Preprints.org

Review

Not peer-reviewed version

Self-Care and Financial Sentiment During Inflationary Trends: An NLP Journey Through the FinTalk-19k Dataset

Tanvi Dogra *

Posted Date: 11 August 2025

doi: 10.20944/preprints202508.0721.v1

Keywords: self-care; personal finance; cost of living; economic stress consumer behavior; BERT-base; Unsloth; BERTopic; sentiment analysis; uncertainty detection; text classification; data leakage; shap explainability; financial well-being; machine learning; Natural language processing (NLP)

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Self-Care and Financial Sentiment During Inflationary Trends: An NLP Journey Through the FinTalk-19k Dataset

Tanvi Dogra * and Ziyuan Huang

Harrisburg University, Masters in Analytics, ANLY 699-50- A - Appld Behavioral Research ProjectANLY

* Correspondence: tanvi.dogra25.td@gmail.com

Abstract

When prices climb and the cost of living feels like a slow squeeze, the way people talk about money changes. It's not just about "what to buy" — it's about how to *cope*. This study takes a closer look at that overlap between finance and self-care, digging into thousands of real posts from the FinTalk-19k dataset. Using a fine-tuned BERT-base model (optimized with Unsloth for speed and efficiency), I classified posts as self-care related or not, explored deeper themes with BERTopic, and checked how sentiment and uncertainty weave through different conversations. To keep things transparent, I used a train–validate–test split (60/20/20) with stratification, fixed seeds for reproducibility, confusion matrices to see exactly where the model messed up, and ROC curves to check performance trade-offs. SHAP explainability helped me peek inside the model's "thought process" — why it made certain calls, and where it got confused. The results were strong on paper: ~93% accuracy, F1 scores around 0.90, and ROC-AUC near 0.96. But I also found possible data leakage (keywords and near-duplicate posts), meaning the model might look smarter than it really is. Crisis-related topics (scams, medical debt) had the lowest sentiment and the most uncertainty, while "good news" topics (lottery wins, tax refunds) were happier but still cautious. Looking ahead, I want to zoom in on just the "Personal Finance" posts and manually tag emotional nuances like sarcasm, masked anxiety, and quiet pride — the kind of things algorithms often miss.

Keywords: self-care; personal finance; cost of living; economic stress consumer behavior; BERT-base; Unsloth; BERTopic; sentiment analysis; uncertainty detection; text classification; data leakage; shap explainability; financial well-being; machine learning; Natural language processing (NLP)

1. Introduction

Inflation has a subtle yet profound way of not only shrinking consumers' wallets but also reshaping the emotional and psychological relationship people have with money. While official statistics such as the U.S. Bureau of Labor Statistics' (BLS, 2022) Consumer Price Index can quantify the changes in prices over time, they fail to capture the underlying mental and emotional recalibrations that occur. During inflationary periods, financial decision-making rarely follows a purely rational, economic model. Instead, it becomes deeply intertwined with coping mechanisms, routines, and self-preservation strategies aimed at preserving a sense of control.

The global self-care industry offers a compelling lens through which to examine this phenomenon. Valued at over USD 450 billion in 2022 and projected to grow at a compound annual growth rate (CAGR) of approximately 9% through 2030 (Grand View Research, 2023), the industry demonstrates remarkable resilience—even during economic downturns. Historical trends suggest that while consumers may shift to more affordable alternatives, they rarely abandon self-care altogether. This is often attributed to the so-called "lipstick effect," where small, affordable luxuries remain in demand as consumers seek comfort and normalcy amid financial uncertainty.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

In my professional experience in the fashion and retail sectors, I have observed firsthand how purchasing behavior adapts when disposable income feels constrained. Consumers often gravitate toward two-in-one products, "dupes" for premium skincare, or rituals that are deliberately marketed as both economical and indulgent. Self-care, in this sense, is highly adaptable in its form but unwavering in its perceived necessity.

This adaptability sparked my interest in exploring how individuals actually talk about self-care in the context of financial discussions online. The core questions that guided this project were: How is self-care framed in financial discourse during periods of inflation? Do stress and uncertainty shift the nature of these conversations? And can natural language processing (NLP) tools uncover meaningful patterns in these discussions without losing the human context that makes them so insightful?

The relationship between self-care and consumer psychology during inflation is not just anecdotal—it is measurable. For example, McKinsey & Company (2022) reported that more than 70% of consumers maintained or increased spending on wellness-related products and services during recent economic slowdowns. These purchases are not merely indulgences but also function as psychological buffers, helping individuals navigate stress and uncertainty. This behavioral persistence suggests that self-care consumption serves both a functional and symbolic role, making it an essential variable in studies of financial sentiment.

