

ANALYSING REFUGEE-HOST COMMUNITY NARRATIVES

ON

SOCIAL MEDIA

in Lebanon



*Empowered lives.
Resilient nations.*

FEBRUARY 2019
PREPARED BY BUILD UP
FOR UNDP LEBANON

BUILD UP



ANALYSING
REFUGEE-HOST COMMUNITY
NARRATIVES ON
SOCIAL MEDIA

BUILD UP 

FEBRUARY 2019

PREPARED BY BUILD UP (JACOB LEFTON, MAUDE MORRISON, MIRA EL MAWLA & HELENA PUIG LARRAURI) FOR UNDP LEBANON

Contents

Executive Summary	4
Key findings	4
Next steps	5
a) Monitoring and analysis	5
b) Engaging polarized users	5
Introduction	6
Methodology	7
Methodological limitations	9
Part 1. Key findings (mapping)	11
A) To what extent are refugee-host community issues discussed on social media in Lebanon?	11
B) On what platforms are these issues most likely to be discussed?	12
C) What are the issues that impact host community – refugee relations that are discussed online?	16
D) Are these issues discussed to an extent whereby meaningful sentiment analysis can be conducted to identify trends and spikes?	20
E) Are there instances when these spikes coincide or lead to violent incidents, or shape public discourse after violent incidents?	25
F) Who are the people / organizations discussing these issues online and to what extent are they ‘influencers’ of others?	26
G) To what extent is it possible to geo-locate where more hostile sentiments are being expressed?	30
H) To what extent are hostile sentiments being expressed by users outside of groups and pages on Facebook?	30
I) Does the current regulatory framework (both national legislation and social media platform rules and regulations) enable the identification of and engagement with polarized users?	30
J) What are the challenges with APIs for social media monitoring and user engagement and how can they be overcome?	31
Part 2. Opportunities and challenges (intervening)	32
A) Based on the findings for questions I-VII, what kind of analysis and tools would be useful in monitoring and analyzing social media?	32
B) Challenges and opportunities of a tool to engage polarized online users on refugee-host community relations	33

Executive Summary

Key findings

- Refugee-host community issues are widely discussed on social media in Lebanon and analysis of these discussions provides significant insight that can complement offline monitoring. Social media has become a 'magnified mirror' of societal tensions in Lebanon and appears to have a real impact on perceptions of refugee-host community issues.
- Key methodological challenges to social media analysis of refugee-host community tensions include access to APIs, Arabic language in automated data collection and sentiment analysis. The report proposes ways to manage these challenges.
- Facebook and Twitter are both popular platforms for discussions, although there are important differences between them. Facebook reaches a broader subset of the Lebanese population, while Twitter serves predominantly political and intellectual elites. Facebook is a space for less filtered commentary and greater interactions between users, whilst Twitter is used predominantly for news bites, uses more formal language and is the site of less direct conversations between users.
- Across these conversations, an ongoing general discourse about Syrian refugees can be readily identified. Issues of return, assistance and crime and violations emerge as key topics. On the whole, assistance provoked more supportive sentiment, whilst return provoked a greater spread of supportive and antagonistic comments. Crime and violations can provoke both reactions, depending on the nature and direction of the crime and violations.
- Whilst coding for sentiment directed towards and from the Syrian refugee community, the majority of data contained attitudes expressed by Lebanese towards Syrian refugees. This reflects the nature of the conversation more broadly, where sentiments expressed by Syrians towards Lebanese appear less often in public fora and are predominantly restricted to private WhatsApp conversations.
- Sentiment on social media appears to spike in response to 1) overall political discourse (eg a Tweet or a public statement by a major political figure); 2) macro-level events that affect the refugee community (eg Storm Norma) and 3) smaller scale incidents involving Syrian individuals (eg an individual crime).
- Where spikes in sentiment (supportive or antagonistic) do occur, they do not represent the absolute level of sentiment among social media users. Rather, they reflect an awakening of individual users' pre-existing sentiments, triggered by the three causes above.
- Whilst responses to events provide interesting insights into sentiment discourse, it is not enough to understand only conversations responding directly to events. To gain a full picture, it is necessary to see the whole range of conversations happening in a given time period. For example, whilst direct responses to Storm Norma were overwhelmingly supportive on Twitter, antagonistic sentiment relating to broader issues of return spiked in the immediate aftermath of the Storm, affecting the overall discourse.
- Whilst this data provides an accurate temperature gauge of offline sentiments, it is very difficult to draw predictive conclusions from the data. In this period, online conversations served to mirror offline events, rather than to trigger them.
- Amid the Twitter conversations, five distinct communities emerged. These communities each operate within their own sphere of influence. The English-speaking community, which includes the majority of UN agencies and INGOs, does not exert much influence over the domestic Lebanese audience.
- Polarization between several Twitter communities appears to be occurring on the basis of political leaning, rather than on the basis of attitudes towards refugees.
- Each of these communities responded differently to different events, demonstrating distinct interests.

