Mining Indonesian Tweets to Understand Food Price Crises

UN GLOBAL PULSE

METHODS PAPER, FEBRUARY 2014

TABLE OF CONTENTS

Executive Summary

Context

Food prices have a direct effect on the purchasing power of a large part of the Indonesian population, and increases pose a threat to household food security, particularly when inflation affects the price of staple foods such as rice or soybeans. Particularly for poor households, food accounts for almost 75 percent of total spending. Government's occasional efforts to reduce fuel subsidies have been known to drive up food prices. It is the government's concern to respond to these shocks and try to mitigate their negative impact, as early as possible.

The use of social media is widespread in Indonesia; the country has the fourth largest Facebook population in the world, and the third largest number of Twitter users worldwide. The 20 million user accounts in Jakarta make it the city with the largest Twitter presence in the world.

Objective

Operating on the premise that online social media conversations might represent a new source of information to monitor food security, this research analyses Twitter conversations related to food price increases amongst Indonesians during the period from March 2011 to April 2013. This research also explores the relations between such conversations, food price inflation and external events.

Methods

Taxonomies (groups of words and phrases with related meanings) relevant to food and fuel price increases were developed in the Bahasa Indonesia language in order to identify relevant content. Using Crimson Hexagon's ForSight software, a classification algorithm was trained to categorize the extracted tweets as positive, negative, confused, or neutral in order to analyze the sentiment of these food price-related tweets. Using simple time series analysis we quantify the correlation between the volume of food-related Twitter conversations and official food inflation statistics, and between food and fuel-related tweet volumes. Spot checks using qualitative method have also been done in several cities.

Results

We found a relationship between retrospective official food inflation statistics and the number of tweets speaking about food price increases ($r=0.42$). We later found, upon analyzing fuel price tweets, that there was a perceived relationship between food and fuel prices. In particular, we found a significant correlation ($r=0.58$) between the two topics suggesting that even potential (rather than realised) fuel price rises affect people's perception of food security.

Discussion

Our research shows that automated monitoring of public sentiment on social media, combined with contextual knowledge, has the potential to be a valuable real-time proxy for food-related economic indicators. In addition, social media analysis can be used to uncover people's reactions to fuel discussions that affect public perception of food issues. If analysis includes geographical mentions, it could help to differentiate the variability among cities/regions.

Current challenges to overcome include how to establish high frequency models of food prices and validate them using official statistics, how to filter out noise due to non-relevant news items and how to harness the potential of inferring demographics.

If social media data mining to model food prices matures to become robust in the future, statistical institutes might consider including social media monitoring into official statistics channels.

Acknowledgements

This paper summarizes the findings and methods from a research project conducted by Pulse Lab Jakarta in 2012-2013. Pulse Lab Jakarta is a joint initiative of the Government of Indonesia, through the Ministry of National Development Planning (Bappenas), and the United Nations, through Global Pulse. The research efforts of Pulse Lab Jakarta focus on testing the viability of using new sources of digital data and real-time analytics to support development goals and strengthen social protection.

This project was conducted in collaboration with the Indonesian Ministry of National Development Planning (Bappenas), UNICEF and WFP in Indonesia, with the support of Crimson Hexagon.

This document has been drafted, edited and produced by Pulse Lab Jakarta in collaboration with the UN Global Pulse team.

For more information on this or other projects facilitated through the Global Pulse Lab network, please visit: http://www.unglobalpulse.org/research

Introduction

Traditional statistics, household surveys and census data have been effective in tracking medium to long-term development trends, but are less effective in generating a real-time snapshot in order for policymakers to develop timely actions to protect vulnerable populations against crises. As the Secretary-General's High Level Panel on Post2015 noted in the report section entitled 'Wanted: A Data Revolution'[1], better data and statistics will help governments to track progress and ensure their decisions are evidence-based.

The High Level Panel's call for a data revolution acknowledges that today there is an ocean of data—generated by citizens in both developed and developing countries—that did not exist even a few years ago. This data is passively generated by people simply by living their daily lives. Mobile phones, social media and Internet searches all leave digital traces that, when anonymized, aggregated and analyzed, can reveal significant insights that help governments make faster and more informed decisions.

One of the major sources of real-time digital data in Indonesia is Twitter. With over 20 million Twitter accounts (1 person in 12), Indonesia ranks third in the world in the number of active Twitter users[2]. Jakarta recently emerged as the "most tweeting" city on earth[3], sending more tweets than London, Tokyo and New York. This wealth of data presents an opportunity to extract real-time insights about publicly shared interests and issues pertinent to the Indonesian population.

In Indonesia, populations have been particularly exposed to food price increases since 2010: the Food Price Index has been growing at a higher rate than the overall Consumer Price Index (CPI)[4]. This inflation is compounded by the price of rice, which has a direct link to Indonesian households' food security and rose 51% from December 2009 to February 2012.

This research project analyzes the volume of Twitter conversations in Bahasa Indonesia about food and fuel price increases and tries to infer real-time information regarding how price increases are perceived by the Indonesian population.

Mining Indonesian Tweets to Understand Food Price Crises presents the context, methods and results of the research. It shares the detailed taxonomy developed and used to monitor and categorize the conversation about food prices in Indonesia and the subsequent quantitative analysis. Finally, suggestions for further research in the field are proposed, based on the research findings.

