How Did People Tweet against Inflation in Japan?*

Toshitaka Sekine† and Tetsuro Wada‡

July 14, 2025

Abstract

During the chronic deflation era starting in the 1990s, Japanese inflation expectations were said to be firmly anchored at a very low level, say, around zero. These expectations seemed to have become something like the social norm. Households were quite against any price hikes, and as a consequence, firms hesitated to raise their prices — when they raised prices, they apologized for their misbehavior. People not only expected that prices *would not* increase, but also believed that prices *should not* increase. That social norm may have changed in response to inflationary shocks after COVID-19 and the Ukraine war. We applied a natural language processing technique to tweets that commented on price hikes and found an increase in posts after 2021 that accepted price hikes for various goods. Some of these posts indicated even positive feelings and mentioned salary hikes.

JEL Classification Number: C0, E31

Keywords: tweet, natural language processing, sentiment analysis, inflation expectation, monetary policy, Japan

*We are heavily indebted to Atsushi Keyaki, who generously guided us on cutting-edge AI applications. We thank seminar participants at the 7th EcoSta International Conference, the 18th CFE-CMStatistics International Joint Conference, the SWET2024, and the TCER Summer Conference for valuable comments and discussions. JSPS KAKENHI Grant Number 23H00050 and 24K04882 supported this work. Any remaining errors are our own.

†Graduate School of Economics, Hitotsubashi University. Email: toshitaka.sekine@r.hit-u.ac.jp (corresponding author)

‡Finatext Ltd. Email: tetsuro.wada@finatext.com

1 Introduction

At last, Japan has left a decades-long chronic deflation since the 1990s (Figure 1). In 2022, Japan's CPI inflation began to soar, reaching around 4 percent yearly. Using millions of tweets posted on the SNS (Social Networking Service), this paper investigates whether or not Japanese household perceptions of price hikes have changed against this backdrop.

The perceptions against price hikes are supposed to hold a key during the chronic deflation era. During that period, households were quite against price hikes, and firms hesitated to raise their prices to avoid customers' anger. As an anecdote, when a Japanese ice cream company, Akagi, raised the price of its ice cream bar from 60 to 70 yen after holding it for a quarter century, the company aired TV advertisements in which its president, with many employees, deeply vowed to show their apology (Watanabe, 2024). People not only expected that prices *would not* increase but also believed that prices *should not* increase. Watanabe (2022) and Nishizaki et al. (2014) argue that these strong perceptions were so firmly embedded in Japanese society to become the social norm. In economics, Okun (1981) emphasized the role of a norm in that the high wage norm was behind the hyperinflation in the United States from the 1970s to the 1980s.

These anti-inflation perceptions may have changed in recent years in response to inflationary shocks after COVID-19 and the Ukraine war, and that change in perception may have led to persistent inflation, in turn. Then-Governor of the Bank of Japan, Kuroda (2022) stated, "As firms adopt an increasingly active price-setting stance, Japanese households' tolerance of price rises has been increasing. This can be regarded as an important change from the perspective of aiming to achieve sustained inflation." His remark incurred negative reactions in SNS,[1] and he took back his remark at the Diet session on June 8th, two days after his original presentation. Although his remark might not have

[1] In our sample posts, "A survey found that over half said they'd still shop at the same supermarket even after a 10% price hike. If that's being used to justify BOJ chief Kuroda's claim that consumers are fine with price hikes, that's just dumb. All supermarkets raised prices—there's no other choice." (posted on 2022-06-07, 23:22:11)

Figure 1: Consumer Price Index

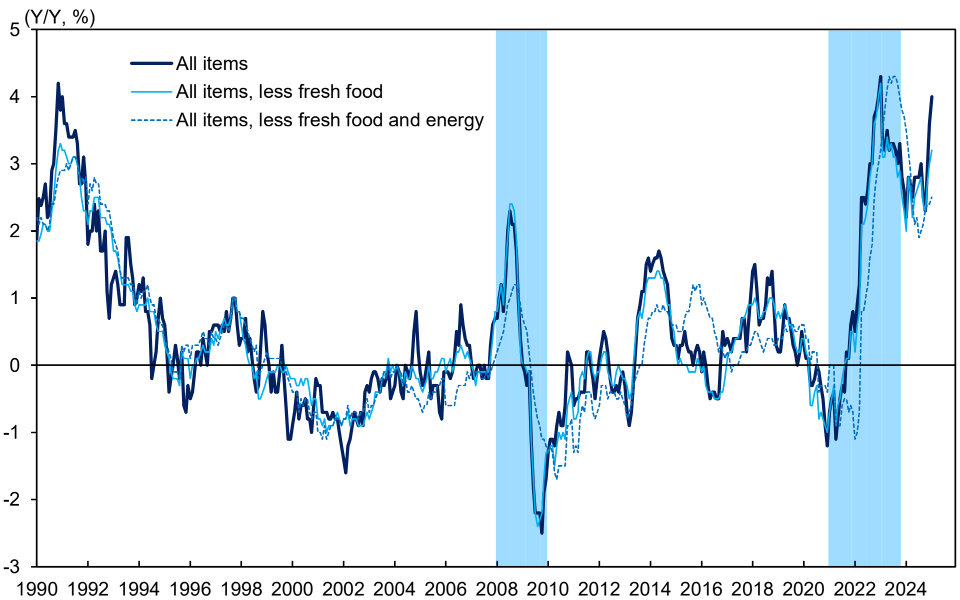

1. Adjusted for the effects of consumption tax rate hikes.

2. Shaded areas correspond to the periods for the tweet analysis below.

Source: Ministry of Internal Affairs and Communications.

been adequate in light of public communication then, what he pointed out might have been correct.

The Bank of Japan's household survey points to a change in anti-inflation perception (Figure 2). In that survey, the Bank of Japan asked 4,000 people (the response ratio is about 50%) how much the present price levels have changed from one year ago. It also asked respondents who answered that the present price levels have increased, which choice (rather favorable/rather unfavorable/difficult to say) is most appropriate to describe their feelings about the price rise. Compared to the 2008-2009 period, when CPI inflation surged owing to commodity price hikes, the perceived inflation was higher in the recent period. Still, the Unfavorable Diffusion Index (DI), the share of respondents choosing "rather unfavorable" minus that of "rather favorable," remains lower. This may suggest an increase in Japanese households' tolerance of price rises.