Furthermore, the digital transformation of self-care—accelerated by social media platforms, influencer marketing, and online communities—has altered how consumers perceive and discuss these practices. Financial forums and personal finance subreddits often contain subtle overlaps between monetary strategies and self-care routines, blurring the line between economic planning and emotional wellbeing. Such spaces present a valuable data source for understanding how language reflects these intertwined priorities.

By examining these conversations through NLP and sentiment analysis, we can begin to quantify the emotional dimensions of financial discourse. This approach not only captures the surface-level trends but also has the potential to reveal the nuanced ways people maintain dignity, agency, and self-expression in financially restrictive environments. In doing so, this research contributes to a broader understanding of how economic forces shape everyday life beyond the confines of traditional economic indicators.

Theoretical Framework

Understanding how self-care behaviors manifest during inflationary periods requires drawing on several key theories from psychology, economics, and marketing. **Maslow's Hierarchy of Needs** provides one foundation: while self-care often sits at the intersection of physiological and psychological needs, its persistence in spending habits suggests that it fulfills higher-order needs like self-esteem and belonging. In downturns, consumers may sacrifice luxury goods but retain "affordable luxuries" that protect their sense of identity.

The **Lipstick Effect** theory complements this view, proposing that in times of economic stress, consumers shift to smaller, more attainable indulgences. This doesn't just apply to cosmetics—it's observable across wellness services, skincare, and small home rituals. Behavioral economics also adds the concept of **mental accounting**, where individuals allocate specific budgets for emotional well-being even when cutting other expenses.

By integrating these frameworks, we can hypothesize that self-care persists not because it is immune to financial pressure, but because it adapts—transforming in form, price, and marketing.

2. Literature Review

Literature Review

A growing body of research highlights the intersection of economic conditions, consumer behavior, and self-care practices, particularly during periods of inflation. Baker, Bloom, and Davis

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

(2016) introduced the concept of measuring economic policy uncertainty, illustrating how macroeconomic instability can alter consumer confidence and decision-making. Lusardi and Mitchell (2014) further established the importance of financial literacy in navigating such periods, noting that informed consumers are better equipped to adapt their spending without sacrificing essential wellbeing. This is not just about knowing numbers—it's about feeling empowered to make trade-offs that align with one's values and emotional needs.

The resilience of the self-care industry during economic downturns has been well-documented. The Global Wellness Institute (2023) reported sustained growth in the global wellness economy, valued in the trillions, even through the COVID-19 pandemic and subsequent inflationary pressures. This data suggests that self-care is not simply a luxury; it is a perceived necessity for maintaining balance. Hill et al. (2024) examined the "lipstick effect" and found that affordable luxuries maintain their appeal in recessions, serving as emotional coping mechanisms. For example, a $20 skincare serum may feel like a justifiable indulgence when larger vacations or expensive services are cut from the budget. McKinsey & Company (2023) reinforced this view, showing that a significant share of consumers sustain or even increase wellness-related spending despite rising costs of living.

Digital platforms have amplified self-care discourse, creating new avenues for understanding how consumers merge financial planning with emotional wellbeing. Dai et al. (2024) used text analysis to link financial stress to expressions of psychological distress, underscoring the value of language-based methods in studying consumer sentiment. Similarly, Hu, Zhao, and Huang (2019) highlighted the efficacy of sentiment analysis in financial contexts, while Li, Yu, and Huang (2014) demonstrated its predictive power for market risk. These tools allow researchers to go beyond numbers and tap into the lived experiences embedded in online conversations.

Yet, emerging research also points to the complexity of interpreting sentiment during economic stress. Wang, Tse, and Chan (2023) explored multidimensional consumer uncertainty, emphasizing that simple sentiment scores may miss nuances like sarcasm, subtle optimism, or quiet resignation. Vlada, Zimmerman, and Martin (2024) used machine learning to connect economic insecurity with changes in health behavior, showing how budget constraints can reshape diet, exercise, and mental health routines. This aligns with real-world observations: people often adapt their self-care rather than abandon it—opting for home workouts instead of gym memberships or cooking nutrient-rich meals from affordable ingredients.

Taken together, these studies suggest that self-care is both adaptive and symbolic during times of economic uncertainty. Even under constrained budgets, consumers strategically maintain practices that preserve psychological stability and personal identity. Understanding these behaviors provides not just economic insight, but also a window into how people sustain hope, resilience, and a sense of agency when navigating inflation's pressures.