Next steps

As a result of these findings, further research is recommended to refine the methodology and gather additional data. As a next step, that research should be built on and efforts made to actively engage with the social media space through strategic communications and depolarization initiatives.

a) Monitoring and analysis

- Overall, quantitative analysis of the kind produced in this report would be a useful complement to traditional monitoring of social tensions. An extended and more sophisticated version of the tools and processes used in this study could be developed to provide relevant results on a rolling basis. The tool could include:
 - ◊ A refined list of keyword combinations, including a full list of Arabizi terms and alternative spellings to support further analysis.
 - ◊ A refined sentiment analysis function, to further divide supportive and antagonistic sentiment into sub-categories.
 - ◊ API access for the specific study.
 - ◊ Inclusion of other refugee populations in the scope of the study.
- Such a tool could be integrated into the UNDP workflow to complement existing tensions monitoring.
- Recognizing that this work is of interest to many organisations and researchers in Lebanon, efforts to engage with other stakeholders working on the issue, including sharing findings from this report, would be beneficial.
- The ability to quantitatively monitor and analyse WhatsApp is unlikely to become a possibility in the medium term.

b) Engaging polarized users

- There is an emerging consensus that more could and should be done to combat the increasingly toxic nature of online conversations.
- Any further interventions should shift away from early warning and move towards conflict prevention through interventions aimed at shifting online narratives.
 - ◊ Concretely, social media analysis could be used to design a strategic communications campaign to influence online attitudes and behaviours relating to refugee-host community tensions.
 - ◊ Social media influencers with a broad reach could be mapped and equipped with the tools to better shape local narratives.
 - ◊ Any strategic communications campaign should be connected to offline dialogues, and integrated into UNDP's existing initiatives. The use of UNDP community committees under the Mechanisms of Social Stability programming could provide one such avenue.



Introduction

In November 2018, UNDP commissioned Build Up to conduct a comprehensive mapping and analysis of Syrian refugee-host community relations in Lebanon (including narratives, sentiments and amount of traffic) on Facebook and Twitter. The report was designed as a feasibility assessment to determine whether it would be viable for UNDP to consider further social media monitoring based on financial, ethical and resource considerations. The report covers findings from a desk review, qualitative interviews and analysis of over 2.6 million Tweets and 14,734 Facebook posts and comments.¹ Building on this analysis, the report also serves to identify challenges and opportunities for UNDP to integrate social media monitoring into its programming. Concretely, we lay out considerations for UNDP to build and operate: a) a tool to regularly monitor online narratives on refugee-host community relations; b) a tool to engage with users who are exposed to polarized refugee-host community narratives.

1 The difference in volume between Twitter and Facebook is due to a difference in API access. A full explanation is given in the methodology section below.

Methodology

The methodology for this research, developed in consultation with UNDP, followed a three-phase approach comprising of a desk review, qualitative interviews and data analysis.

The desk review covered an analysis of reports relevant to the topic of Syrian refugee-host community tensions and social media narratives in Lebanon. This was complemented by Key Informant Interviews (KIIs) with 12 stakeholders. These included four representatives of UN agencies, two from academic institutions, three from NGOs and civil society organisations, two independent researchers and one representative of a research and consulting firm. The objective of these interviews was two-fold: 1) understand the context in which social media is affecting Syrian refugee-host community tensions; 2) gather inputs for the social media analysis methodology. Following the desk review and KIIs, the following methodological decisions were made:

First, the decision was made not to analyse narratives on WhatsApp due to ethical concerns over privacy: the only way to quantitatively analyse WhatsApp conversations is for researchers to personally enter groups and download data from those groups. The quantitative scope of this study is thus restricted to Facebook and Twitter, although some qualitative findings relating to WhatsApp and Instagram are presented in section 1B below.

Second, the scope of the study was restricted to the Syrian refugee population. As a result, the term 'refugee' refers to the Syrian refugee community for the remainder of this report.

Third, Build Up sought access to APIs from Facebook and Twitter to enable comprehensive data collection, citing the objectives and nature of the research. However, our applications for access were rejected despite multiple attempts and outreach to contacts at both companies. This made it impossible to fully automate data collection on Facebook, limiting the scope of the data collected. This is a challenge pertinent to much social media analysis. As a result, data collection was conducted in English and Arabic, from January 1st - January 31st 2019, as follows:



Twitter: We used personal API access in order to analyse a large volume of relevant tweets on an automated basis. Based on the desk review, manual Twitter scans and interviews, we developed a list of 168 keywords or phrases considered relevant (in Arabic and English), as well as a seed list of 79 accounts that regularly tweet about the topic, whose content informed further refinement of the keywords. Our system automatically downloaded tweets that contained any of these keywords or phrases from both seed users and other users. From January 1st - 31st, we gathered 2,641,107 Tweets. Following regular monitoring of data and current events, we refined the final keyword list to select a subset of highly relevant data.

Facebook: Given that we did not have access to the Facebook API and concerns surrounding the alternative web-scraping approach, manual data collection was considered the only viable option for Facebook. To assist manual data collection, the SingleFile tool was used for data-recording.² Based on the desk review, we developed a list of 54 relevant Facebook pages and profiles that we visited every day during the data collection period. All 54 pages were specific to Lebanon. 41 posted entirely in Arabic, 13 included English language posts (comments were often in both languages). We would manually read each post from the day and determine if it was of interest to the study according to a set of criteria. If it was deemed relevant, we used SingleFile to copy the relevant data into a spreadsheet, which was then automatically transferred to an SQL database. Only public posts were copied into the database. Throughout the data collection period, we gathered 14,734 posts, comments and replies. All posts were anonymised in the database.

² <https://chrome.google.com/webstore/detail/singlefile/mpiodijhokgodhhofbcjdecppffjipkle?hl=en>

Categorisation: Throughout the data collection period, we sought to further categorise the data by a) sentiment b) issue type and c) events. We developed sets of keywords and phrases for Twitter and Facebook datasets to denote each of these categories. Sentiment keywords were developed for the overall dataset (general sentiment) and for specific events (event sentiment). We focused the event categorisation on four events that occurred during January that had relevance to the topic:

Storm Norma

On January 6th, a major storm hit Lebanon, affecting large numbers of refugees in the Beqaa valley and triggering a widespread humanitarian response.

Aarsal incident

On January 18th, a Syrian man stabbed a Lebanese woman in a shop in Tannourine. The man was arrested on January 26th.