Figure 2: BoJ Opinion Survey

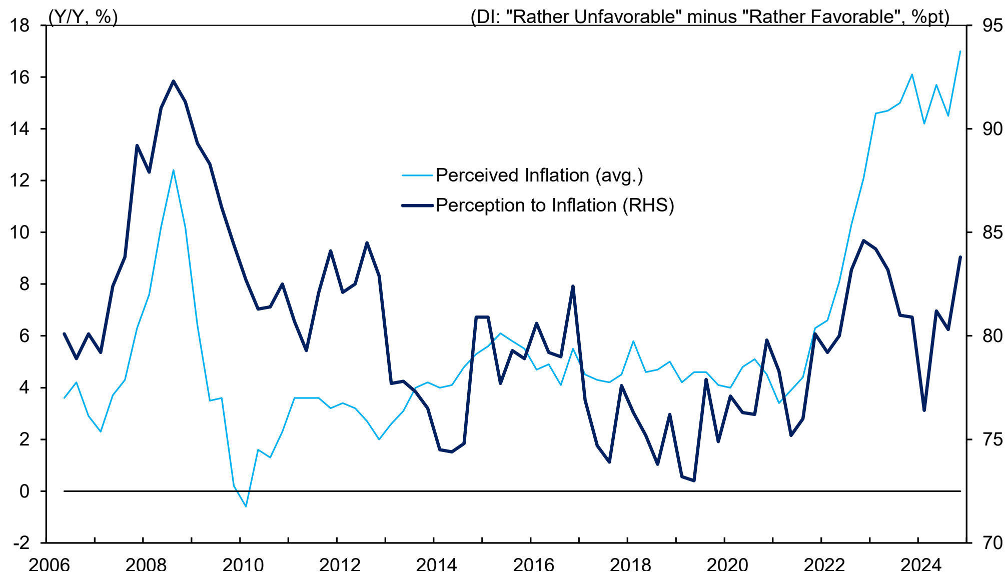

Source: Bank of Japan.

The change in perceptions is also visible in other sources. First, the Nikkei newspaper recently reported price hikes more positively than the 2008-2009 period (Figure 3). The lexicon approach, details of which will be elaborated later (Section 3), is applied to the texts of all Nikkei articles containing words related to price hikes in the titles of the articles (3,738 articles in total). The average sentiment scores of these articles became higher, especially after 2022. This means that the tone of the articles reporting price hikes became more positive. This may reflect and/or affect the general public's perception of price rises. Second, the same tendency is found in the Economy Watchers Survey.[2] The same lexicon approach is applied to the comments of economy watchers on current conditions, which contain words related to price hikes (3,901 comments in total). Like in the case of the Nikkei newspaper, the calculated average sentiment scores were higher in recent periods.

[2]The Economy Watchers Survey is an economic survey conducted by the Cabinet Office of Japan. It collects opinions from people closely observing economic activities in their daily work, such as retail store managers, taxi drivers, and restaurant owners. These respondents, known as "economy watchers," provide insights into current economic conditions based on their firsthand experiences.

Figure 3: Nikkei Newspaper and Economy Watchers Survey

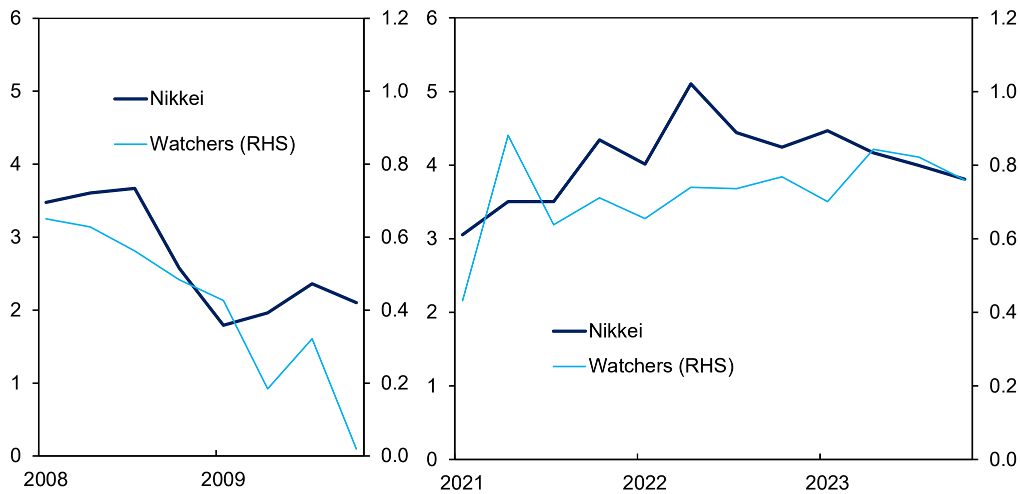

Note: Tone of Nikkei newspaper articles and comments of Economy Watchers against price hikes. The higher figure corresponds to a more positive tone.

Source: Authors' calculation.

To the extent that these economy watchers represent the general public, this may indicate that the Japanese became more tolerant of price hikes.

This paper uses tweets on the SNS to determine whether Japanese households have changed their perceptions of price hikes. It collects more than two million posts mentioning price hikes for that purpose. Applying advances in Natural Language Processing, the paper finds a change in the sentiments or tones of these posts from the 2008-2009 period to the 2021-2023 period. Although more posts indicated anger, those accepting price hikes have increased at the same time. While these posts were associated with only specific goods (such as tobacco) in 2008-2009, there was no specificity in 2021-2023—people seemed to accept price hikes for various goods. Posts showing even pleasant feelings (valence) marginally increased, and some mentioned salary increases.

The structure of the paper is as follows: After touching upon related literature in the rest of Section 1, Section 2 introduces used tweet data. Section 3 explains various approaches to natural language processing. Section 4 summarizes the main results. Section

5 concludes the paper.

1.1 Related Literature

As mentioned above, this paper is closely related to the zero inflation norm discussed by Watanabe (2022, 2024) and Nishizaki et al. (2014). The importance of the norm is also emphasized by Bank of Japan (2024) in its review of unconventional monetary policy. Aoki et al. (2019) models the households' reluctance to accept price hikes affects firms' price-setting behavior through a kinked demand curve.

The paper is also related to the research into household inflation expectations in Japan, examining what their characteristics are and how they are formed. There are many papers in this domain, such as Hori and Kawagoe (2011), Kamada (2013), Kamada et al. (2015), and Diamond et al. (2020), to mention a few.

Methodologically, the paper can be seen as another application of Natural Language Processing (NLP) in Economics. Economists have already used NLP extensively. For instance, Ahrens and McMahon (2021), Shapiro et al. (2022), Nakajima et al. (2021), and Heddaya et al. (2024, 2025) extract economic signals, inflation expectations, or narratives from text data like newspaper articles, central bankers' speeches, and Economy Watchers' comments.