---

[1] (2012) High level report on post-2015 development agenda http://www.post2015hlp.org/the-report/
[2] Felix Richter (2013). Twitter's Top 5 Markets Account for 50% of Active Users - Statista. Retrieved November 21, 2013, from http://www.statista.com/topics/737/twitter/chart/1642/regional-breakdown-of-twitter-users/.
[3] (2012). The World Cities That Tweet the Most - Richard Florida - The Atlantic Cities. Retrieved November 21, 2013, from http://www.theatlanticcities.com/arts-and-lifestyle/2012/08/world-cities-tweet-most/2944/.
[4] WFP (2012). Monthly Price and Food Security Update, Indonesia, March 2012. Retrieved from http://home.wfp.org/stellent/groups/public/documents/ena/wfp246211.pdf.

Social Media for Development

The rise of social media has been accompanied by a plethora of research on the techniques of mining social media to detect opinions, trends and consumer patterns. *Salathe et. al* recently completed research mining Twitter data for anti-vaccination sentiment, in an effort to understand how negative sentiment can spread via online communities[5]. UNICEF published a paper in April 2013 on anti-vaccine sentiment on social media across Eastern European, including Facebook, Twitter, forums and blogs[6]. It aimed to monitor specific concerns related to vaccines, identify influencers in online communities and develop strategies to counter anti-vaccination campaigns. Several researchers have mined Twitter and other social media for opinions on movies and so box office revenue[7] and to predict future stock price behaviour [8].

Certain topics are better suited to social media analysis than others. *Asur* and *Huberman* list the conditions for a topic to be a good candidate for analysis in their paper studying Twitter's predictive value for box office revenues[9]. Namely, the topic has to be widely discussed on Twitter (or in some social media outlet) and real-world outcomes have to be easily verifiable. In detecting opinions *Pang and Lee* discuss steps in identifying texts (in this context a 'text' represents any human generated linguistic content such as a tweet or a blog entry) that are of interest[10]. First, relevant texts can be filtered based on topic, followed by an assessment of whether the texts are objective or subjective, after which their polarity (i.e. whether a text is negative, neutral or positive) can be assessed and finally the intensity of their opinion.

There are two broad approaches to automated sentiment analysis of texts; unsupervised learning approaches and supervised learning approaches. In its simplest form, the former approach uses sets of single words with known positive and negative meanings such as

> Positive: 'great', 'good', 'improvement', 'happy'
> Negative: 'terrible', 'poor', 'sad', 'tragic'

The number of words in each text with positive meaning is then compared to the number of words with a negative meaning to give an overall sentiment score. Such an approach, however, struggles to correctly classify 'not great' as having negative sentiment since it simply sees 'great' in the list of positive words and infers positive sentiment. Other subtleties include slang meanings of words that have an alternative sentiment compared to their formal use. One such term in Bahasa Indonesia is "nganggur" or "don't have job", which when used casually can also mean "doing nothing" in the positive context of relaxing during free time. Hence the algorithm, when training the machine, should

---

[5] Salathé, M., Vu, D. Q., Khandelwal, S., & Hunter, D. R. (2013). The dynamics of health behavior sentiments on a large online social network. *EPJ Data Science*, *2*(1), 1-12.
[6] (2013). Tracking anti-vaccine sentiment in Eastern Europe - Unicef. Retrieved November 21, 2013, from http://www.unicef.org/ceecis/Tracking_anti-vaccine_sentiment_in_Eastern_European_social_media_networks.pdf.
[7] Joshi, M., Das, D., Gimpel, K., & Smith, N. A. (2010). Movie reviews and revenues: An experiment in text regression. *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
[8] Bollen, J., Mao, H., & Pepe, A. (2011). Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. *ICWSM*.
[9] Asur, S., & Huberman, B. A. (2010). Predicting the future with social media. *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*. IEEE.
[10] Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, *2*(1-2), 1-135.

incorporate the context of their use. In Indonesia there are 300 local dialects so analysis shouldavoid examining terms in isolation but must also look at the context of their use.

More broadly, while well-established collections or 'corpora' of words with known positive or negative sentiment exist in English and other major languages (e.g. Linguistic Inquiry and Word Count[11]) these are much less developed in other languages of interest for development work.

Supervised learning, on the other hand, requires human classification of example texts as positive or negative. Computer algorithms then 'learn' how to determine if a new text is positive or negative from these examples. While supervised learning requires some human effort in training the algorithm unlike its unsupervised counterpart, the analysis has the advantage of being generally more context specific and therefore more accurate.

In evaluating polarity, *Pang and Lee* discuss features to search for, including keywords that indicate emotion, position of key words and parts of speech[12]. Joshi et al. discuss the use of n-grams (a string of n words) that are topic specific or indicate emotion in their work on detecting opinions on movies via text mining[13]. *Pang, Lee* and *Vaithyanathan* attempt a more statistical approach to sentiment classification, using models to quantify the probability of a text's polarity given the presence of various features[14].

Current research discussed above, supports the notion that Twitter can be used to analyze public sentiment on food prices in real time. Food price fluctuations are widely discussed and the effects are easily observable (e.g. protests). Furthermore there is a large and growing body of research on techniques to categorize relevant tweets and/or use supervised training methods to automatically mine social media texts.

[11] Linguistic Inquiry and Word Count www.liwc.net
[12] Ibid.
[13] Joshi, M., Das, D., Gimpel, K., & Smith, N. A. (2010). Movie reviews and revenues: An experiment in text regression. *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
[14] Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up?: sentiment classification using machine learning techniques. *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10.* Association for Computational Linguistics.