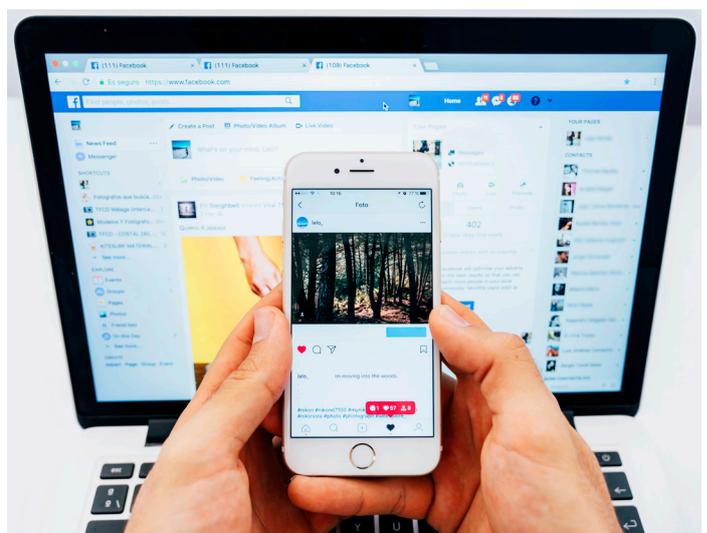
The death of Ahmad Zoabi

On January 19th, Ahmad Zoabi, a Syrian child who worked as a shoe shiner on the streets on Beirut, was found dead in a building after being chased by Beirut municipal police.

Tannourine stabbing

On January 21st, a group of youth destroyed Syrian shop fronts, cars and properties in Aarsal. A video of the incident was widely shared on social media

		
Quantity of data	2,641,107	14,734
Arabic language posts	761,147	-
English language posts	1,286,668	-
Tweets from users located in Lebanon ³	56,900	-
Highly relevant data ⁴	25,310	8,134
Number of Retweets	1,376,474	-
Number of distinct users	883,781	10,309



³ This is likely to be an underestimate of Lebanon-specific tweets since many users do not list their location (see section G)

⁴ For Twitter - included a term from our final keyword list and stated their location as Lebanon. For Facebook - the parent post included a term from our final keyword list

⁵ This is likely to be an underestimate due to technical challenges associated with the user id identification

Methodological limitations

This was a scoping study, in part an evaluation of what is possible in terms of narrative analysis on social media in Arabic and English. As a result, several learnings have emerged about the limitations of this approach. These learnings inform the recommendations at the end of this report.

- Restricted scope.** Due to limited time, the study was restricted to narratives relating to Syrian refugees in Lebanon. This ignores other important refugee populations in Lebanon and issues of social tensions relating to them. Recommendations for further research thus include an expansion to additional populations, such as the Palestinian refugee community.
- Access to Twitter data.** Lack of access to APIs created considerable ethical challenges for the study. Whilst personal API keys for Twitter are currently allowed, the specific project was not approved, making the methodology limited in its long-term replicability unless API access is secured specifically for the project. Despite these challenges, the ethical guidelines laid out at the outset of this study were met.⁶
- Sampling of Facebook data.**
 - Manual data collection:** The lack of API access significantly reduced the quantity of Facebook data that we could collect. In addition, such a manual process is inevitably open to human error, resulting in the potential for relevant data from selected pages to be missed. The use of keyword filters in the database enabled us to search the Facebook dataset for relevant content in the same way we were able to search the Twitter database. Regardless of these mitigating measures, the Facebook data should not be considered a fully representative sample of the Facebook sphere, but rather a potentially biased snapshot of relevant conversations.
- Keyword use in online conversation:** Although searching for keywords is effective for Twitter, where users are more likely to mention keywords directly through hashtags, Facebook users are less likely to mention topics within statuses or comments or use hashtags in their posts. This means that on a post related to Storm Norma for example, only a main post would mention the word Norma itself, while comments tend to be reactionary and generic. In order to mitigate this we categorised the top post as relevant to an event, but included all responses to that post as relevant, even where they did not include event-specific keywords. We followed the same approach for issue categorisation.
- Arabic in automated data collection.**

Inclusion of “Arabizi”.⁷ The use of English letters to spell Arabic words is common practice for Lebanese users on social media, particularly on Facebook and among younger users. Although technically able to code for Arabizi in the database, the lack of uniform spellings made it impossible to accurately capture Arabizi terms within the scope of this study - as a result we did not include them in our keywords list. Our qualitative analysis revealed that such spellings are much less common on Twitter than on Facebook: when searching on Twitter for the word ‘lej2in’ to denote refugee, we found almost no results. On Facebook, whilst such searches were not included as keywords in the database queries, we were able to download posts that included Arabizi and do include several examples of such posts in our qualitative analysis. The below table illustrates the challenges of identifying uniform Arabizi spellings.

⁶ See Annex II: ethical guidelines

⁷ Also known as Arabglish, Arabglizi, or Arabic Chat Alphabet

Arabic word/ English translation	Arabizi spellings	Text example from data
عاصفة/storm	3asfe	Awal 3asfe esma norma
	3asifa	#Live #Snow #Norma #3asifa
بكرة/tomorrow	bokra	akid lyom 2abel bokra
	boukra	Akid alyom abl boukra
	bkra	Lym abel bkraa

Standard spelling. This challenge was coupled with further challenges related to Arabic language, such as the use of similar roots to denote different words, or the same word with different short vowel sounds which are not usually shown when writing colloquially. The Arabic word Tannourine **تنورين**, for example, can also be pronounced *tnawwreen* with the same spelling. The latter means to illuminate, usually with one's presence. This resulted in a need to constantly review the automated data collection and revise keywords on a rolling basis.