In NLP, sentiment analyses of tweets pose particular challenges, as discussed below (Section 3), and there are many attempts in the literature such as Giachanou and Crestani (2016), Mohammad (2016), Zimbra et al. (2018), and Braig et al. (2023). Bollen et al. (2011) use Twitter sentiment analysis to predict stock prices. Ehrmann and Wabitsch (2022) and Wabitsch (2024) analyzed tweets mentioning ECB monetary policy.

The paper exploits the recent advance of Deep Learning or Large Language Models (LLM) for NLP analyses. Dell (2024), Korinek (2023, 2024) and Kwon et al. (2024) discuss potential usage of LLM in Economics. At the time of writing this paper, applications of

LLM are still limited in the economic literature. However, given the very rapid speed of technical advances—Korinek (2023) cites that the amount of computational power employed in training cutting-edge LLMs has doubled, on average, every six months—economists will take advantage of it in their applications and more papers are sure to come soon. This paper can be regarded as one of the early attempts in that direction. Given the speed of technological developments, the method used in this paper may become obsolete quickly. However, it is worth taking a still shot to record and share the current methodology with economists who are interested in this field.

2 Data

We extensively use tweets posted on X (Twitter in its old name). We collect those written in Japanese that contain words related to price hikes (excluding reposts or retweets).[3] These tweets were posted from January 1, 2021, to October 16, 2023; the sample end corresponds to when we started this project. For comparison, we also collect tweets posted from January 1, 2008, to December 31, 2009. During this period, CPI inflation increased, reflecting commodity price hikes (Figure 1). However, the surge was short-lived as inflation turned negative after the rough waves of the Great Financial Crisis reached the Japanese economy. It would be interesting to compare the tweet sentiment between this short-lived inflation surge and the recent inflation episode.

Table 1 shows the number of tweets for our analysis, together with the number of articles from Nikkei Newspaper and the comments of Economy Watchers discussed above. One thing that is immediately clear is a dramatic increase in tweets from 2008-2009 to 2021-2023. This is the case even though the sampling rate of the former sample period is 100% viz-a-viz 18.7% in the latter sample period. This reflects the higher inflation rate in the latter sample period, which may also increase the number of Nikkei articles and

[3] The following keyword search extracts the relevant posts.
「値上」OR ((「価格」OR 「値段」) AND (「引き上」OR 「上昇」OR 「高ま」OR 「高く」)).

	2008	2009	2021	2022	2023	Total
Tweets	20,253	40,361	398,483	987,116	704,356	2,150,569
o/w BERT selected	1,104	2,010	20,048	61,377	48,648	133,187
Nikkei Newspaper	1,194	516	501	871	656	3,738
Economy Watchers	1,089	36	146	1,256	1,374	3,901
Headline CPI Inflation (%)	1.4	-1.4	-0.2	2.5	3.2	

Table 1: Number of Posts/Articles/Comments

Economy Watchers' comments. However, this is much more because the SNS has become more prevalent in Japanese society. This leaves a caveat for our analysis: people posted in the former sample period might be more restricted than those in the latter.

Although the inflation rate was negative in 2009 and 2021, we include these years in the sample to see to which extent the sentiment in 2008 and 2022-2023 prevailed before and after. That said, we also have in mind that the comparison should be made more narrowly between 2008 and 2022-2023.

3 Approaches

Rapidly advancing NLP technology has evolved through the following three approaches.

First, the lexicon-based approach relies on a predefined dictionary or lexicon of words, where each word is associated with specific semantic, syntactic, or sentiment-related attributes. This approach is commonly used in sentiment analysis. For instance, words in a text are matched against a sentiment lexicon to determine their polarity (positive, negative, or neutral). The overall sentiment of a sentence or document is then computed by aggregating the sentiment scores of individual words. While lexicon-based methods are simple and interpretable, they struggle with handling context, negation, sarcasm, and domain-specific language, making them less effective than modern deep learning-based NLP techniques. Ahrens and McMahon (2021) and Shapiro et al. (2022) use this approach to analyze the sentiments prevailing in newspaper articles and central bank speeches.

Second, the classical machine learning approach involves training models to recognize patterns and make predictions based on textual data using traditional machine learning methods such as Naive Bayes, Support Vector Machines, Decision Trees, and Random Forests. These models often rely on handcrafted features such as term frequency, n-grams, and TF-IDF, making them effective for tasks like text classification, sentiment analysis, and named entity recognition. Nakajima et al. (2021) uses a Naive Bayes model to derive firms' inflation expectations from comments of the Economy Watchers Survey.

Third, the deep learning approach uses deep learning architectures, including Generative AI models. Large Language Models (LLMs) are a subset of Generative AI models, which are built using transformer architectures, such as GPT (Generative Pre-trained Transformer) and BERT (Bidirectional Encoder Representations from Transformers), trained on massive datasets to understand and generate human-like text. They can learn contextual representations of words, enabling superior performance on complex tasks, including sentiment analysis, with minimal fine-tuning. However, these models require massive computational power and pose challenges like bias, factual inaccuracies, and hallucinations. Kwon et al. (2024) apply LLM (Llama) to identify perceived drivers of stock market prices using news reports.

This paper uses the first (lexicon-based) and the third (deep learning) approaches.

The procedure for the lexicon-based approach follows the standard NLP for Japanese texts, which consists of the following two steps:

Lexicon-based Approach

1. Decompose sentences in the examined text into words. We use Mecab (a standard Japanese morphological analyzer) with the NEolgd dictionary (a customized system dictionary for MeCab containing many new words extracted from many language resources on the Web).

2. Calculate the sentiment/tone score of sentences by aggregating the polarity of words in the sentences using a list of words annotated by the polarity. We utilize the polarity list from Ito et al. (2018), which is suitable for market sentiment analysis.

We apply this procedure to Nikkei newspaper articles and Economy Watchers' comments containing words related to price hikes, the results of which are demonstrated in Figure 3.

While the lexicon-based approach performs reasonably well in analyzing standardized texts like newspaper articles, detecting the sentiment of tweets is a non-trivial task, as emphasized by Giachanou and Crestani (2016). They list the challenges of the tweet sentiment analysis stemming from the following characteristics, which are particular to tweets.

- Text length: The short texts up to 140 characters (except for some users after about 2023) make it difficult to detect the context.
- Topic relevance: Many posts are unrelated to the topics of interest.
- Incorrect language: Emphatic upper-casing, emphatic lengthening, abbreviation, slang, neologisms.
- Data sparsity: There is a lot of noise due to incorrect language and misspellings.
- Negation: Positive becomes negative or vice versa.
- Stop words: Often filtered out words (e.g., "like") have meanings.

