < 0$ for all j. Since $G_j'(\gamma) = -f_j(\gamma)\gamma^\theta < 0$, each integrand in N' is positive, hence $N' > 0$. For $\gamma_{jt}^* > 1$, $H_j'(\gamma) = f_j(\gamma)[1 - \gamma^\theta] < 0$, so each integrand in D' is also positive, implying $D' > 0$.

To sign $N'D - ND'$, note that $\omega_t^o = N/D \in (0,1)$ because $G_j(\gamma) < H_j(\gamma)$ for all γ. Moreover, the integrand of N' exceeds that of D' pointwise for $\gamma_{jt}^* > 1$:

$$f_j(\gamma_{jt}^*)(\gamma_{jt}^*)^\theta \left(-\partial \gamma_{jt}^* / \partial \eta_t \right) > f_j(\gamma_{jt}^*) \left[(\gamma_{jt}^*)^\theta - 1 \right] \left(-\partial \gamma_{jt}^* / \partial \eta_t \right), \tag{75}$$

since $(\gamma_{jt}^*)^\theta > (\gamma_{jt}^*)^\theta - 1 > 0$. Therefore $N' > D'$. Combined with $\omega_t^o < 1$, we obtain $N' > D' > \omega_t^o D'$, i.e. $N' - \omega_t^o D' > 0$. Since $D > 0$ and $N'D - ND' = D(N' - \omega_t^o D')$, we conclude $N'D - ND' > 0$, hence:

$$\frac{\partial \omega_t^o}{\partial \eta_t} > 0, \tag{76}$$

establishing Proposition 3.

B.4 Proof of Proposition 4

Thus the share of expenditure spent on type j goods is given by

$$s_{jt} = \frac{P_{jt} C_{jt}}{P_t C_t} = \int \left[\mathbf{1}_{jt} \gamma^\theta + (1 - \mathbf{1}_{jt}) \right] P_{jt}^{1-\theta} P_t^{\theta - 1} dF(\gamma) \tag{77}$$

$$= P_{jt}^{1-\theta} P_t^{\theta-1} \left[\int_{\gamma_{jt} \geq \gamma_{jt}^*} \gamma^\theta dF(\gamma) + F\left(\gamma_{jt}^*\right) \right] \tag{78}$$

Soon we show that the consumer will purchase the good online when $\gamma_{jt} > \gamma_{jt}^*$ and offline otherwise. Therefore, the average share of expenditure on type j goods is given by

$$\omega_j = \frac{P_j C_j}{P_t C_t} = \left[F_j\left(\gamma_j^*\right) + \left[1 - F_j\left(\gamma_j^*\right)\right]\gamma_{jt}^{\theta}\right] P_{jt}^{1-\theta} P_t^{\theta-1}, \tag{79}$$

where $F_j(\cdot)$ is the cumulative distribution function (CDF) for γ_{jt} and γ_j^* is the steady-state value of γ_{jt}^*.

Proof: We have

$$\frac{\partial^2 S_t^o}{\partial \eta_t^2} = v C_t f'\left(\tilde{\gamma}_t^*\right) \frac{\partial \tilde{\gamma}_t^*}{\partial \eta_t} = -\left(v C_t\right)^2 f'\left(\tilde{\gamma}_t^*\right) < 0 \tag{80}$$

B.5 Proof of Proposition 7

We derive $\frac{\partial^2 S_t^o}{\partial \eta_t \partial C_t}$ here.

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial C_t} = f'\left(\tilde{\gamma}_t^*\right) v C_t \frac{\partial \tilde{\gamma}_t^*}{\partial C_t} + f\left(\tilde{\gamma}_t^*\right) v = -f'\left(\tilde{\gamma}_t^*\right) v^2 \eta_t C_t + f\left(\tilde{\gamma}_t^*\right) v, \tag{81}$$

where the sign of $f''(\tilde{\gamma}_t^*)$ depends on the specific distribution of $\tilde{\gamma}$. However, since v is very small, the sign of $\frac{\partial^2 S_t^o}{\partial \eta_t \partial C_t}$ is dominated by the sign of the second term, $f\left(\tilde{\gamma}_t^*\right) v$, which is positive. Thus, we have

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial C_t} > 0, \tag{82}$$

which immediately follows that

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial Y_t} = \frac{\partial^2 S_t^o}{\partial \eta_t \partial C_t} \frac{\partial C_t}{\partial Y_t} > 0 \tag{83}$$

as $\frac{\partial C_t}{\partial Y_t} > 0$ by Proposition 1. Thus we prove Proposition 7.

B.6 Proof of Proposition 8

We derive $\frac{\partial^2 S_t^o}{\partial \eta_t \partial v}$ now. We have

$$\frac{\partial \tilde{\gamma}_t^*}{\partial v} = -\eta_t C_t. \tag{84}$$

Therefore, we have

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial v} = f'\left(\tilde{\gamma}_t^*\right) v C_t \frac{\partial \tilde{\gamma}_t^*}{\partial v} + f\left(\tilde{\gamma}_t^*\right) C_t = -f'\left(\tilde{\gamma}_t^*\right) v \eta C_t^2 + f\left(\tilde{\gamma}_t^*\right) C_t, \tag{85}$$

where the sign of $f'(\tilde{\gamma}_t^*)$ depends on the distribution of $\tilde{\gamma}$. However, since v is very small, the sign of $\frac{\partial^2 S_t^o}{\partial \eta_t \partial v}$ is dominated by the the sign of the second term, $f\left(\tilde{\gamma}_t^*\right) C_t$, which is positive. Thus, we have

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial v} > 0. \tag{86}$$

Thus we prove Proposition 8.