Sentiment analysis. Coding for sentiment is key to being able to analyse narratives on social media semi-automatically. However, most topics - including Syrian refugee-host community tensions - do not lend themselves to straightforward coding. In this case, a term that can be considered negative towards refugees such as hate or **كره** can be used to denote positive sentiment when placed within a larger frame. For example, see the following Tweet relating to Storm Norma:



Follow

بكره كنتزي ودفايتي... وبكره لون النار وريجة الدفا.
وبكره ضحكنا وراحتنا ولحافنا وقت يكون في اهل
وولاد وختايرة عايشين بجحيم الثلج.
مخيم عرسال للمظلومين السوريين.

Translate Tweet



Translation: "I hate my sweater and my heater. I hate the sound of fire and the smell of warmth. I hate our laughter and comfort and covers when there are parents and grandparents living in the hell of snow. [Below is] the AAarsal camp of oppressed Syrian refugees."

Furthermore, individual events tend to provoke different phrases and keywords that denote positive or negative sentiment specifically for that event, but are not used again in broader conversations. For example, the term 'racism' was widely used in the Ahmad Zoabi case to show solidarity towards Syrian refugees and to criticise the actions of the Beirut municipality, whilst in other cases the term was used to denote negative sentiment. Even within responses to a particular event, sentiment coding cannot be defined as zero-sum - there is much ambiguity within expressions of opinion that such black and white coding is unable to accommodate. To mitigate for this challenge, we developed specific coding for each major event, as well as a more general list of supportive and antagonistic sentiment keywords, and complemented the automated analysis with an in-depth qualitative review of the data. More details of this approach are outlined in section D.

Key findings (mapping)

A) To what extent are refugee-host community issues discussed on social media in Lebanon?

Social media in Lebanon is undoubtedly an outlet for discussion of issues related to refugee-host community relations. The volume of content drawn from Twitter over a one-month period (2.6 million Tweets worldwide) reveals the extent of these discussions. Major political, media and entertainment personalities engage in these conversations (see section F) and individual posts on these topics have a wide reach.

Qualitative interviews reflected unanimous views that social media has become a 'magnified mirror' of societal tensions in Lebanon.⁸ Due to the volume of relevant traffic, difficulties in discerning facts from falsity online, and limited culture of critically interpreting sources, social media is widely thought to have a real impact on individuals' perceptions and serves to reinforce existing narratives.

In a UNHCR survey from 2016, over two thirds of respondents acknowledged that discussions in the media (separate but considered closely linked to social media by multiple interviewees) trigger reactions against refugees.⁹ Interviewees noted the particular importance of social media in reinforcing the 'us and them' narrative, recognised in the UNDP Bar Elias report as a narrative which 'increasingly pits one homogenous Syrian community against an equally homogenous Lebanese community'.¹⁰ At the same time, social media can and does serve as a platform for expressions of inter-communal solidarity, as evidenced both during and in the aftermath of Storm Norma.

Where refugee-host community issues are discussed, the conversation does not always explicitly refer to keywords defined as relevant to the topic. Where a topic is mentioned explicitly, reactions (comments and replies) often do not repeat the relevant keywords, as it is demonstrably clear to readers what they are talking about. This is particularly relevant on Facebook, where conversations trend towards reactions to individual posts, rather than Twitter, where retweets include the same language as the original tweet, including hashtags. Of over 14,000 Facebook posts and comments, only 436 referred to the term 'refugee', or لاجئي or نازحين. However, we know that this is a clear underestimation of the amount of posts and comments that refer to the topic of refugee-host community relations, as many of these references had many comments that did not mention the specific language.

On Twitter, the percentage of tweets referring to these terms reflects a greater level of conversation using direct keywords (205,870 of 2.6 million Tweets). However, this too is likely to be an underestimate. Twitter does not give good guidance on the total number of relevant tweets that are drawn from this kind of streaming methodology. They do however offer a paid service providing 10% of all relevant content, suggesting that our access to content is lower than 10%. However, experience suggests that less used keywords deliver a higher percentage of results.

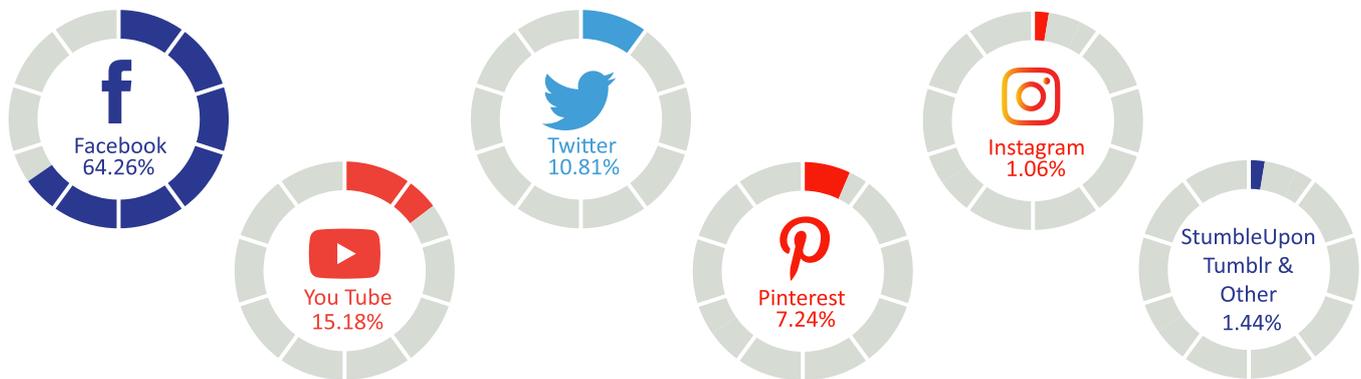
⁸ Interview with Nasser Yassin, Director of Research, Issam Fares Institute for Public Policy and International Affairs, American University of Beirut

⁹ Dylan O Driscoll, *Communication interventions supporting positive civic action in Lebanon*, March 2018

¹⁰ UNDP, *Losing control: Results of a WhatsApp Survey of Syrian Refugees and Host Communities in Lebanon*, 2018

B) On what platforms are these issues most likely to be discussed?