A casual look at our tweets reveals that these claims apply to our case — Online Appendix A demonstrates some examples of our tweets in Japanese to give a sense of them. The topic relevance is our particular concern. So many posts do not contain sentiments against price hikes.

To select relevant posts for our analysis, we utilize a deep learning model, BERT, a language model widely used in natural language processing. We prepared 2,500 training data for BERT to select relevant posts (plus 500 posts for tests), which state price hikes and express sentiment against them. For that sake, we hired seven university students. We conducted three trial runs and two online meetings to develop an annotation guideline (Online Appendix C) that would establish a common understanding of selecting posts mentioning recent price changes and containing positive or negative sentiments (including emoticons and slang). Using these inputs, BERT attained a precision of 80% for validation data and 67% for test data. Then, we asked BERT to select posts that satisfied the criteria, ending with 133,187 tweets, as shown in Table 1.

Then, we apply the above lexicon-based approach to the selected tweets. We also use the deep learning approach. This is because a simple positive/negative sentiment analysis may not capture an increase in households' tolerance of price rises. Naturally, very few consumers are happy to see a price rise but accept it as there is no other choice, mumbling "Oh well, it is what it is." The second example of tweets in Online Appendix A can be translated as "The bento (lunch box) shop I always order from came by. They asked to raise prices due to inflation... Can't be helped." We use an LLM to see whether the examined posts contain this can't-be-helped feeling.

We use GPT (gpt-4o-2024-05-13), another deep learning model, to judge the sentiment of the posts by assessing whether or not the corresponding post suggests an attitude of acceptance or tolerance for price changes (the used prompt is in Online Appendix D). Then, we calculate the share of these posts among BERT-selected posts. Although Korinek (2023) shows an example of using GPT for sentiment analyses, the approach has

Figure 4: Russell's Circumplex Model

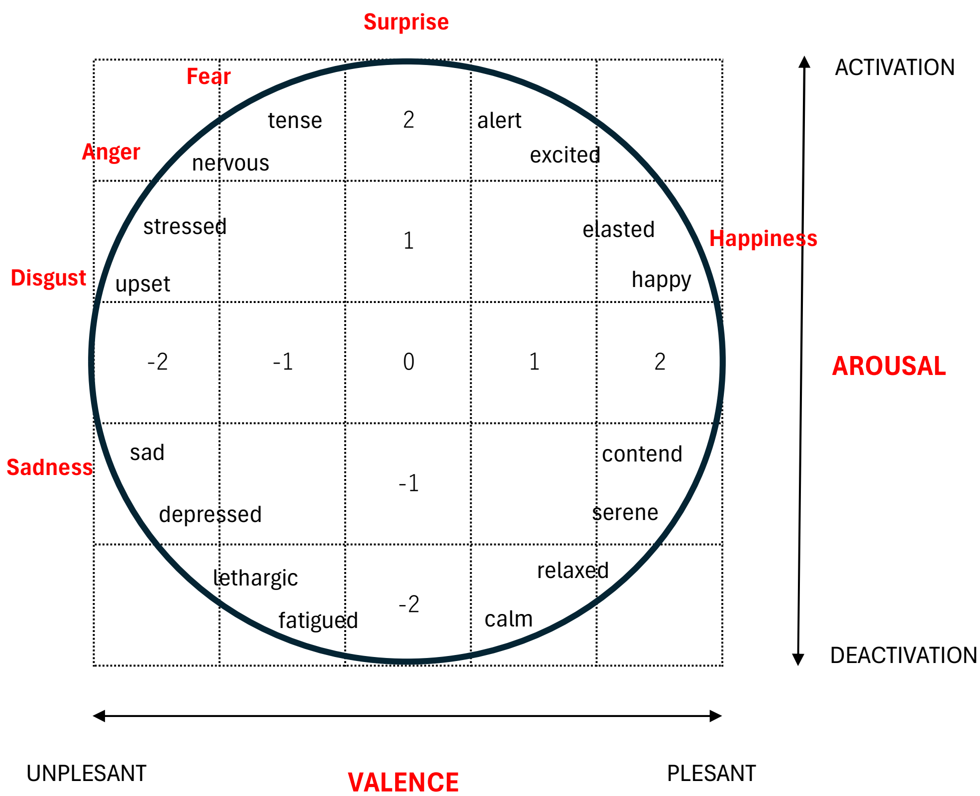

Note: Modified by authors from Russell and Barrett (1999).

not yet seen many applications in economics. GPT substitutes the above annotation process of seven students who prepare trained data for BERT, assuming that GPT has been pre-trained by digesting a gigantic amount of text. Dell (2024) compares two approaches: GPT has the advantages of low startup costs and no training data. However, it has the disadvantages of less fine-grained control, no reproducibility, and a tendency to be a black box. When writing this paper, it is still premature to say that economists will accept this approach in their standard toolbox, but we see great potential, as shown below.

We also use GPT to apply Russell's circumplex model of emotion (Russell and Barrett, 1999). Psychologists developed the model to understand and map out various human emotions on a circle rather than using categories like "happy" or "angry" alone. It shows how emotions relate to each other based on two simple dimensions: one is "valence" (how pleasant or unpleasant an emotion feels), and the other is "arousal" (how physically or mentally activated you feel). The model suggests that various emotions can be placed in a circle space represented by these two dimensions, such that excitement and joy are in the territory of positive valence and positive arousal; anger, fear, and anxiety are negative valence and positive arousal; sadness, boredom, and resignation are negative valence and negative arousal; and relaxation and calm are positive valence and negative arousal (Figure 4). There are some attempts to make AI learn this model and analyze consumers' sentiments for marketing purposes.

More specifically, we instruct GPT to use Russell's circumplex model and score each post on a scale of -2 to +2 for arousal and valence, respectively (see Online Appendix E for the used prompt). This means that GPT places a post in one of 25 cells in Figure 4. We also ask about the reasons for the respective evaluation.

In sum, our deep learning approach can be summarized as the following procedure. Using three different procedures (steps 3.1, 3.2, and 3.3 below) may serve as robustness checks.