Research Questions

This project attempted to provide answers to three main questions:

a. Are people talking about food price increases on Twitter? If so, how?

To answer this, we evaluated the extent to which food price increases are discussed on Twitter and what are the sentiments (positive or negative) around such conversations.

b. How does this information compare with ground truth information?

In particular, the project focused on understanding the volume of Twitter conversations about food price rises compared with actual food price hikes as confirmed by official statistics.

c. Are people also talking about fuel price increases on Twitter? If so, how does it relate to the Twitter conversations about food prices rising?

We evaluated the extent to which fuel price increases are mentioned and compared the number of messages against the ones related to food prices.

Data

Twitter

For this study, the Twitter firehose – that is the entirety of tweets as they are posted in real-time – was mined using Crimson Hexagon's social media monitoring platform, ForSight. The platform grants access to historical Twitter content and to various tools for searching, analyzing and reporting on the data (keyword filtering, supervised classification and a dashboard for visualizing graphs, time series and word clouds).

The project focused on tweets about food price increases, however tweets about fuel price increases were also monitored to analyze possible links between fuel price-related and food price-related conversations.

Official Statistics

Through the Indonesian State Ministry of National Development Planning (BAPPENAS) and the World Food Program (WFP), we collected official prices and used them as a baseline for comparison with the Twitter data. In particular we used general foodstuff CPI data from the Indonesian Office of Statistics (BPS) and used milk and rice price data from the WFP. Since local soybean shortages have led to the import of American soy, US soybean inflation data was collected from the World Bank as a proxy.

In addition, we identified the timeline of relevant events during the period of our investigation, between March 2011 and April 2013:

1. 24th-27th July: Soybean and its derivative products shortage (Tempeh and Tofu)
2. 31st July-30th August 2011, 19th July-18th August 2012: Ramadan (Holy month for Muslims)
3. 18th-19th November: Government of Indonesia's initial fuel subsidy cut plan

Methodology

Crimson Hexagon continuously collects tweets that are then stored in a database, which allows their users to investigate content retrospectively. The software includes a classification algorithm that provides a means to measure the proportions of specific opinions or themes that are present in large, text-based data sets.

The general tools used for data collection, categorization and analysis are called monitors. The following steps outline the process of setting up, running, and using the monitors.

Step 1: Define the Data Set

The dataset used in this study includes all publicly available tweets coming from the Twitter firehose from March 2011 until April 2013. The massive growth in Twitter use over this period is reflected in an increasing volume of tweets over time in each of the monitors.

Food Price increase-related Tweets

Public reaction to food price increases (harga makanan naik), such as staple food (sembako), rice (beras), as well as eggs (telur) and milk (susu) were considered.

Fuel Price increase-related Tweets

Public reaction to price increases of fuel (harga bahan bakar naik) for both cooking and transport, including kerosene (minyak tanah), as well as cuts in government fuel subsidies were considered. As a result of cuts to fuel subsidies, funds were generated that were intended to improve social protection programs such as BLSM (unconditional cash transfer), JAMKESMAS (health insurance for the poor) and BSM (cash stipends for poor students to improve access to basic education). Mentions of cuts to fuel subsidies were also considered relevant.

Step 2: Filter the Data Set

While the metadata for users coming from the Twitter firehose provides language and some location tagging, it still requires further filtering. For example, a bilingual user might set the default language for her account as Bahasa but frequently tweet in English adding noise to the analysis. Therefore we further filtered the tweets to isolate those in the Bahasa Indonesian language. Indonesia is a multilingual country but since Bahasa Indonesia is the predominant tweeting language, we can assume the sample to be representative of Twitter traffic originating from Indonesia.

Next, to determine content that might be relevant to the subject of analysis, a broad keyword filter was used to identify which tweets from the firehose are on topic. The condition for relevance based on keywords is outlined below. Each word is given in Bahasa Indonesian with the English translation in brackets. All words are nouns except where indicated. In order to be considered relevant a tweet must contain "harga" (price) and "naik" (rise) along with one or more food commodities.

harga (price)

AND

naik (rise) v.

AND

sembako (groceries) OR makanan (food) OR pangan (food) OR beras (rice) OR gula (sugar) OR minyak goreng (cooking oil) OR daging ayam (chicken) OR daging (meat) OR daging ayam (chicken) OR daging sapi (beef) OR telur, telor (egg) OR susu (milk) OR cabe/cabai (chilli) OR tepung (flour) OR kedelai, kedele (soy beans)[1] OR tahu, tempe (tofu, tempeh)[2]

[1] "kedele" is variation of pronunciation of "kedelai" in Indonesia spoken language
[2] "tahu" and "tempe" are both derivative products from "kedelai"

Table 1: Food Price Rise Taxonomy. Words are nouns unless otherwise specified.

The taxonomy for fuel price rises was also developed, it is similar to the previous taxonomy developed for food prices; to be considered relevant the tweet must contain "harga" (price) and ("naik" (rise) or "kenaikan" (increase) or "mahal" (expensive)) together with the words that refer to the type of fuel commodities (diesel, fuel, or LPG).

harga (price)

AND

naik (rise) v. OR kenaikan (increase) OR mahal (expensive/high) adj.

AND

bensin (gasoline) OR premium OR pertamax (premium) OR minyak tanah (kerosene) OR mitan (kerosene) OR solar (diesel) OR BBM (fuel) OR bahan bakar (fuel) OR gas OR elpiji (LPG) OR LPG

Table 2 Fuel Price Rise Taxonomy. Words are nouns unless otherwise specified.