Overview: key differences between social media platforms



StatCounter Global Stats

Social Media Stats Lebanon from Feb 2017 -Jan 2018

Source: <http://gs.statcounter.com/social-media-stats/all/lebanon/2018>

Key findings



Facebook

- Widely used by Syrian and Lebanese communities - the most popular social media platform in Lebanon, accounting for 45.26% of market share February 2017-2018 (data does not include WhatsApp).¹¹
- More informal language used (including Arabizi), making it harder to analyse quantitatively.
- Used for practical purposes e.g. bureaucracy, job hunting, reconnecting with friends and family who have been displaced etc. This use is of particular relevance to the refugee community, although Facebook is also used in this way by Lebanese communities.¹²
- Popular for local pages, such as those focusing on one municipality, or community level updates.
- Posts result in more engagement than Twitter.
- Preferred platform for social media 'influencers' in Lebanon.¹³
- Used by organizations, primarily UNHCR, to reach refugees through private groups and pages.



Twitter

- Rarely used by Syrian refugee community and less intellectual Lebanese.
- Dominated by more intellectual, political and social actors and activists.
- Discussion is less emotional and less violent imagery is shared.
- Tends to have less reach and produce lower levels of engagement than Facebook.
- More formal language used, makes it easier to analyse.
- Platform where Lebanese citizens interact with politicians and policymakers.

¹¹ Ibid

¹² <https://mbasic.facebook.com/lajeenlebanon>, <https://www.facebook.com/-320800848034033> سوريين- في-لبنان

¹³ Interview with Haramoun Hamieh, media researcher



WhatsApp

- Used widely (84% of refugees according to UNDP WhatsApp report) but not included in this study due to data challenges / privacy.
- Although widely used as a means to communicate within the refugee community, it is not commonly used as a means of communication with host communities.¹⁴ People tend to operate within their groups but not as frequently across them. Relied on for news updates, as well as for personal conversations with friends and family. As a result, can be used for more candid and sensitive conversations between close associates.
- Also serves more practical purposes eg Shawishes using it to distribute aid.¹⁵ Groups tend to be formed around specific topics such as aid or local area updates.
- Used by organizations, primarily UNHCR, to reach refugees through private groups.
- User interface and ease of communication also make it a popular choice for local government authorities to give updates at the grassroots level, usually using WhatsApp groups or message chains coordinated by municipality employees and sometimes shawishes of ITSs. Updates here can span from news on the local/village level and distributions from humanitarian organizations, to security updates and ceremonial announcements.
- Often the fastest way for information to spread around a particular event, despite the recent limit in forwarding (to max 10 people). When information does spread, it is often in small bite-size forms (eg a url), providing less space for reactions (unlike Facebook whose design enables lengthy reactions and debate on specific posts).
- Used as a tool for collective action, including sharing invitations to sit-ins, boycotts etc.
- Often preferred platform by refugees, whose financial standings restrict their access to data (often choosing to buy cheaper WhatsApp-only data bundles).
- Also enables illiterate populations to use voice messaging- relevant for the Syrian refugee population, many of whom had never completed primary education.¹⁶



Instagram

- Increasingly important platform for Lebanese audiences, particularly for younger users. 1.4 million active users in Lebanon in 2017.¹⁷
- Used to express visual content.
- Starting to be used by some news sites and influencers to share more political updates.
- Was used during Storm Norma to raise awareness and promote solidarity (see example below).
- Worth exploring further in future research.

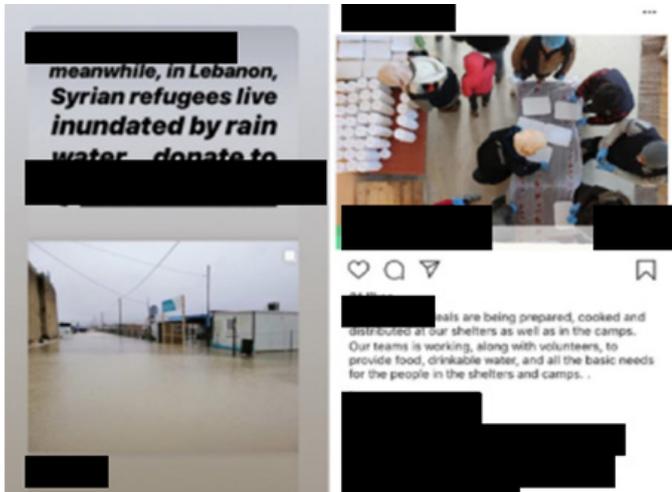
¹⁴ Interview with UNHCR Lebanon, Communications with Communities Unit

¹⁵ UNDP *Losing Control: Results of a WhatsApp Survey of Syrian Refugees and Host Communities in Lebanon, 2018*

¹⁶ Markus Goransson, *Apping and resilience: How smartphones help Syrian refugees in Lebanon negotiate the precarity of displacement*, Clingendael: Netherlands Institute of International Relations, July 2018

https://www.clingendael.org/sites/default/files/2018-07/PB_Mobile_phones_July_2018.pdf