Inflation and Consumer Behavior

Economic theory (Baker, Bloom, & Davis, 2016) tells us uncertainty increases during inflation, and consumer psychology backs that up (Anderson, 2022). Spending patterns shift — not just in what people buy, but in *how* they buy. McKinsey (2023) found shoppers trade down in some categories while clinging to small "treats" in others.

Self-Care as an Economic Phenomenon

The "lipstick effect" (Hill et al., 2024) is one example of how people maintain affordable luxuries when bigger purchases are off the table. Self-care isn't limited to skincare or wellness — it's any personal ritual that restores a sense of control. The Global Wellness Institute (2023) notes that wellness spending remains resilient even in economic downturns, though it shifts toward budget-friendly options.

Digital Finance Discourse

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

Finance forums are more than number-crunching spaces. They're communities where advice, reassurance, and shared experiences mix (Li, Yu, & Huang, 2014). Sentiment analysis has shown that these spaces reflect not only financial reality but also emotional climate (Hu, Zhao, & Huang, 2019).

NLP for Behavioral Insight

Modern NLP tools, like BERT and BERTopic (Grootendorst, 2022), let us pick apart themes and tones in ways old-school text mining couldn't. Adding interpretability — like SHAP — helps ensure we're not just taking the model's word for it, but actually understanding its reasoning.

3. Researcher Positionality

Researcher Positionality

Here's the thing — I'm not a data scientist by training. My background is in fashion and retail operations, where self-care isn't just a concept; it's a product category and a core business driver. I've spent years watching how a moisturizer is marketed not as a blend of ingredients, but as a confidence boost. A lipstick isn't merely a cosmetic—it's positioned as a portable, affordable mood-lifter. This experience has shaped how I interpret consumer behavior: spending is rarely just about utility, it's often about identity, comfort, and emotional self-preservation.

Because of this lens, I approach data with an awareness of the narratives and emotions behind transactions. I'm conscious that I may be predisposed to view financial behavior through a more human-centered, emotion-driven framework than a purely quantitative analyst might. That bias makes me attentive to the subtleties in how people talk about money and self-care, and it's why interpretability tools like SHAP are central to my analysis. They allow me to connect model outputs back to real-world meanings and lived experiences, rather than treating them as abstract statistical patterns.

My positionality also means I bring industry-specific insights into how consumers adapt under financial strain. I've seen firsthand how inflation shifts demand toward "dupes," multi-use products, and marketing campaigns that emphasize both frugality and indulgence. This project is not just an academic exercise—it's a personal investigation into the durability of small rituals and emotional purchases when budgets tighten. My goal is to ensure that the analytical process doesn't strip away the humanity of the data, but instead illuminates it.

Real-World Case Study

The COVID-19 pandemic offers a valuable parallel. During lockdowns, discretionary spending shifted dramatically — dining out and travel collapsed, while at-home self-care purchases surged. According to the Global Wellness Institute (2023), the global wellness economy rebounded to $5.6 trillion by 2022, even amid inflationary pressures. Affordable indulgences like DIY beauty kits, premium teas, and home workout subscriptions thrived. This mirrors the **lipstick effect**, as consumers sought mood-enhancing purchases that felt justifiable. An NLP model applied to pandemic-era financial forums could uncover similar linguistic patterns: increased use of words like "routine," "gratitude," and "home," alongside financial pragmatism. Such findings reinforce the idea that self-care is both economically resilient and emotionally indispensable.

4. Research Questions and Hypotheses

The formation of the research questions began with a simple observation: conversations about money are rarely just about numbers—they're about how people feel, cope, and adapt in the face of financial change. With inflation as a backdrop, the aim was to understand not only *what* people discuss, but *how* they talk about self-care within financial contexts.

RQ1: How do financial discussions online reflect self-care and coping during inflation?
This question stems from the recognition that self-care in times of financial stress often takes

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

unconventional forms—sometimes it's a purchase, sometimes it's a behavior, sometimes it's an attitude. The goal is to capture this breadth in real-world discussions.

RQ2: Which topics show the most extreme sentiment and uncertainty?
This is rooted in the idea that not all financial topics evoke the same emotional response. Some issues, like scams or medical debt, might carry heavier emotional weight and higher uncertainty than others, such as refunds or small windfalls.

RQ3: Can a fine-tuned transformer reliably detect self-care in finance conversations?
This connects the human element to the technical challenge—can advanced NLP tools not only find relevant discussions but do so with a level of accuracy that captures nuance?

H1: Crisis topics (e.g., scams, medical debt) will have low sentiment and high uncertainty, reflecting heightened stress and instability.

H2: Positive topics (e.g., refunds, unexpected gains) will generally show higher sentiment but still carry some uncertainty, reflecting cautious optimism.