Deep Learning Approach

1. Clean up posts (exclude posts containing URLs).

2. Use BERT to select posts related to the topics of interest.

3.1 Use the lexicon-based approach to have sentiment scores of selected posts.

OR

3.2 Use GPT to judge the selected posts' sentiments (accepting price changes or not).

OR

3.3 Use GPT to judge the selected posts' sentiments (applying Russell's circumplex model).

	2008	2009	2021	2022	2023
1. Tweet (Lexicon)	0.64	0.44	0.56	0.72	0.66
2. Tweet (GPT)	5.21	8.75	12.31	12.03	12.29
3. Nikkei Newspaper	3.47	2.11	3.64	4.51	4.13
4. Economy Watchers	0.61	0.30	0.69	0.74	0.78

Note:

1. Rows 1, 3, and 4 represent sentiment scores derived from the lexicon approach.
2. Row 2 is the share of posts accepting the price hike.

Table 2: Sentiment Scores

4 Results

Tweet posts suggest that the Japanese may have become more tolerant of price hikes. Table 2 summarizes the results of the above steps 3.1 and 3.2 of the deep learning approach. It also indicates the sentiment scores of the Nikkei Newspaper and the Economy Watchers Survey, shown in Figure 2 previously. Similar to the case of the Nikkei Newspaper and the Economy Watchers Survey, the lexicon approach reveals that the tone of tweets mentioning price hikes has become more positive, especially from 2022 to 2023 (Row 1). The share of tweets that GPT judges as accepting price hikes has also increased for the 2021-2023 period (Row 2).

Russell's circumplex model also points to higher tolerance for price hikes. Table 3 is the outcome of Russell's circumplex analysis conducted by GPT (the above step 3.3), which shows the shares of the levels of arousal and valence each year. The share of the high arousal (2) and the low valence (-2) has increased from 2008-2009 to 2021-2023. As the combination of both corresponds to anger in Russell's circumplex (Figure 4), more tweets reveal anger against price hikes in the recent period. At the same time, the share of the positive valence (1 and 2) has marginally increased. This means that slightly more tweets take price hikes positively.

To see more background on these changes in the Twitter sentiment, we further analyze

Arousal

Level	-2	-1	0	1	2	Total
2008	0.0	20.2	6.2	59.6	14.0	100.0
2009	0.0	20.5	8.3	59.9	11.4	100.0
2021	0.0	15.0	4.6	61.1	19.2	100.0
2022	0.0	16.3	3.9	60.7	19.0	100.0
2023	0.0	16.5	3.9	60.2	19.4	100.0

Valence

Level	-2	-1	0	1	2	Total
2008	20.8	72.0	3.1	3.3	0.7	100.0
2009	18.4	68.2	5.0	7.4	1.0	100.0
2021	28.5	56.6	3.8	8.7	2.4	100.0
2022	32.5	56.8	3.0	6.5	1.3	100.0
2023	33.2	56.0	2.9	6.5	1.4	100.0

Note: The results of Russell's circumplex analysis.

Table 3: Arousal and Valence (%)

(1) tweets posts that GPT judges as accepting price hikes in 2008-2009, (2) those in 2021-2023, and (3) tweets posts that GPT judges as having positive valence in 2021-2023. The number of these posts is 228, 15,663, and 10,721, respectively. As mentioned above, Twitter was less widely used in the 2008-2009 period. As a result, the number of tweets posts accepting price hikes is very small.

Word clouds of these three posts reveal two distinguished features (Figure 5). First, "tobacco" (the Japanese language does not distinguish between tobacco and cigarettes, but uses kanji, hiragana, and katakana as indicated in original Japanese word clouds in Online Appendix B) stands out in the word cloud of posts accepting price hikes in 2008-2009 (top panel). An example of a post reads "They should just jack up the price of cigarettes to 1,000 yen already." In contrast, there are no specific goods in 2021-2023 (middle panel). Price hikes were accepted almost exclusively for tobacco in 2008-2009, whereas people became more accepting of price hikes for various goods and services in 2021-2023.

Second, the word cloud of posts accepting price hikes in 2021-2023 (middle panel) resembles that of positive valence in 2021-2023 (bottom panel). “cheap”, “good”, and “buy” are also frequently mentioned (presented by larger fonts) in 2008-2009 (top panel). Still, many other words are commonly found in the middle and bottom panels. The Weighted Jaccard Similarity is 0.67 between the middle and bottom panels, which is much higher than that between the top and middle panels and the top and bottom panels (Table 4).[4] This may suggest that some common elements made people more accepting of the price hikes and feel a higher valence in 2021-2023.

	(1)	(2)	(3)
(1)	-		
(2)	0.33	-	
(3)	0.30	0.67	-

Note: (1)-(3) correspond to those in Figure 5.

Table 4: Weighted Jaccard Similarity

To gain insight into the narrative behind the change in perception, we conducted a keyword search on the above sets of posted tweets (Table 5). For instance, Japanese people may accept price hikes if uncontrollable exogenous factors, such as the COVID-19 outbreak or the war in Ukraine, triggered these price increases. They may not be upset if their salaries also increase. If that were the case, the above-selected tweets should contain the corresponding words in the text.

The first thing to note is that the limited number of tweet posts contain the searched keyword. The share of posts including the possible causes of the 2021-2023 inflation surge, such as “Coronavirus”, “Ukraine”, “Yen’s depreciation”, “Rising raw material prices”, or “Bank of Japan” is less than 1 percent of the corresponding set of posts (Columns (2) and

[4]The weighted Jaccard (J) between word sets A and B is calculated as

$$J(A,B) = \frac{\sum_{i \in A \cap B} \min(w_A(i), w_B(i))}{\sum_{i \in A \cup B} \max(w_A(i), w_B(i))},$$

where $w_A(i)$ and $w_B(i)$ are frequency of word i in set A and B, respectively. Between 0 and 1, the higher J means the higher similarity between two word sets.

Figure 5: Word Cloud

(1) 2008-2009 (accept price hikes)

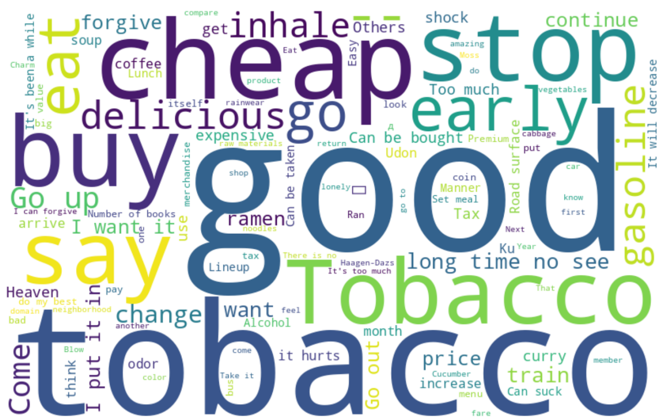

(2) 2021-2023 (accept price hikes)

(3) 2021-2023 (positive valence)

Note: English translation of word clouds in Online Appendix B.