Essential language expertise and local knowledge was provided by Pulse Lab Jakarta, BAPPENAS, and an extensive network of collaborators. Some example tweets are listed below.

"Waduh bbm naik harga makanan naik saya bs agak kurusan ntar"
"Whoa fuel price rise, then food price rise, then I will become slimmer"---Mar 26 2012

"BBM naik, harga makanan juga naik x_x"
"Fuel price has hiked, now food follows x_x "---Mar 27 2012

"Bagi rakyat kecil bkn masalah u/ mngurangi pnggunaan BBM, tp efek domino dr knaikan BBM (harga pangan, trnsportasi umum naik) yg memberatkan"
"For the poor, reducing fuel usage is not the problem, but the domino effects of fuel price rise is (food, public transportation price rise)"---Mar 30 2012

Step 3: Categorization and Analysis

Before analysis, relevant categories were defined by a domain expert. Based on the manual classification of some sample data ("training"), an algorithm then analyzed the proportions of data that fall into the previously defined categories. The classification process therefore involves both manual and automated processes. The first step is for a researcher to manually classify randomly selected posts. Posts that are not clear or can fit into more than one category are skipped during training. When each category has sufficient training posts the monitor is run and the algorithm automatically classifies each further tweet collected by the monitor.

The categories created are given below:

- Positive (tweet that indicates positive emotion)
  Example:
  "Mungkin satu-satunya manusia yang suka jikalau harga cabai naik adalah istriku"
  "Maybe the only person that happy whenever chili price goes up is my wife"
  ---7th April 2013

- Negative (tweet that indicates negative emotion)
  Example:
  *"Harga bensin naik.. Harga makanan pasti naik juga.. sedihnya mahasiswa adlh uang bulanan gak naik.. *hiks"*
  "Fuel price is rising..definitely food price will go up as well.. It's sad for students because their monthly allowance doesn't follow"
  ---30th April 2013

- Confused/wondering (tweet that indicates some confusion)
  Example:
  "Harga kebutuhan pokok kini samakin meningkat, apa usaha pemerintah???"
  "Food price is rising, what is the government effort to tackle that problem???"
  ---2nd July 2012

- Realised price rise/high-no emotion
  Example:
  "Nah, kalo gaji PNS naik, BBM naik, dampaknya adl kenaikan harga berbagai komoditi kebutuhan pokok seperti bawang, cabai, gula, daging, dsb."
  "So, if the government employee rise, the fuel price rise, the effect will be on the increase of several staples commodities such as onion, chili, sugar, meat, etc."
  ---30th Apr 2013

Results

This section presents the research results. The analysis is of Twitter conversations relevant to food price rises, identifying events that trigger conversations, fuel price rise conversations and analysis of their relation to food price rise conversations, and finally measurement of correlation between food conversation with the official food price inflation data.

In total 113,386 tweets were collected, with 12% classified as being of positive sentiment, 32% as negative, 33% as confused and the remaining 23% with no sentiment.

Description of Twitter Conversations about Food Price Rise

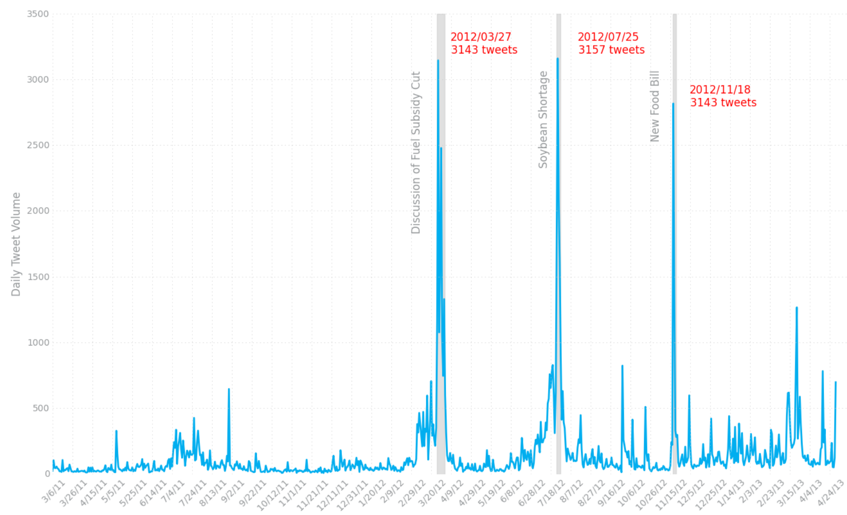

Figure 1: Daily Tweet Volume Related to Food Price Rise (March 2011 - April 2013)

Figure 1 shows the daily number of food price rise related tweets. During the timeframe considered, it demonstrates a significant range in the volume of conversation related to food prices rise, between virtually zero to more than 3,000 tweets per day, with 3 significant spikes occurring in 27th March 2012, 25th July 2012, and 18th November 2012.

Government of Indonesia's Discussion Regarding Fuel Subsidy Cut (26th March - 2nd April 2012)
The clear increase in the volume of food price increase-related tweets between 26th and 31st of March 2012 coincides with discussion around potential 33% fuel subsidy cuts by the Indonesian government at the end of February 2012 —which in 2011 had accounted for 20% of total government expenditure. This led to large protests in response to which the Indonesian government did not implement these proposals.