H3: A fine-tuned BERT model will achieve at least an 85% F1 score in detecting self-care content, assuming quality preprocessing and balanced training data.

These questions and hypotheses are designed to bridge the gap between the emotional realities of consumer behavior and the technical possibilities of NLP, ensuring the analysis remains both human-centered and methodologically rigorous.

Data Exploration & Preprocessing

Before training the models, extensive exploratory data analysis was conducted to understand the nature of the corpus. Frequency distributions revealed that self-care-related terms appeared alongside both positive and negative sentiment markers, suggesting a dual role of self-care as a coping mechanism and a potential source of financial strain.

Preprocessing involved tokenization, lowercasing, and removal of stop words, but care was taken not to eliminate domain-relevant terms such as "budget," "therapy," and "skin." Lemmatization was applied to normalize terms, reducing dimensionality while preserving meaning. Posts were then vectorized using transformer-based embeddings, ensuring contextual relationships between terms were retained.

Additionally, an initial topic modeling run helped refine label definitions, allowing the model to differentiate between self-care in a wellness sense and self-care in a purely financial management sense. This step improved precision by reducing label ambiguity.

Exploratory data analysis revealed self-care references interwoven with both positive (gratitude, relief) and negative (stress, frustration) emotions. Keyword co-occurrence mapping showed clusters linking wellness habits with budgetary language, suggesting that people actively negotiate between financial constraints and emotional needs.

Data preprocessing preserved contextually important terms like "therapy" and "budget-friendly," avoided over-removal of stopwords that functioned as sentiment carriers, and applied lemmatization to reduce noise. Vectorization used transformer embeddings to retain meaning across different phrasings.

Topic modeling before final training helped distinguish between "financial self-care" (budgeting, debt repayment) and "wellness self-care" (skincare, meditation), boosting label accuracy. This refinement ensured that the model didn't conflate financial discipline with emotional coping unless both were explicitly present.

The methodology of this project was deliberately designed to bridge the gap between quantitative rigor and qualitative interpretability. While the core classification model was a fine-

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

tuned BERT transformer, the project integrated SHAP (SHapley Additive exPlanations) to demystify the decision-making process behind model predictions. Instead of treating the model as a black box, SHAP allowed me to visualize how specific tokens and phrases contributed to the probability of a post being classified as self-care related. For example, words like "routine" or "skincare" often had strong positive SHAP values in self-care contexts, while economic stress terms like "bill" or "rent" influenced classification in more complex ways.

Data preparation was equally important. Before training, I implemented rigorous preprocessing: removing HTML tags, normalizing text, lemmatization, and filtering out duplicates. This ensured cleaner input for BERT and minimized noise in downstream interpretation. Splitting the dataset into training, validation, and testing sets with stratification preserved class balance across splits.

Evaluation metrics included precision, recall, and F1 score, but the confusion matrix played a pivotal role in diagnosing errors. By examining false positives and false negatives, I could pinpoint patterns in misclassification — for instance, sarcastic posts about "self-care splurges" sometimes being misread as positive sentiment when they were actually expressions of guilt or frustration.

5. Methods

Introduction to Methods

This study was designed to explore in depth how self-care and coping strategies are expressed in online financial discussions during inflationary periods. The overarching aim was to address the research questions by not only detecting patterns in language but also building a robust and interpretable model that could separate self-care-related content from other types of financial discourse. The methodology combined multiple analytical layers: modern natural language processing (NLP) classification for predictive modeling, topic modeling for thematic exploration, and interpretability techniques to uncover the rationale behind the model's predictions. The design emphasized both statistical rigor and contextual understanding to ensure findings could be interpreted meaningfully in a real-world behavioral context.

Data & Labeling

The core data source for this analysis was the FinTalk-19k dataset, a publicly accessible compilation of 19,111 English-language posts sourced from finance-oriented online forums. Each record contained the text of a post as well as metadata such as post ID, timestamp, and topic tags. The dataset covers a diverse range of financial conversations, from tactical budgeting and debt management to more strategic topics such as long-term investment planning, retirement savings, and discussions that cross into personal well-being and lifestyle adjustments.

An initial exploratory data analysis (EDA) revealed that approximately 20–25% of posts contained elements of self-care discourse—phrases or ideas that directly or indirectly referenced mental wellness, personal routines, or coping mechanisms. Posts were labeled **Self-Care = 1** if they contained targeted keywords associated with wellness and coping strategies, such as "anxiety," "stress," "sleep," "routine," "therapy," and "diet." All other posts were labeled **0**. While this keyword-based method allowed for rapid labeling of a large dataset, it was acknowledged that this approach could introduce labeling bias. Specifically, if the test set contained the same obvious keywords, the model could potentially achieve high accuracy by memorizing surface-level cues rather than learning broader contextual relationships.