	(1) 2008-2009 (accept hikes)	(2) 2021-2023 (accept hikes)	(3) 2021-2023 (positive valence)	(4) 2021-2023 (all)
Coronavirus	0 (0.0)	107 (0.7)	37 (0.3)	792 (0.6)
Ukraine	0 (0.0)	3 (0.0)	1 (0.0)	80 (0.1)
Yen's depreciation	0 (0.0)	131 (0.8)	65 (0.6)	1,405 (1.1)
Rising raw material prices	1 (0.4)	34 (0.2)	8 (0.1)	111 (0.1)
Bank of Japan	0 (0.0)	2 (0.0)	1 (0.0)	24 (0.0)
Salary	0 (0.0)	535 (3.4)	135 (1.3)	4,911 (3.8)
It can't be helped	25 (11.0)	3,108 (19.8)	207 (1.9)	3,990 (3.1)

Note:

1. The number of appearances of specific words. The share in the respective total posts is in parentheses.
2. (1)-(3) correspond to those in Figure 5. (4) corresponds to posts selected by BERT (step 2 of the Deep Learning Approach).
3. "It can't be helped" corresponds to "shikata-nai", "itasikata-nai" or "shouga-nai" in Japanese.

Table 5: Keyword Search

(3)). It is not surprising to observe that there are no posts citing the first three keywords during the 2008-2009 period (Column (1)), as the phenomena of these words was not prevalent at that time, such that the Yen was 90-110 against the US Dollar in 2008-2009, whereas it was 100-150 in 2021-2023. However, even "Rising raw material prices," which was evident during the 2008-2009 period, was seldom cited. Given that these keywords are not included much in all selected posts from 2021 to 2023 (Column (4)), the tweets' short sentences do not have sufficient word length to describe the reasons behind their sentiment, including the background of the price hikes.

That said, we may still be able to detect the narrative in these limited tweet posts. Reading through individual posts categorized in Column (2), we find some of them telling the story that the Japanese accepted the price hike, as they were caused by something uncontrollable (at least for individuals). Examples of these tweets are as follows (admittedly, the sixth item of the Bank of Japan has a different tone):

1. "Honestly, with COVID and all, raising prices was the only way to stay afloat." [Coronavirus]
2. "Kotobukiya's reissued model kits are getting more expensive, huh. Well, it can't be helped with how things are these days. Still, it's tough when money's already tight.

Honestly, I resent the countries that spread COVID and started wars." [Coronavirus]

3. "So the Russia-Ukraine conflict is hitting us here too (emoticon). I understand the price hikes, but in the end, peace is what we need most." [Ukraine]

4. "Fried chicken getting more expensive? Painful, but I get it. Inflation from the weak yen and wars is just unavoidable." [Yen's depreciation]

5. "Price increases make sense with raw materials getting more expensive, but it'd be nice if paychecks kept up." [Rising raw material prices]

6. "Uniqlo isn't the only one raising prices—lots of stuff will be going up soon. It's not like what the BOJ Governor said; it's just that fuel and other raw materials are getting more expensive, so there's no avoiding it." [Bank of Japan][5]

Second, a change in wage formation may be part of the narrative of changing anti-inflation perception. Although still a limited number, "Salary" came to be cited during the 2021-2023 period, whereas it was not mentioned at all in the 2008-2009 posts. Reading individual posts in Columns (2) and (3) of Table 5, we find that most of the tweets citing salary say that price hikes would be acceptable or welcome (positive valence), if wages were also increased (the fifth item above and the first and the second items below). At the same time, we observe that some posts mentioned that price hikes were acceptable because wages had already been increased (see the third to the fifth items below). Although the number of these posts is still limited, we believe this conveys a significant marginal change from the chronic deflation era.

1. "Fine, keep raising prices—just raise my salary too!"

2. "Not seeing price increases feels like a win, but that's what holds back inflation and wage growth. Price hikes aren't always the enemy."

3. "Gas and electricity costs skyrocketed out of nowhere. (emoticon) It's rough, but my wages went up too, so I'm managing. (emoticon)"

4. "Oh! My salary's going way up next month! For once, I'm happy about the price hikes! (emoticon)"

5. "Lots of stuff is going up in price, but hey, my paycheck's up 9%, so it's fine by me. (emoticon)"

[5] See footnote 1 for another example.

Third, “It can’t be helped” is more frequently cited by posts accepting price hikes (Columns (1) and (2) in Table 5) than those of positive valence (Column (3)). This is by design, as part of our deep learning analysis. Some posts do not explicitly show the intention of accepting price hikes, but have a positive tone, such as “Price hikes everywhere today — but hey, my paycheck went up too! (emoticon) Pretty much matches what I got back in my overworked manager days. (emoticon)”.

5 Conclusion

This paper applies a natural language processing technique to tweets that comment on price hikes to see whether there has been any change in the Japanese anti-inflation perception or zero inflation norm (Watanabe, 2022, 2024). Using three different approaches (one lexicon and two deep learning), the paper finds that, after inflation shocks of COVID-19 and the Ukraine war hit Japan, more tweet posts have revealed positive tone, willingness to accept price hikes, and even pleasant feelings (valence) compared to the previous inflation episode between 2008 and 2009. Together with evidence from the Nikkei newspaper and the Economy Watchers’ Survey, it is very likely that, overall, the Japanese have increased tolerance to price hikes for various goods, while some have increased anger.

In an attempt to uncover narratives behind the shift in anti-inflation perception, the paper finds that (i) people have come to accept price hikes when these price hikes were triggered by uncontrollable exogenous factors, such as the pandemic and the war, and (ii) they have also become more accepting of price hikes and developed a positive attitude when their salaries have increased.

That said, the number of these posts was limited. This could be because tweet posts are too short to convey the narratives effectively. Alternatively, it could also be because the change in inflation tolerance has been so marginal. It is worthwhile to continue monitoring SNS to see which (or both) is the case, using more powerful computation and advanced techniques that are sure to come in the future.

References

Ahrens, M. and M. McMahon (2021): "Extracting Economic Signals from Central Bank Speeches," in *Proceedings of the Third Workshop on Economics and Natural Language Processing*, Association for Computational Linguistics.

Aoki, K., H. Ichiue, and T. Okuda (2019): "Consumers' Price Beliefs, Central Bank Communication, and Inflation Dynamics," Bank of Japan Working Paper Series 19-E-14, Bank of Japan.