¹⁷ <https://www.slideshare.net/EveryLeader/digital-in-lebanon-2018>



Example of Instagram use during Storm Norma response

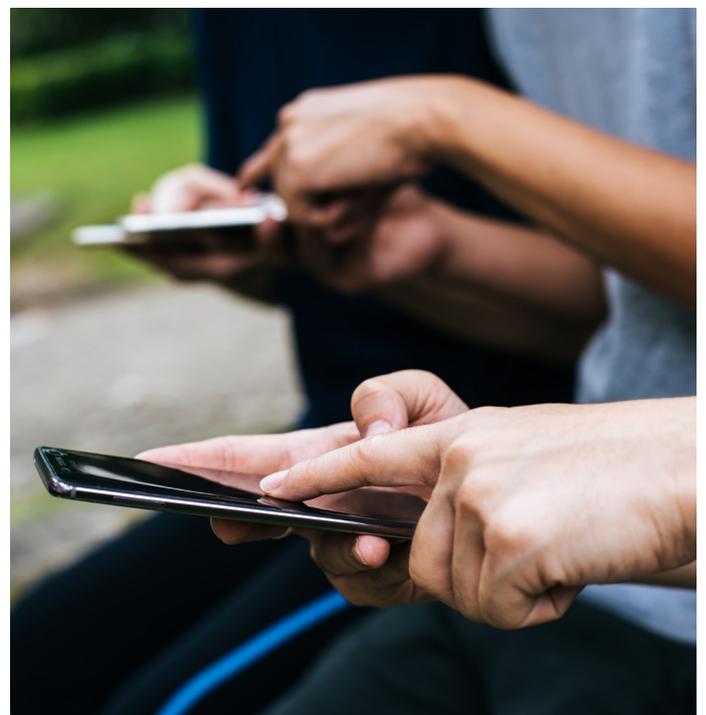
Detail: differences between Twitter and Facebook on the issue of Syrian refugees

Facebook and Twitter are both platforms on which these issues are extensively discussed. However, the way in which discussions unfold and the audiences of these discussions are quite different.

Differences in reach. Facebook reaches a much broader subset of Lebanese and Syrian populations, with individual Facebook posts tending to reach more people than individual Tweets. Although differences in data collection methods limit our ability to compare sheer volumes of Twitter and Facebook conversations, we can compare the extent to which the conversations that do occur spread. While Tweets are spread through Retweets, Facebook posts are spread through comments on posts which results in more users becoming involved in one post.

Differences in tone. Whilst both Facebook and Twitter provide interesting insights, Facebook appears to be a platform for less filtered and therefore more telling commentary around refugee-host community tensions. It is a space that provides for more reactions and longer conversations and debates than Twitter and is more widely used across Lebanon.

Interaction on Facebook. Facebook is also a space in which communities intersect with each other more often through interaction on longer comment threads and public pages that have a broader audience than Twitter accounts. Comment threads are populated with a more informal, discussion-style commentary that is accessible to a wider subset of the population. The function of “mentioning” other users (friends or not) in comments also shows how users engage with other users rather than replying or sharing a post. This contrasts with the Retweet function on Twitter, where users rarely comment on a post but Retweet it from their own account. The Facebook user interface shows comments and replies, and has no character limit, which allows for easier back-and-forth interactions, and a more horizontal relationship between influencers/opinion shapers and their followers. The exchange below, extracted from a post about a Lebanese man taking Syrian refugees into his house for shelter after Storm Norma, shows how users argue more informally on Facebook.





Main Facebook comment followed by replies.

English translation

لماذا ينجبون وهم في حالة حرب و تهجير بلا مؤى و لا اكل ولا شرب و فقر!!!! ماهذا التسيب هم يظلمون اولادهم و يظلمون اطفالهم و يقتلوهم بالجوع و البرد لا بد من تحكيم العقل

Why do they have children in midst of war, eviction, without shelter or food or drink and poverty!!!! Why this disruption- they are wronging their children and killing them with hunger and cold- we need arbitration for their minds.

الحرب عمرا سنين اذا بدن يوقفو [user name mentioned] عن الانجاب ما بضل حدا بكرا شو مخبالن ما حدا بيعرف مانك اكرم من الله مثل ما رب العالمين بعنن هل الرجل الطيب بيعنن غيرو اني شو ما كانت طائفك او دينك فقط بمبدأ الانسانية بكفي انك تدعين اذا مارح تساعدي باي شي لو كان معنوي

The war is 9 years old if they want to stop having children no one would be around later, no one knows what tomorrow brings... Regardless of your sect or religion, start with humanity first, it's enough to pray for them, you don't have to help in any other way.

لونك متزوجه ابتعري ليش [user mentioned]

If you were married, you would know why.

Le 3m btechetmo hal2ad bi [user mentioned] ma3a 7a2 l benet che2to am abayto . Oumoume w ma oumoume , 5alloune elkoun eno l oumoume l 7a2i2iye hiye bas l 2em ma te2bal tchouf ebna 3eyich bi zrouf bech3a . Mn 7a2 l walad li byouwlad yet3allam w yetdaffa w yet7akkam w y3ich toufoule , mch hek bhal te3tir w obviously hl wled mch 7aslin 3a chi , fawtit l walad l soure 3l mestachfayet bi lebnen 3m betkoun ktiir so3be . Fa ana law lej2a souriye knt bfaddil ma jib wled ta ma 3azzeboun . W bl nesbe la yalli 3m bi oul l walad byeje w bi jib reze2to ma3o mch mazbout w l dalil 7alit hol l wled . W alla sa7 bi 7ot l rou7 bl walad bas l ahel hene li bi arero yjibou fa law 3anjad badoun masla7it wledoun ma keno bi 5alfo . Bas 3a kl 7al mni7 li fi hl eben l 7alel li edir yesta2beloun bi bayto

Why are you attacking [user mentioned] she is right whether you like it or not. You're telling her to think about motherhood, let me tell you that real motherhood is when a mother doesn't accept to have a child in ugly circumstances. It's a child's right to be born and learn and be warm and get medical care and live their childhood, not in these living conditions. Getting a Syrian child into Lebanese hospitals is really difficult. If I were a Syrian refugee I would prefer not to have children to not torture them...