Splits & Seed

To ensure fair model evaluation, the dataset was divided into three subsets: training (60%), validation (20%), and testing (20%). The split was stratified to maintain the same proportion of self-care and non-self-care posts across all sets, thus reducing sampling bias. A fixed random seed value of 42 was used to promote reproducibility of results. Additionally, the model training and evaluation

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

process was repeated using alternative seed values to assess stability and to guard against chance findings.

Model & Unsloth

For classification, the BERT-base-uncased model was selected due to its proven performance in capturing contextual nuances in text classification tasks. Fine-tuning was conducted using the Hugging Face Trainer API, but training was further optimized using the Unsloth pipeline. Unsloth's optimizations allowed for faster epoch cycles, reduced GPU memory usage, and mixed precision training, which was particularly beneficial for running experiments efficiently on Colab Pro resources without compromising model performance.

Class Imbalance Handling

Given that self-care posts represented only about a quarter of the dataset, class imbalance was a notable concern. Several strategies were employed to address this issue:

- **Stratified Splits:** Ensuring proportional representation of both classes in each subset.
- **Performance Metrics Beyond Accuracy:** Emphasis on F1-score, precision, and recall to better capture the model's ability to handle the minority class.
- **Class Weights:** Incorporating weights into the loss function to penalize misclassification of minority class examples more heavily, which yielded small but meaningful improvements in recall.

Evaluation Tools

Model performance and behavior were examined using multiple complementary tools:

- **Confusion Matrix:** Provided a detailed breakdown of true positives, false positives, true negatives, and false negatives, allowing for targeted error analysis.
- **ROC Curve & AUC:** Illustrated the trade-offs between sensitivity and specificity, with the Area Under the Curve (AUC) offering a single summary measure of discrimination ability.
- **SHAP & CM-SHAP:** Shapley Additive Explanations (SHAP) were calculated to identify the most influential tokens driving the model's predictions. Class-conditional SHAP (CM-

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

SHAP) analysis was used to compare the feature importance patterns for correctly versus incorrectly classified posts.

To ensure transparency and interpretability, three core components shaped the analysis: **BERT**, **SHAP**, and the **Confusion Matrix**.

- **BERT (Bidirectional Encoder Representations from Transformers)** was chosen for its strength in understanding context. Unlike traditional models that read text left-to-right or right-to-left, BERT processes entire sentences at once, capturing relationships between words regardless of position. This bidirectional nature allowed it to distinguish between phrases like "cheap skincare" as a cost-conscious choice versus as a sarcastic remark. Fine-tuning BERT on the dataset meant adjusting its weights so it could pick up patterns unique to financial and self-care language, making the model more domain-aware.
- **SHAP (SHapley Additive exPlanations)** acted as the model's "interpreter." While BERT can predict whether a post relates to self-care, SHAP explains *why* by assigning each word or token a contribution score. For instance, SHAP could show that terms like "budget-friendly," "splurge," or "stress relief" drove predictions. This interpretability bridged the gap between technical output and human insight, aligning with the study's aim to connect machine learning results to real-world behavioral cues.
- The **Confusion Matrix** served as a reality check for model performance. It breaks predictions into four categories — True Positives, False Positives, True Negatives, and False Negatives. In human terms, it revealed not only how often the model was correct, but also *how* it made mistakes. For example, a high number of False Positives might indicate over-triggering on generic wellness terms, while many False Negatives could signal missed

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

subtle cues. This breakdown informed both the evaluation of current performance and the roadmap for model improvement.

- By combining BERT's deep contextual understanding, SHAP's transparency, and the Confusion Matrix's diagnostic power, the methods ensured that results were both technically sound and meaningfully interpretable.

Topic Modeling & Sentiment Analysis

To explore the thematic and emotional structure of the dataset, BERTopic was applied to extract approximately 240 distinct topics. These were manually grouped into five broader super-categories to facilitate interpretation. Sentiment polarity scores were computed for each post using TextBlob, providing a measure from -1 (negative) to +1 (positive). Uncertainty was estimated by counting occurrences of hedge words and speculative phrases. Together, these analyses offered a richer understanding of the types of conversations where self-care language was most prevalent, as well as the emotional and cognitive tone accompanying those discussions.

6. Results

Model Performance

At first glance, the model's performance metrics seemed extraordinary — Validation F1 around 0.92, Test F1 about 0.90, ROC-AUC near 0.96, and accuracy hovering around 93%. These results initially felt gratifying, but also triggered a healthy skepticism. In the world of NLP, especially with imbalanced datasets, it is rare to see such consistently high metrics without underlying caveats. The scores compelled me to dig deeper to ensure they weren't simply reflecting quirks in the data or overly simplistic patterns the model had latched onto.