B.7 Proof of Proposition 5

For durable goods, the consumer's net marginal benefit of shopping type-j good via channel $i \in \{o, f\}$ is given by

$$V_{jt}^{i,d} = \frac{1}{C_t} C_t^{\frac{1}{\theta}} C_{jt}^{-\frac{1}{\theta}} \gamma_{jt}^i + \beta \frac{1}{C_{t+1}} C_{t+1}^{\frac{1}{\theta}} C_{jt+1}^{-\frac{1}{\theta}} \gamma_{jt}^i - P_{jt}\Lambda_t + (v^i \eta_t - z^i) C_t^{\frac{1}{\theta}} C_{jt}^{-\frac{1}{\theta}} \gamma_{jt}^i, \tag{87}$$

where the superscript d denotes "durable".

Thus the consumers compare $V_{jt}^{o,d}$ and $V_{jt}^{f,d}$ and purchase the good online if

$$V_{jt}^{o,d} \geq V_{jt}^{f,d}. \tag{88}$$

Since $\gamma_{jt}^o = \gamma_{jt+1}^o$, $\gamma_{jt}^f = \gamma_{jt+1}^f$ and $P_{jt} = P_{jt+1}$, then we have

$$\frac{C_{jt}}{C_t} = \frac{C_{jt+1}}{C_{t+1}}, \tag{89}$$

for all types of goods $j \in [0, 1]$. Then the condition (88) can be rewritten as

$$\begin{aligned}
\frac{\gamma_{jt}^o}{\gamma_{jt}^f} &\geq \frac{1 + \beta + \eta_t v^f C_t - z^f C_t}{1 + \beta + \eta_t v^o C_t - z^o C_t} && (90)\\
&= \frac{1 + \frac{1}{1+\beta}\left(\eta_t v^f - z^f\right) C_t}{1 + \frac{1}{1+\beta}\left(\eta_t v^o - z^o\right) C_t} && (91)\\
&\approx 1 - \frac{1}{1+\beta}\left[\left(v^o - v^f\right)\eta_t + \left(z^f - z^o\right)\right] C_t && (92)\\
&= 1 - \frac{1}{1+\beta}\left(v\eta_t + z\right) C_t. && (93)
\end{aligned}$$

That is, the threshold $\tilde{\gamma}_t^*$ for durable goods is given by $1 - \frac{1}{1+\beta}\left(v\eta_t + z\right) C_t$ instead of $1 - \left(v\eta_t + z\right) C_t$ as for nondurable goods.

Therefore, for durable goods, we have

$$\frac{\partial \tilde{\gamma}_t^*}{\partial \eta_t} = -\frac{1}{1+\beta} v C_t \tag{94}$$

which causes

$$\frac{\partial S_t^o}{\partial \eta_t} = \frac{\partial \left(1 - F\left(\tilde{\gamma}_t^*\right)\right)}{\partial \eta_t} = -f\left(\tilde{\gamma}_t^*\right) \frac{\partial \tilde{\gamma}_t^*}{\partial \eta_t} = f\left(\tilde{\gamma}_t^*\right) \frac{vC_t}{1+\beta}. \tag{95}$$

We know that $\frac{\partial S_t^o}{\partial \eta_t}$ is smaller for durable goods than for non-durable goods, that is, the response of online expenditure is less sensitive for durable goods than for non-durable goods.

B.8 Proof of Proposition 6

$$\begin{aligned}
\frac{\partial^2 S_t^o}{\partial \eta_t \partial \omega} &= \left[\Gamma'\left(\tilde{\gamma}_t^o\right)\Gamma\left(\tilde{\gamma}_t^f\right)\frac{\partial \tilde{\gamma}_t^o}{\partial \eta_t} - \Gamma'\left(\tilde{\gamma}_t^f\right)\Gamma\left(\tilde{\gamma}_t^o\right)\frac{\partial \tilde{\gamma}_t^f}{\partial \eta_t}\right] \frac{\partial}{\partial \omega_t}\left[\frac{\omega\left(1-\omega\right)}{\left[\omega\Gamma\left(\tilde{\gamma}_t^o\right) + \left(1-\omega\right)\Gamma\left(\tilde{\gamma}_t^f\right)\right]^2}\right] \\
&\approx \left[\Gamma'\left(\tilde{\gamma}_t^o\right)\Gamma\left(\tilde{\gamma}_t^f\right)\frac{\partial \tilde{\gamma}_t^o}{\partial \eta_t} - \Gamma'\left(\tilde{\gamma}_t^f\right)\Gamma\left(\tilde{\gamma}_t^o\right)\frac{\partial \tilde{\gamma}_t^f}{\partial \eta_t}\right] \frac{\left(1-2\omega\right)}{\left[\omega\Gamma\left(\tilde{\gamma}_t^o\right) + \left(1-\omega\right)\Gamma\left(\tilde{\gamma}_t^f\right)\right]^2}.
\end{aligned} \tag{96}$$

Since we have assumed that $\omega \in (0,1)$ which aligns with the data, we have $\frac{\partial^2 S_t^o}{\partial \eta_t \partial \omega} > 0$.
We first show that

$$\frac{\partial \tilde{\gamma}_t^*}{\partial z} = -C_t. \tag{97}$$

Thus, we have

$$\frac{\partial^2 S_t^o}{\partial \eta_t \partial z} = f'\left(\tilde{\gamma}_t^*\right) vC_t \frac{\partial \tilde{\gamma}_t^*}{\partial z} = -f'\left(\tilde{\gamma}_t^*\right) vC_t^2 < 0, \tag{98}$$

which means that in regions with greater accessibility to offline shopping, the response of online shopping is stronger.