[user mentioned] الله لا يطعمك ولاد و اذا عندك الله يحرمك منهم

May God never give you children and if you have them may God take them away from you.

[user mentioned] يبقا بس تتشردى والله كريم تزوجى ولا تخلفي مو مشان شي مشان ماتكون الخلفة متلك قليلة فهم وشكرااااا

When you're stranded and get married don't have children just so that your kids aren't as dumb as you are and thank you



Broadcasting on Twitter. Twitter, although still relevant to the conversation, remains a more high-level platform, providing news headlines and key findings, but more rarely serving as a platform for debate or dialogue across community lines. Of the highly relevant tweets, almost half (12,098 of 25,310) were Retweets, demonstrating the nature of Twitter as a space to share news updates, rather than to share original personal commentary. The first two screenshots below show an exchange of replies between two people, as commentary on a decision by the Minister of Foreign Affairs. It demonstrates the regular practice of expressing responses through a retweet from one's own account, rather than a reply on a long comment thread with multiple users on another user's account.

Screenshot



English translation

The main post, which is a commentary on a retweet in itself, reads,

User 1: “And have you tried as foreign minister to solve the issue of those in Syrian prisons for the good of the people?” User 2: “No no no. You have it all wrong, he is the minister of luxury and travel.”

User 1: “You’re right how could I ask this of him.”

User 2: “Be careful, next time they’ll accuse you of misconstruing the position.”

User 1: [laugh emojis]

C) What are the issues that impact host community – refugee relations that are discussed online?

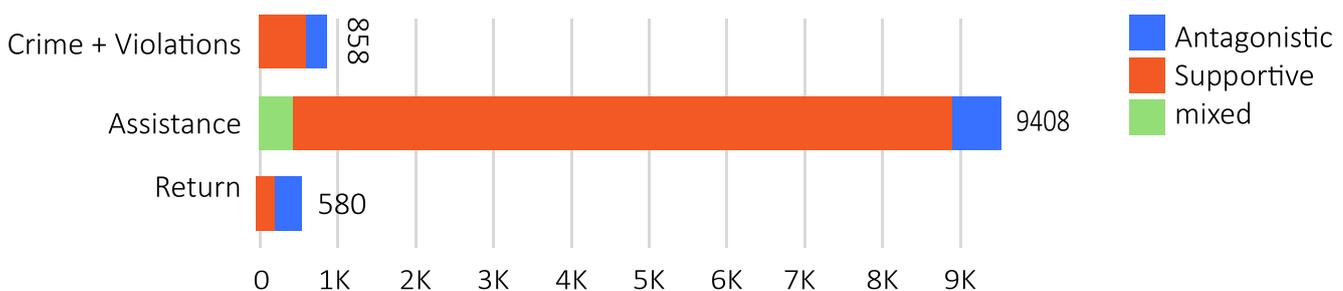
The issues discussed online tend to mirror offline events, with topics trending in the immediate aftermath of the events and often lasting only a short time. However, within discussions around events, and outside a timeline of relevant news stories or ad hoc reactions, **an ongoing general discourse about Syrian refugees can be readily identified.** These can be seen in both the framing of the reporting of particular events, as well as the comments and reactions to posts.

For the purposes of our research, we identified specific keywords that denoted the **topics of return, assistance and crime and violations**, as these emerged as key topics dominating the general discourse, and were also reflected in the response to major events occurring in January. In addition to these keywords, we categorised all content relating to Storm Norma as ‘assistance’ and all content relating to Ahmad Zoabi, Tannourine and Aarsal incidents as ‘crime and violations’. The inclusion of Storm Norma in this dataset accounts for the large number of Tweets relating to assistance, making it easier to categorise posts as directly relevant to assistance.

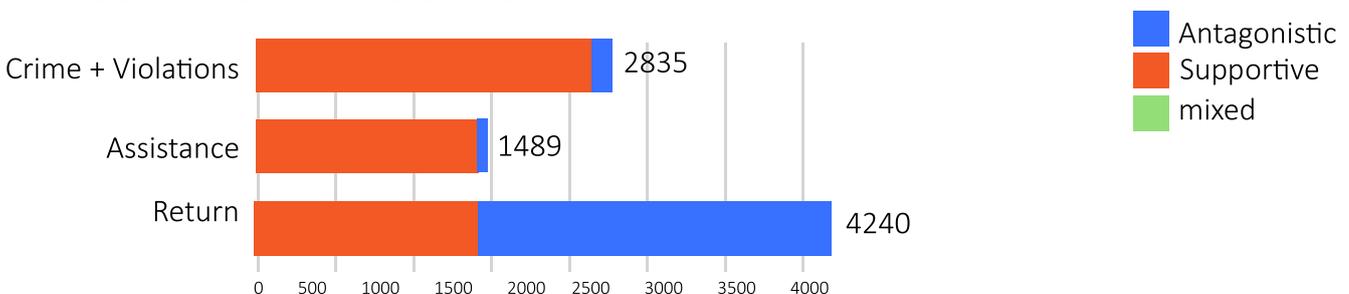
The below graphs demonstrate the **spread of sentiment within each topic**. Assistance provoked a majority of supportive sentiment, seen in the show of solidarity and calls for mobilisation following Storm Norma. Return is a more controversial issue, provoking a more even spread of sentiment on both platforms. This corroborates qualitative analysis that shows return to be a topic of conversation between multiple perspectives. Crime and violations provoked significant supportive sentiment, despite being widely considered as a potential trigger of tensions. Whilst a crime committed by a Syrian perpetrator (eg Tannourine) does indeed provoke antagonistic reactions (see section D), crimes committed against Syrians by Lebanese actors (or state actors in the case of Ahmad Zoabi) triggered predominantly supportive sentiments in the form of solidarity, sympathy and calls for action by those supportive of the Syrian refugee presence in Lebanon.