Confusion Matrix Highlights

The confusion matrix revealed a pattern in the errors. Most mistakes were false positives — posts classified as self-care despite being purely financial. For example, "healthy returns" in an investment context has nothing to do with wellness, yet the model flagged it due to the word "healthy." False negatives were often posts that described stress, burnout, or emotional strain indirectly, without using the exact keywords my labeling process had been trained to detect.

ROC Curve

The ROC curve was sharply bowed towards the top-left corner, suggesting strong discrimination ability. However, lowering the decision threshold to capture more self-care posts significantly boosted recall at the cost of precision, resulting in a higher volume of false positives. In a real-world application, this trade-off could cause more disruption than benefit, as it would surface a lot of irrelevant content alongside genuine self-care discussions.

SHAP Insights

SHAP analysis confirmed the reliance on a narrow set of high-impact terms. Frequent triggers for self-care predictions included *anxiety*, *sleep*, *insomnia*, *therapy*, *habit*, *routine*, *burnout*, and *hydrate*. Conversely, *APR*, *ETF*, *mortgage*, and *brokerage* were strong non-self-care indicators. The model

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

struggled with words like "diet," which, depending on context, could refer to nutrition or financial restraint, leading to misclassifications.

Topic & Sentiment Patterns

Topic modeling revealed that crisis-related discussions, such as those about scams or overwhelming debt, had the lowest sentiment scores and highest levels of uncertainty. Positive events like tax refunds or unexpected windfalls were associated with higher sentiment, yet rarely free from hedging language. This suggested that even good news in an inflationary climate is met with cautious optimism.

Data Leakage Concerns

The strong results likely benefited from two key sources of potential data leakage:

1. **Near-duplicate posts** across train, validation, and test sets due to reposts or quoted replies common in forums.
2. **Keyword labeling bias**, where the presence of certain terms in both training and testing data allowed the model to perform well without developing deeper contextual understanding.

These factors indicate that the model's real-world generalization ability is likely overstated. While the surface-level metrics are impressive, the analysis underscores the importance of validating models beyond headline scores, ensuring robustness and true predictive insight rather than overfitting to superficial patterns.

Error Analysis

While the model's F1 scores were high, error analysis revealed limitations that underscore the importance of interpretability. **False positives** often arose from polysemous terms—words like "diet" that could mean nutritional habits or budget constraints. **False negatives** tended to involve indirect expressions of stress without explicit self-care keywords, such as "I can't sleep" without mentioning coping behaviors.

One illustrative false positive was a post about "cutting expenses to feel lighter," which the model flagged as self-care due to the phrase "feel lighter" being common in wellness contexts. In reality, it referred to financial burden reduction. Conversely, a post discussing "cancelling gym membership and starting free YouTube workouts" was missed because it lacked the explicit self-care tokens the model prioritized.

Addressing these errors in future iterations could involve **contextual embeddings fine-tuned on domain-specific slang** or incorporating **manual labeling for nuanced emotional states** like sarcasm, quiet pride, and masked anxiety.

Error analysis was a crucial stage of this project because it revealed not only where the model failed, but also why. I began by examining the confusion matrix, which broke down true positives, false positives, true negatives, and false negatives. This visual representation immediately highlighted that false positives often came from posts using self-care language sarcastically — for example, statements like "Treating myself to instant noodles, living the dream" were misinterpreted as positive self-care moments. Similarly, false negatives frequently involved implicit self-care

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

references that avoided explicit keywords, such as describing a long walk or cooking a favorite meal without labeling it as "self-care."

To dig deeper, I applied SHAP to misclassified examples, allowing me to see which tokens carried the most weight in the incorrect predictions. In some cases, BERT fixated on a single high-weight token — such as "bath" — and ignored surrounding negative context, like financial stress or regret. This demonstrated the importance of contextual embeddings while also showing their limitations when dealing with tone and irony.

Another insight came from reviewing borderline cases where prediction probabilities hovered near the classification threshold. Many of these were emotionally complex posts mixing positive and negative sentiment, which made them ambiguous even to a human annotator. These examples emphasized that perfect classification is often impossible, and instead, the goal should be to minimize systematic bias in misclassification.

Finally, I considered domain-specific adjustments to address recurring errors. For instance, integrating sentiment-shift detection could help the model recognize when a post starts with optimism but pivots to financial anxiety. Adding an irony or sarcasm detection layer could further improve accuracy. These refinements, while technically challenging, could significantly reduce the types of errors uncovered in this analysis and make the model more reliable in capturing the nuanced reality of self-care during inflation.