In particular on March 30th 2012, the Jakarta Post reported that "more than 5,000 workers from industrial areas in and around Jakarta staged a mass demonstration at the front gates of the legislative compound, demanding the House of Representatives (DPR) turn down the government's plan to increase fuel and electricity prices in April and May respectively." The protest in Jakarta was part of several others happening in the country's main cities, which were successful in halting the policy, and according to the Financial Time's blog, Beyond Brics, "pushed Indonesia's opposition to reject the government's plan to cut spending on fuel." [15]

Soybean Shortage (24th - 27th July 2012)
In July 2012, driven by sharp rises in soybeans imported from the US, the government introduced the emergency measure of reducing import taxes. Despite this, many households suffered from the increase in the price of this staple and many businesses suffered. Homegrown production in Indonesia has been unable to keep pace with demand.

New Food Bill (18th-19th November 2012)
This chatter is coincidental with a law proposal, finally passed on November 18, to establish a new food agency in Indonesia with policymaking authority[16]. The announced goal of the agency was to facilitate the decision-making process of the different ministries and government bodies involved in food issues, ultimately helping Indonesia to reach self-sufficiency in staple foods, including rice and soybeans.

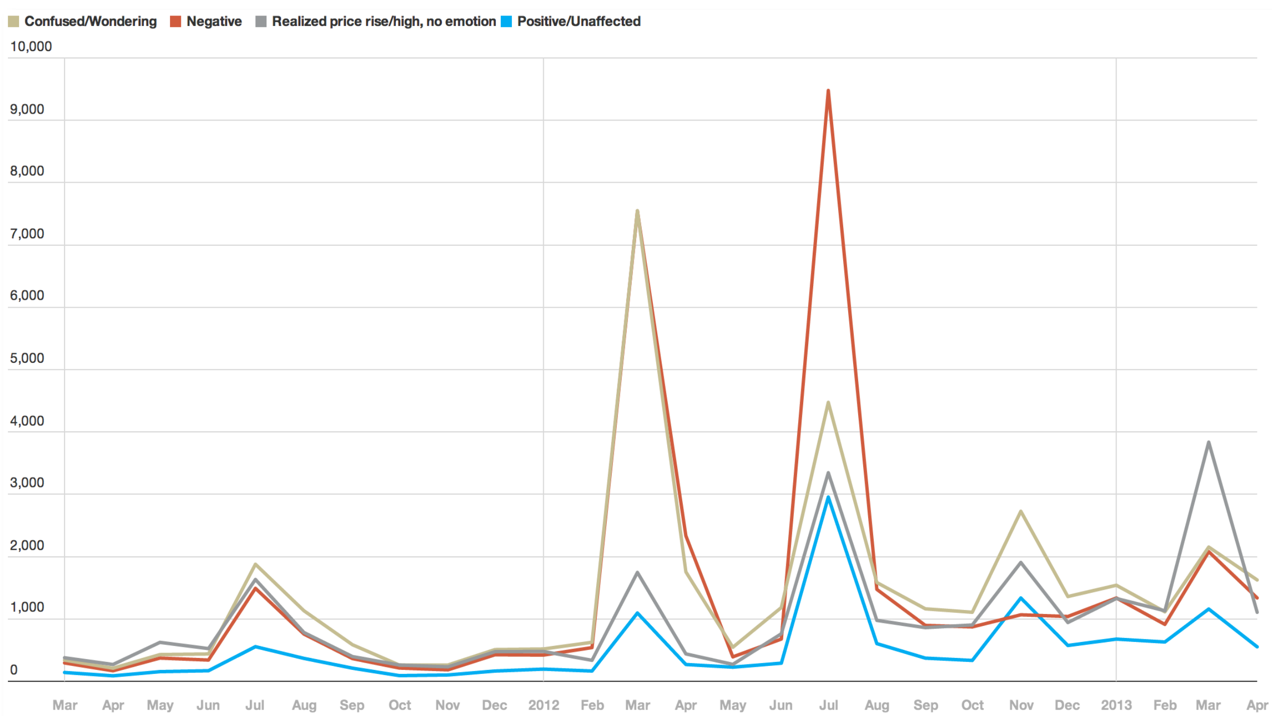

Figure 2: Public Sentiments for Food Price Rise in Twitter Conversation (March 2011 - April 2013) Volumes of food related tweets classified as displaying positive, negative, confused or neutral sentiment. Monthly averages are shown for clarity.

---

[15] (2012). Indonesia: fuel subsidy cut runs into protest and politics | beyondbrics. Retrieved November 21, 2013, from http://blogs.ft.com/beyond-brics/2012/03/30/indonesia-fuel-subsidy-cut-runs-into-protest-and-politics/.
[16] (2012). High hopes pinned on new food agency | The Jakarta Post. Retrieved November 21, 2013, from http://www.thejakartapost.com/news/2012/10/22/high-hopes-pinned-new-food-agency.html.

Next the sentiment of the food-related tweets was analysed as shown in Figure 2. Predictably, the majority of tweets related to food price increase show confused or negative sentiments, particularly around the soybean shortage and fuel subsidy announcement.

Relation Between Twitter Conversations and Official Food Price Inflation Data

Figure 3 below shows the monthly averaged time series of food price-related tweets along with the monthly food CPI inflation statistics provided by BPS. The gray shaded region corresponds to a clear outlier that is identified as the initial announcement from Government of Indonesia about fuel subsidy cut (March 2012) which triggered massive Twitter conversations. The correlation was calculated both on the full time range as well as excluding this datapoint.