Bank of Japan (2024): "Review of Monetary Policy from a Broad Perspective," https://www.boj.or.jp/en/mopo/mpmdeci/mpr_2024/k241219b.pdf.

Bollen, J., H. Mao, and X. Zeng (2011): "Twitter Mood Predicts the Stock Market," *Journal of Computational Science*, 2, 1–8.

Braig, N., A. Benz, S. Voth, J. Breitenbach, and R. Buettner (2023): "Machine Learning Techniques for Sentiment Analysis of COVID-19-Related Twitter Data," *IEEE Access*, 11, 14778–14803.

Dell, M. (2024): "Deep Learning for Economists," NBER Working Papers 32768, National Bureau of Economic Research, Inc.

Diamond, J., K. Watanabe, and T. Watanabe (2020): "The Formation of Consumer Inflation Expectations: New Evidence from Japan's Deflation Experience," *International Economic Review*, 61, 241–281.

Ehrmann, M. and A. Wabitsch (2022): "Central Bank Communication with Non-Experts – A Road to Nowhere?" *Journal of Monetary Economics*, 127, 69–85.

Giachanou, A. and F. Crestani (2016): "Like It or Not: A Survey of Twitter Sentiment Analysis Methods," *ACM Computing Surveys*, 49, 1–41.

Heddaya, M., C. Tan, R. Voigt, Q. Zeng, and A. Zentefis (2025): "A Century of Inflation Narratives," https://papers.ssrn.com/abstract=5259107.

Heddaya, M., Q. Zeng, C. Tan, R. Voigt, and A. Zentefis (2024): "Causal Micro-Narratives," ArXiv:2410.05252 [cs].

Hori, M. and M. Kawagoe (2011): "Inflation Expectations of Japanese Households," *Economic Review*, 87, 993–1008.

Ito, T., H. Sakaji, K. Tsubouchi, K. Izumi, and T. Yamashita (2018): "Text-Visualizing Neural Network Model: Understanding Online Financial Textual Data," in *Advances in Knowledge Discovery and Data Mining*, ed. by D. Phung, V. S. Tseng, G. I. Webb, B. Ho, M. Ganji, and L. Rashidi, Cham: Springer International Publishing, 247–259.

KAMADA, K. (2013): "Downward Rigidity in Households' Price Expectations: An Analysis Based on the Bank of Japan's 'Opinion Survey on the General Public's Views and Behavior'," Bank of Japan Working Paper Series 13-E-15, Bank of Japan.

KAMADA, K., J. NAKAJIMA, AND S. NISHIGUCHI (2015): "Are Household Inflation Expectations Anchored in Japan?" Bank of Japan Working Paper Series 15-E-8, Bank of Japan.

KORINEK, A. (2023): "Generative AI for Economic Research: Use Cases and Implications for Economists," *Journal of Economic Literature*, 61, 1281–1317.

——— (2024): "LLMs Level Up—Better, Faster, Cheaper: June 2024 Update to Section 3 of "Generative AI for Economic Research: Use Cases and Implications for Economists," Published in the," *Journal of Economic Literature*, 61, 1–38.

KURODA, H. (2022): "The Bank's Thinking on Monetary Policy: Toward Achieving the Price Stability Target in a Sustainable and Stable Manner," Speech at the Kisaragi-kai Meeting in Tokyo, June 6, 2022.

KWON, B., T. PARK, F. PEREZ-CRUZ, AND P. RUNGCHAROENKITKUL (2024): "Large Language Models: a Primer for Economists," *BIS Quarterly Review*, 37.

MOHAMMAD, S. (2016): "A Practical Guide to Sentiment Annotation: Challenges and Solutions," in *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, ed. by A. Balahur, E. van der Goot, P. Vossen, and A. Montoyo, San Diego, California: Association for Computational Linguistics, 174–179.

NAKAJIMA, J., H. YAMAGATA, T. OKUDA, S. KATSUKI, AND T. SHINOHARA (2021): "Extracting Firms' Short-Term Inflation Expectations from the Economy Watchers Survey using Text Analysis," Bank of Japan Working Paper Series 21-E-12, Bank of Japan.

NISHIZAKI, K., T. SEKINE, AND Y. UENO (2014): "Chronic Deflation in Japan," *Asian Economic Policy Review*, 9, 20–39.

OKUN, A. M. (1981): *Prices and Quantities: A Macroeconomic Analysis*, Washington, D.C.: Brookings Institution.

RUSSELL, J. A. AND L. F. BARRETT (1999): "Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant," *Journal of Personality and Social Psychology*, 76, 805–819, place: US Publisher: American Psychological Association.

SHAPIRO, A. H., M. SUDHOF, AND D. J. WILSON (2022): "Measuring News Sentiment," *Journal of econometrics*, 228, 221–243.

WABITSCH, A. (2024): "The Messenger Matters," https://papers.ssrn.com/abstract=5238987.

WATANABE, T. (2022): *On the Meaning of Prices (Bukka to ha Nanika)*, Tokyo: Kodansha, in Japanese.

——— (2024): *Understanding Inflation and Deflation (Bukka wo Kangaeru)*, Tokyo: Nikkei Business Publications, Inc., in Japanese.

ZIMBRA, D., A. ABBASI, D. ZENG, AND H. CHEN (2018): "The State-of-the-Art in Twitter Sentiment Analysis: A Review and Benchmark Evaluation," *ACM Trans. Manage. Inf. Syst.*, 9, 5:1–5:29.

Online Appendix

A Examples of Posted Tweets

1. 値上げをして給料が上がる、好循環までもう少しだ(emoticon of fire) いつまでもデフレマインドじゃダメだ(emoticon of fire)
2. いつも注文頼んでるお弁当屋さんがきた。物価高騰の影響で値上げの依頼。仕方ないよねぇ。。。
3. ガソリンがいきなり5円値上がり....。勘弁していただきたい。
4. ガソリン30円値上げて。クルマ通勤の私はもう血吐きそうなんですけど。(TдT)
5. ガソリン値上げ決定かあ。。。。廃屋は180円台だぞ。。。
6. なんか存した気分になるなぁ・・・
7. シャウエッセンサイレント値上げした？
8. もう、Sにしようとかにしようとか思わない... BとかCでもいい... サインあったら、Sとか思ったけど... 値上がりしてて、サインもないなら... （愚痴ってすまん）
9. んー、なんか正直この値上げはエアプな気もしてきたな ただでさえ初期からUAV持ちスキンが有利だったのに更にそれを加速させそう まぁ財布システムがある以上UAVは保険としてのアイテムに過ぎないかな UAVだってたかが数十秒しか効果無いし