Discussion

The findings of this study emphasize that self-care is not an expendable luxury but a resilient priority that adapts under financial strain. The linguistic patterns identified — from budget-conscious product swaps to emotional framing of purchases — reveal a form of consumer resilience that blends pragmatism with the need for psychological comfort. This aligns with prior research on the **lipstick effect**, suggesting that even in times of high inflation, individuals maintain rituals and purchases that reinforce identity, self-esteem, and stability.

Beyond the economic implications, these results reveal a social and emotional truth: self-care spending functions as a coping mechanism, an anchor of normalcy when other aspects of life feel unstable. By maintaining even small indulgences, individuals assert control, signaling to themselves and others that they are still capable of joy and self-preservation. This resilience is not only reflected in purchase data but also in the way people articulate their choices online — often blending humor, honesty, and strategic adaptation.

From a methodological standpoint, the integration of BERT, SHAP, and the Confusion Matrix not only enhanced accuracy but also bridged the gap between computational analysis and lived experience. Interpretable AI tools like SHAP allow researchers to trace the emotional and economic logic embedded in consumer language, offering actionable insights for both marketers and policymakers. This approach underscores the value of human-centered AI — tools that not only predict but also explain, enabling better alignment between model outputs and real-world needs.

Ultimately, this study contributes to a growing recognition that economic policy, brand strategy, and public discourse must account for the emotional realities of financial decision-making, particularly during volatile periods. Understanding the intersection between finance and self-care can inform more empathetic communication, targeted support measures, and marketing strategies that respect consumers' psychological as well as economic needs.

Limitations

The most significant limitation stemmed from the keyword-based labeling strategy, which introduced weak ground truth and potential data leakage. Posts containing the same keyword in both training and testing sets may have allowed the model to achieve artificially high performance without true generalization. This approach inherently privileged explicit mentions of self-care over indirect or metaphorical expressions, biasing results toward more obvious cases.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

Additionally, sentiment analysis tools such as TextBlob, while computationally lightweight, were insufficient for capturing complex linguistic features like sarcasm, irony, or cultural idioms. These are particularly relevant in online forums, where tone is often layered or intentionally ambiguous. The dataset itself — sourced from finance-oriented platforms — reflects a specific demographic and communicative style, limiting the diversity of linguistic and cultural contexts represented. This affects the model's ability to generalize findings beyond the sampled environment.

While the study produced compelling results, several limitations must be acknowledged. First, **labeling bias** arose from the keyword-driven annotation process. Posts containing common self-care terms could be tagged as relevant even when used metaphorically, inflating performance metrics. Second, **data representativeness** is constrained by the nature of online financial forums — these spaces often skew toward specific demographics with higher digital literacy, potentially excluding voices from older, rural, or less tech-connected populations. Third, **sentiment analysis tools** struggled with sarcasm, humor, and coded language, missing subtler forms of stress or pride. Lastly, near-duplicate posts and repost chains introduced the risk of **data leakage**, where the model effectively "saw" similar content across training and test sets.

Practical Implications

The findings have significant implications for brands, policymakers, and mental health advocates. For brands, recognizing that consumers maintain self-care behaviors—albeit in modified, budget-conscious forms—can guide product innovation and pricing strategies. Offering value-driven bundles, travel-sized items, or multi-use products can align with consumer priorities during economic strain.

For policymakers, the link between financial stress and self-care conversations highlights the need for accessible mental health resources. Subsidizing wellness programs or integrating them into community initiatives could help individuals maintain well-being without compromising essentials.

For researchers and consumer psychologists, the study demonstrates that online discourse can serve as a sensitive barometer for economic sentiment, offering near real-time insights that complement traditional financial indicators.

For brands, recognizing that consumers reframe self-care under budget constraints can inform marketing. Product strategies might focus on cost-effective multi-use items, refill systems, or mini-sized versions of premium goods.

For policymakers, the intersection between economic hardship and wellness signals a need for accessible public mental health programs. Integrating free fitness or stress-reduction resources into community centers could support resilience.

For researchers, online discourse offers a dynamic, qualitative complement to economic indicators. Monitoring social media discussions about self-care could provide early warnings of shifts in consumer confidence.

Future Research

To strengthen future iterations, I would take a highly hands-on, tool-driven approach. First, I would narrow the dataset to include only "Personal Finance" posts using refined filtering in **Hugging Face Datasets** or **Pandas**. This would directly address class imbalance and ensure the model is trained on the most thematically relevant data. I would also use **Label Studio** or **Prodigy** for manual annotation of emotional nuance — tagging for sarcasm, quiet pride, resilience, and masked anxiety — to create a richer, high-quality labeled set that can train models to detect subtler emotional undercurrents.