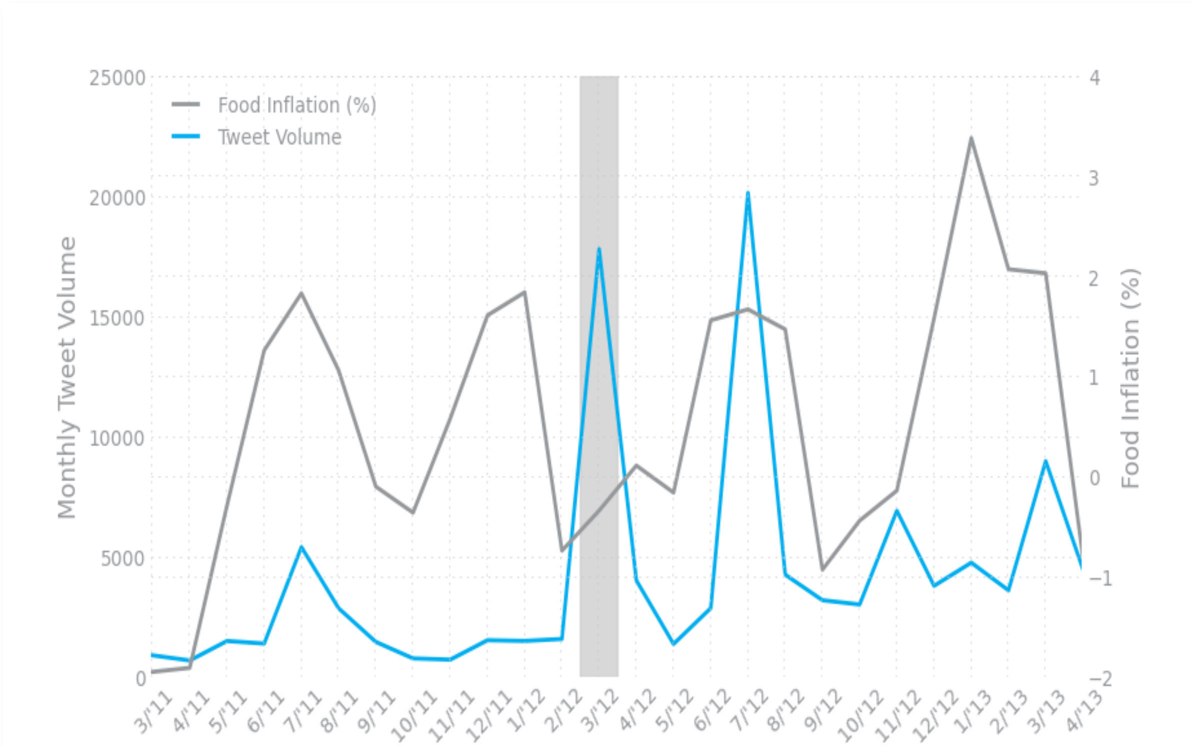

Figure 3 Plot of Monthly Food Price-Related Tweet Volume with Official Food Price Inflation Statistics. The grey highlighted area marks the month of March 2012 during which proposals on fuel subsidies cuts were under consideration by the Government of Indonesia.

For further analysis, correlation was measured between the number of tweets classified as demonstrating different emotions for each sentiment category and the food CPI. Values calculated including the outlying month of March 2012 are shown in brackets. The correlation coefficient quantifies the degree to which the two time series move up together or down together; it lies in a range between -1 (moving in exactly opposite directions) and 1 (moving in exactly the same direction) with a value of 0 representing no correlation. The p value essentially quantifies the probability that the same r value could be found using random data[17].

The weakest correlation, and the only high p-value, was seen with tweets classified as being of negative sentiment, while the strongest correlation was seen with tweets of neutral sentiment, potentially showing that neutral tweets are more factual.

[17] 0.05 is the accepted threshold for statistical significance in the literature.

Emotional Dimension	R	P
Positive	0.41 (0.39)	0.04 (0.05)
Negative	0.26 (0.18)	0.21 (0.37)
Confused	0.48 (0.55)	0.01 (0.004)
Neutral	0.57 (0.55)	0.003 (0.003)
All	0.42 (0.32)	0.04 (0.12)

Table 3: Correlations between Public Sentiment on Food Price Rise and Official Food CPI Data

We also examined the volume of tweets related to specific food items and their individual inflation indicators. We found that movements in global soy prices correspond with social media traffic regarding a wide variety of foodstuffs; milk, rice, soy and general foodstuffs all correlate significantly (Pearson's correlation coefficient lying in the range 0.42-0.66).

Figure 4 shows the time series of each quantity, the tweet volumes have been rescaled by their maximum. In July 2012 the volume of soy related tweets reached nearly 15,000, roughly 10 times greater than the peak of the other tweet volumes. The rise in US soy prices had a knock-on effect on conversations around soy in Indonesia. Tweets related to other foodstuffs also experienced a significant peak in the same month. Interestingly, this suggests a degree of interconnectedness not only between the roles of different foods as local households invoke coping strategies but also through global supply chains[18]. As well as the ability to factor international price movements into their local price calculations, consumers relate the increase of soy prices to potential future movements in different foodstuffs through coping strategies.