B Word Clouds in Japanese original

(1) 2008-2009 (accept price hikes)

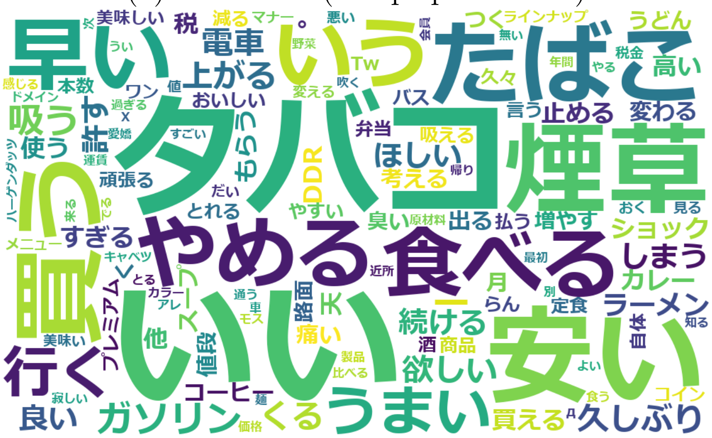

(2) 2021-2023 (accept price hikes)

(3) 2021-2023 (positive valence)

C Annotation Guideline (excerpt)

1. 主題が投稿時点の直近（前後の1ヶ月以内）に起きた・起こる、自身が価格影響力を持たない消費財・サービスの価格変化について言及している可能性が高い。
2. 消費財・サービスの価格の変化について、ポジティブ・ネガティブな批評的表現(言及された価格変化を受容しているかどうかが判別できる表現)や要素（Emoji、ネットスラングを含む）が含まれる。

上記の２つの要素を満たすPostを今回の対象とし、選別していきます。Postの内容をもとに、対象に当てはまるかを次の3つの確信度から選択してください。

- 明確に1と2の両方に当てはまる（確信度：高）: 主題は最近の消費財・サービスの価格変化である可能性が高く、明確に中立でない、批評的表現（ポジティブ・ネガティブ）を含む。
- 1と2の両方に当てはまると明確には言い切れない（確信度：中）: 比喩・婉曲によって明確とは言い切れない;「驚く」「やばい」「w」などの両義的な言葉; 中立的な表現にとどまる; 固有名詞などで、トピックを把握できない; コンテキストによって文章の含意が変化する。
- 1と2の両方に当てはまるとは言い難い（確信度：低）: 主題が明確に1と異なる;「値上げします」のような価格交渉的なやりとり; 投資・資産運用に関するニュース; ゲームキャラに関する評価; 批評的な表現が含まれていない; 値下げに関するニュース; 「値下がりしたから買う」「値下がりを待つ」という表明のみ。

D Prompt to GPT

以下の指示に従ってください。
まず、ポストのテキストを注意深く読み、値上がりに関する言及部分を抜粋します。
その後、抜粋した表現が「受け入れ」または「許容」を示しているかどうかを文章全体から判断します。もし該当する表現が「受け入れ」や「許容」を示している場合は「LABEL_1」を、それ以外の場合は「LABEL_0」をラベルとして割り当てます。
以下にいくつかの例を示します。

\#\#\# 例1
post: するめうまい!値上がりしてたのは許せんが、うまいので良し。
expression: 値上がりしてたのは許せん
label: LABEL_0
reason: 値上がりを許せないという表現のみで、値上がりを受け入れたり許容したりする表現が見つからないため。

\#\#\# 例2
post: CPエボの値上がりマジやばたにえん......45,000キロ修復なしワンオーナーとかいう奇跡みたいな個体を160万で買えてよかった......
expression: CPエボの値上がりマジやばたにえん
label: LABEL_0
reason: 「マジやばたにえん」という表現のみで、値上がりを受け入れたり許容したりする表現が見つからないため。

\#\#\# 例3
post: マックまた値上げかあ...もう買ってないからどうでもいいけど...(˘ω˘)
expression: マックまた値上げかあ
label: LABEL_0
reason: 値上がりを受け入れたり許容したりする表現が見つからないため。

\#\#\# 例4
post: Netflix、値上げは全然構わないから哲仁王后の配信を...お願い...#Netflix
expression: 値上げは全然構わない
label: LABEL_1
reason: 「全然構わない」という表現から許容的なニュアンスが読み取れるため。

\#\#\# 例5
post: あっぽーみゅーじっく値上げすんのよな...どうしよう...でも還元率高いのはあっぽーなのよね...
expression: あっぽーみゅーじっく値上げすんのよな
label: LABEL_0
reason: 値上がりを受け入れたり許容したりする表現が見つからないため。

\#\#\# 例6
post: 順当に値上げしててわろた
expression: 順当に値上げしててわろた
label: LABEL_0
reason: 「わろた」という表現のみで、値上がりを受け入れたり許容したりする表現が見つからないため。

これらの例を参考に、与えられたポストを分析し、次のJSON形式で結果を提示してください:

\‘\‘\‘json \\

\{ \\
"expression": "<抜粋された表現部分>", \\
"label": "LABEL_0" or "LABEL_1", \\
"reason": "～という表現から許容的なニュアンスが読み取れるため。" or "値上がりを受け入れたり許容したりする表現が見つからないため" etc... \\
\} \\
\`\`\`

受容や許容を示す表現が見つからない場合でも、値上がりに言及しているだけの表現は「LABEL_0」としてください。
この指示についてユーザーと話し合う必要はありません。あなたの唯一の目標は、指示に従って正確なJSONレスポンスを提供することです。

E Prompt to GPT (Russell)

```
XのPostについて、Russellの感情円環モデルを用いて分析します。
提示されたPostを読み、値上がりに対する感情を、覚醒度と快・不快の2つの観点から-2，-1，0，1，2の5段階で評価してください。
その際、以下の指針に従ってください。
1．値上がりについての言及があるかどうかを確認し、ない場合には覚醒度、快・不快のどちらも0としてください。
2．Post全体の感情の評価ではなく、値上がりに対する感情の評価に集中してください。
3．覚醒度と快・不快の両者について、そのように評価する理由を記入してください。
4．結果を以下のJSON形式で出力してください。
‘‘‘json
{
    "arousal": -2 or -1 or 0 or 1 or 2,
    "valence": -2 or -1 or 0 or 1 or 2,
    "reason": "値上がりに対する感情の評価の根拠を記入してください。"
}\
‘‘‘
この指示についてユーザーと話し合う必要はありません。あなたの唯一の仕事は、上記の指示に従って文章を分析し、正確なJSONを出力することです。
```