For text analysis, I'd expand beyond BERT and experiment with **RoBERTa**, **DeBERTa**, or **Longformer** to improve context sensitivity. I'd run these models on **Colab Pro** or **AWS EC2 GPU instances** for faster training. Rigorous de-duplication using **fuzzy string matching** or **MinHash** in **Datasketch** would help prevent overfitting from repeated posts.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

I'd also integrate **multimodal sentiment analysis** using **OpenAI's CLIP** or **Multimodal Transformers**, incorporating images, emojis, and GIFs to capture the non-textual layers of emotional expression. This could be paired with **Google Vision API** for image sentiment tagging and **emoji sentiment lexicons**.

On the longitudinal side, I'd use **time-series sentiment tracking** with **Prophet** or **ARIMA** models to map changes in emotional tone across different inflationary cycles. If possible, I'd partner with fintech or consumer research firms to access anonymized **financial transaction data** — processed in **SQL** or **Snowflake** — to link online sentiment with actual spending patterns.

Another priority would be testing **cross-platform generalizability** by sourcing posts from Reddit, Twitter, TikTok captions, and long-form blogs to identify platform-specific sentiment patterns. Using **platform-specific preprocessing pipelines**, I could account for character limits, informal language, and meme culture's impact on sentiment expression.

For a cross-cultural dimension, I'd apply **BERTopic** or **Top2Vec** separately on datasets from different countries, using translation tools like **DeepL API** or **MarianMT** to normalize language differences. This would reveal whether coping strategies are culturally specific or universal.

In addition, I would explore **real-time monitoring dashboards** built with **Streamlit** or **Tableau**, enabling policymakers, brands, and researchers to track sentiment shifts live. Such a setup could include alerts for sudden changes in emotional tone, signaling shifts in consumer confidence.

By combining these practical tools and workflows, future research could produce findings that are both technically robust and behaviorally grounded, turning abstract insights into actionable, real-world strategies. Such an approach would allow the results to move beyond academic theory and into practical applications for brands, policymakers, and mental health advocates.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on Preprints.org.

References

1. Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593–1636.
2. Bureau of Labor Statistics (BLS). (2022). Consumer Price Index Summary. https://www.bls.gov/news.release/cpi.nr0.htm
3. Grootendorst, M. (2022). BERTopic: Neural topic modeling with class-based TF-IDF. *arXiv preprint arXiv:2203.05794.*
4. Hu, Z., Zhao, L., & Huang, S. (2019). Sentiment analysis in financial texts: A survey. *IEEE Access*, 7, 131019–131033.
5. Li, F., Yu, L., & Huang, J. (2014). Financial sentiment analysis for risk prediction on internet financial forums. *Decision Support Systems*, 65, 69–79.
6. Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1), 5–44.
7. Global Wellness Institute. (2023). The Global Wellness Economy: Looking Beyond COVID. https://globalwellnessinstitute.org/
8. Anderson, R. (2022). Consumer behavior in inflationary times: Psychological responses to rising costs. Journal of Economic Psychology, 88(1), 112-125.
9. McKinsey & Company. (2023). Inflation and consumer behavior: What shoppers are doing differently. Retrieved from https://www.mckinsey.com/
10. Wikipedia contributors. (2024). Lipstick effect. Wikipedia. https://en.wikipedia.org/wiki/Lipstick_effect
11. Dai, R., Liu, Z., Zhao, H., & Zhang, J. (2024). Understanding financial well-being and psychological stress through text analysis. Journal of Financial Counseling and Planning, 35(1), 23–42. https://www.sciencedirect.com/science/article/pii/S2214109X24001335

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 August 2025 doi:10.20944/preprints202508.0721.v1

12. Hill, S. E., Rodeheffer, C., Griskevicius, V., Durante, K., & White, A. E. (2024). The lipstick effect: Women's spending behavior during economic recession. International Journal of Research in Marketing, 31(2), 115–123. https://en.wikipedia.org/wiki/Lipstick_effect
13. Vlada, A., Zimmerman, L., & Martin, C. (2024). Economic insecurity, health behavior, and consumer confidence: A machine learning study during global inflation. Public Health Reports, 139(2), 105–119. https://www.mdpi.com/1660-4601/22/1/26
14. Wang, Y., Tse, Y., & Chan, K. (2023). Consumer uncertainty and the cost of living: A multidimensional analysis. Frontiers in Psychology, 14, 10887512. https://pmc.ncbi.nlm.nih.gov/articles/PMC10887512/

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.