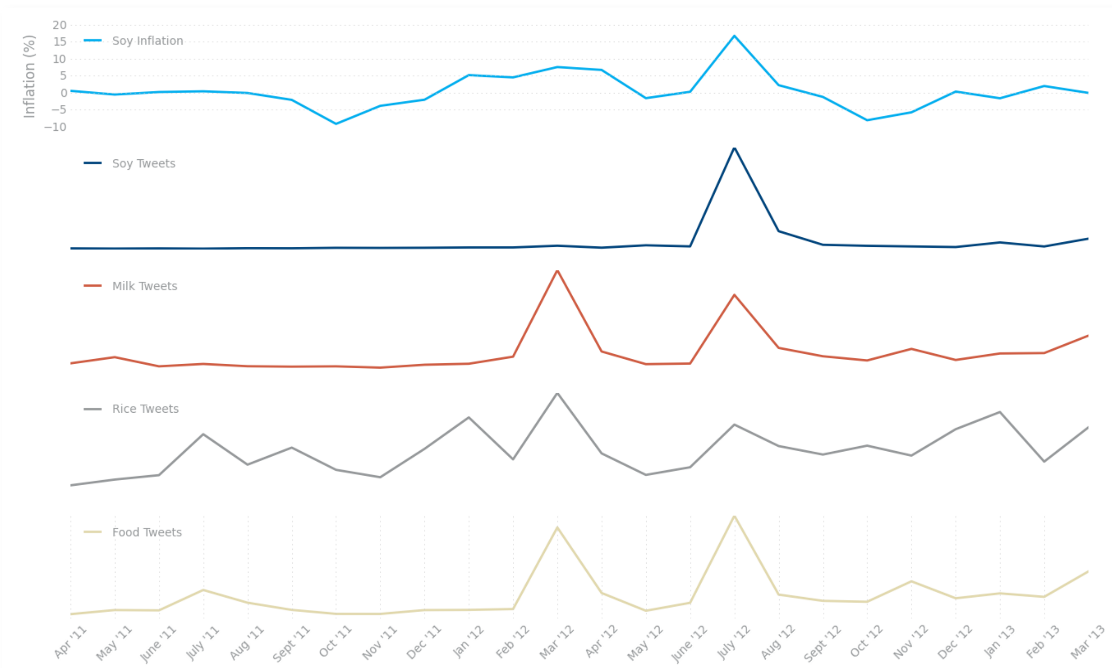

Figure 4 Plot of Normalized Monthly Tweet Volumes for Specific Foodstuffs and Soy Inflation Data

[18] (2005) The Globalization of Food Systems: A Conceptual Framework and Empirical Patterns, The Food Industry Center, University of Minnesota (retrieved 27th november 2013 http:
[18]//ageconsearch.umn.edu/bitstream/14304/1/tr05-01.pdf)

Twitter Conversations about Fuel Price Rises

We see that a significant spike in fuel price tweets coincides with a spike in food price tweets. We therefore investigated the relation between these two series. Together with food price, a fuel price rise monitor was also launched using taxonomy in Table 2. After the monitor produced the results, the correlation between the two was measured to investigate the relation between food price and fuel price hike.

Figure 5 also shows the relationship between food price rise and fuel price rise related tweets (see fuel-related taxonomy in Table 2). Interestingly we see a moderate correlation between the daily tweet volumes relevant to food and fuel; $(r,p)=(0.58, p<10^{-10})$ suggesting that the prices of the two commodities are related. Clearly the conversations about the fuel subsidy announcement led to an increase in fuel-related tweets. It is possible that people were able to make the likely causal connection that the predicted fuel price increase would be reflected in the price of food. However, the opposite is not true: spikes in food traffic were not matched by an increase in fuel traffic.

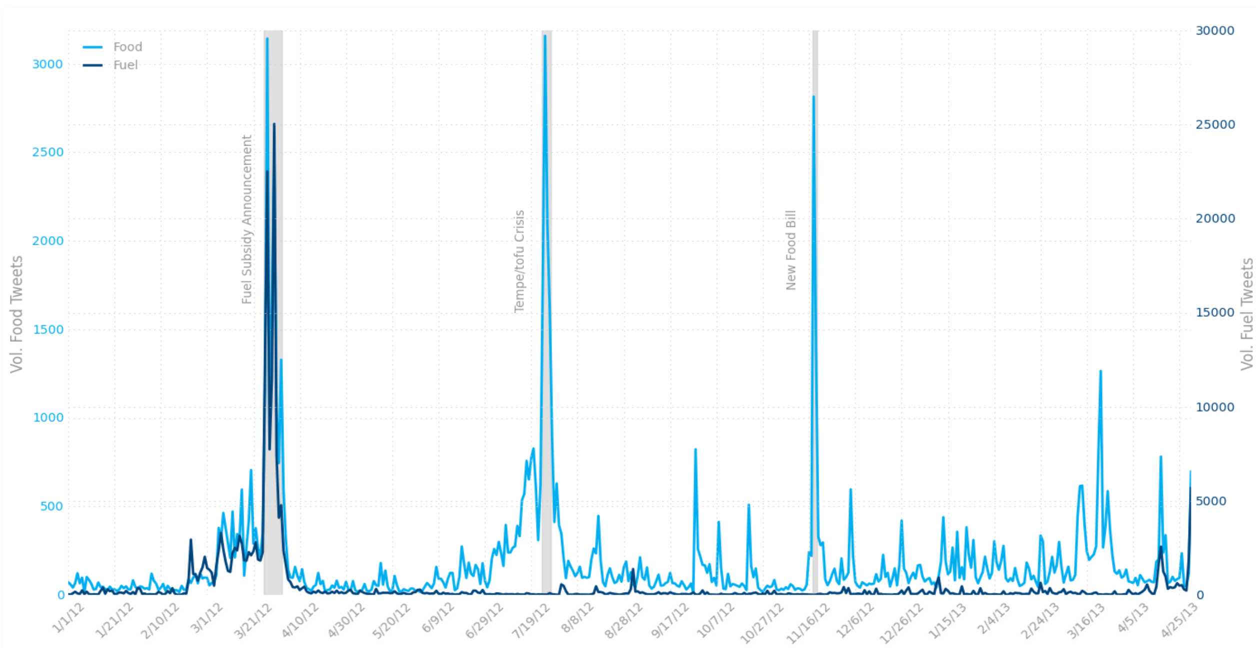

Figure 5 Plot of daily food and fuel related Tweet Volume Related to the Food and Fuel Price Rise (January 2012 and - April 2013)

Conclusion, Recommendations and Further Research

In this study we have investigated how Twitter use in Indonesia reflects changes in food prices. In particular, we have seen some indications that real price movements are reflected in conversations on the topic of food. Further, our taxonomy has shown how different food staples are discussed and how these different conversations reflect official statistics.

We have shown that even a basic analysis of the volume of tweets related to food price rises shows a relation with official statistics on CPI. In our analysis we have found a moderate Pearson correlation coefficient ($r=0.32$, $p=0.12$) between the two time series. While the promise of such an analysis is compelling, we also find evidence that such automated mining of social media streams must continue to be combined with 'smart' domain specific knowledge. For instance, we observe a clear 'false positive' in our data; spikes in Twitter traffic with no corresponding underlying increase in inflation. This occurs around the publication of a high profile news article (26th March) related to fuel that led people to speculate about potential future food price increases. Omitting this clear outlying data point from our analysis increases the correlation noticeably ($r=0.42$, $p=0.04$).

The research presented here represents a proof-of-concept demonstration that semi-automated sentiment analysis of social media streams can demonstrate significant correlation with official, ground truth statistics. Now that the potential of such techniques has been verified, further work is necessary both to improve the accuracy of the category classification, 'nowcasting' food prices from Twitter conversation, and also to refine the technique to provide more fine-grained analysis. Future developments should allow for strengthening of early warning systems and predictive models. Furthermore, techniques are emerging to investigate trends with demographics, such as filtering users by age, gender, and locations.

Somewhat ironically, more fine-grained official statistics would be necessary to conduct a more detailed calibration. We have a daily record of Tweet traffic, but since our food inflation 'ground truth' data was aggregated monthly it is necessary to throw away much of the detail in Twitter content by aggregating monthly (down-sampling) in order to compare the two. A finer temporal resolution would not only give the advantage of giving more agile policy recommendations, on the scale of days rather than months, but would also allow for more sophisticated time series analysis.

We presume that daily Twitter volumes have a well-defined baseline or 'normal' number each day and that any deviation is either due to (1) some underlying event, such as a sharp increase in food price, or (2) small fluctuations within a well defined range; this is the assumption of stationarity. However, due to the increasing popularity of Twitter, it is likely that over the studied period of several years that this baseline rate is increasing. That is to say, with more tweets on all topics over time, we will observe more tweets on food price increases over time and this trend should be accounted for.

Further, we implicitly assume that Twitter conversations respond linearly to increases in food prices, that is to say an increase of X in the price of food leads to an increase of Y tweets and that a further increase in X will lead to a further increase of Y tweets.

It may be that very large jumps in prices will lead to a disproportionate increase in Twitter traffic or even qualitatively different manifestations of negative sentiment i.e. protests[19]. A further non-linear effect comes from the presence of 'influencers' in the network of Twitter users; a user with a larger following or more authority will likely give rise to a larger degree of negative sentiment than a user with a smaller following.

The idiosyncratic nature of Tweet content, e.g. using emoticons, slang and other cultural references, also requires the application of context-specific knowledge in the human training stage. As the rewards of automated Twitter analysis become clearer it is likely that efforts to develop techniques specifically tuned to extract meaning from Twitter content will increase.

Having clearly demonstrated the responsiveness of social media streams to underlying changes in food prices, we recommend that policymakers continue to build on this research and refine the methodology in several key ways. Firstly, our findings are remarkably accurate given that we have considered a country-level aggregation of social media conversations. In a decentralised country such as Indonesia, there is a clear need to spatially and temporally disaggregate content. This requires robust 'geolocation'; the process of mapping of a user's offered textual description of their location i.e. 'Jakarta' to a latitude/longitude coordinate; (-6.2, 106.8).

An alternative mechanism to analyse food price changes is to directly extract numerical price values mentioned in human generated content such as "I just paid $4 for a loaf of bread! What's going on". Another key aspect of food security is identification of coping strategies - substituting expensive items with cheaper alternatives. While both of these techniques require more sophisticated textual analysis, there is the clear advantage of a more direct means of evaluating the food stress within households. Thus, there is the potential for a real-time map of food prices and food stress, which would be invaluable for policymakers.

Building these capabilities inside governments and the public sector will require specific training to selected public service officials. Finally, if this kind of analysis becomes robust and mature in the near future, statistical institutes might consider including social media monitoring into official statistics channels.

For more information on Global Pulse's research please visit: http://www.unglobalpulse.org/research

[19] Lagi, M., Bertrand, K., & Bar-Yam, Y. (2011). The food crises and political instability in North Africa and the Middle East. *Available at SSRN 1910031*.