There is a clear difference between levels of antagonism relating to returns between the two platforms. Facebook saw a much greater level of antagonistic sentiment, in part reflecting several polls on returns which provoked a large number of reactions calling for Syrians to return to Syria at the earliest opportunity. This also reflects the more informal, emotional types of conversation typical of Facebook, versus the more factual sharing on Twitter. It also reflects the greater prevalence of untagged posts on Twitter (ie neither supportive nor antagonistic posts) against the greater proportion of tagged posts on Facebook.

Volume & sentiment of Issues for twitter



Volume & sentiment of Issues for Facebook



Crime and violations:

As noted above, the response provoked by crime and violations posts depends significantly on the type of crime, as well as the direction of crime between communities.

a) Crimes committed by a Syrian individual

Here, three trends can be observed. First, posts tend to emphasize the nationality of the culprit. Second, provocative imagery is used in posts about the crime. Third, responses to the crime tend to bring out underlying antagonistic sentiment where the crime is seen as a symptom of the larger Syrian refugee burden, rather than as an isolated incident.

For example, Tweets about the Tannourine incident regularly included the term Syrian in their reporting of the incident. This contrasts with other crimes reported through simple captions such as 'man arrested for x crime'. Whilst all individuals framing this event in this way may not hold antagonistic sentiments, the online context appears to encourage this distinction based on nationality. As a result, the narrative of the two communities being treated differently prevails. This portrayal of foreign nationalities as 'the other' is also played out towards other communities, including the Palestinian refugees, although this was not a focus of this study.



Female citizen stabbed with a knife in Tannourine, Saad: The Syrian presence is a ticking time bomb Syrian attacks female citizen and stabs her with a knife...Here's what happened to "Siham" in #Tannourine #Lebanon

In addition to the framing of these crimes, those committed by Syrian individuals seemed to involve more violent imagery and expressions, as shown in the below image from a major news outlet.

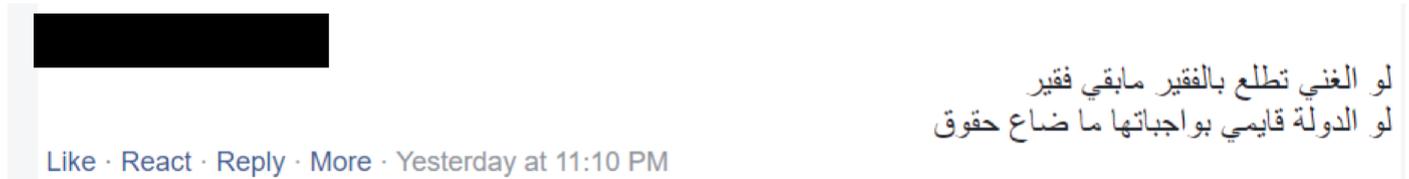


The headline here reads 'He stabbed her in the back to rob her in Tannourine'. Below, the commentary is evident of the extrapolation from the crime to the broader refugee issue (comment 1), but also of the violence used in the discussion of such events (comment 3).

- Comment 1:** Don't you dare take him back to his country, because UN, EU, and humanitarian organizations will attack us. We should ask him nicely if he wants to return voluntarily or run free here and butcher our "security men" (ISF) with the excuse of robbery. If he doesn't face justice here he should be taken back to his country to get tried.
- Comment 2:** How come he has prior convictions and you're leaving him running free in the country, at least take him back to his country when he's out of prison
- Comment 3:** Execution for this dog!

b) Crimes or violations committed against Syrian individuals.

These cases resulted in supportive sentiment, for example in response to the Ahmad Zoabi incident. Whilst supportive in the broad sense, these comments referred less directly to wanting to help Syrian refugees, and were more about expressing solidarity and condemning the actions of the Beirut municipality or government more broadly. See example below.

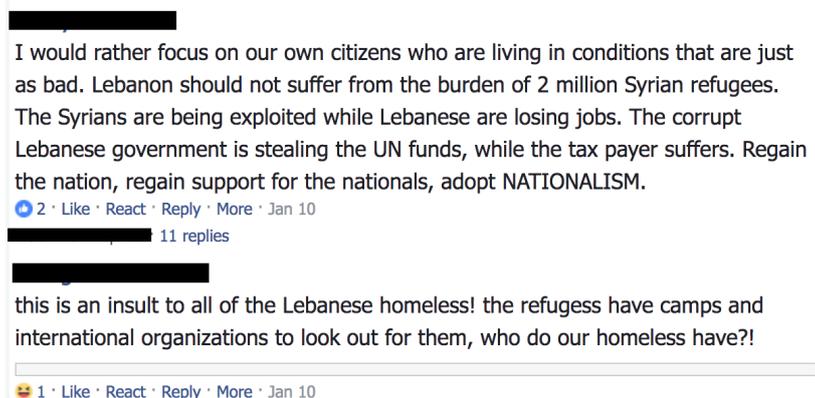


If the rich would pay attention to the poor, there would be no more poverty.
If the government was doing its job from the beginning, there would be no lost rights.

Assistance

Assistance allows more space for positive sentiment than crime or return, as links to natural disasters and poor living conditions for Syrian refugees provoked sympathy and calls for solidarity among Lebanese. Across both platforms, widespread mobilisation and calls for action were evident. At the same time, some backlash was created from this mobilisation, seen through the narrative that Syrian refugees are disproportionately benefiting from international aid.

For example, this comment was found on a post calling for funding to support the storm response:



Even where responses to assistance are overwhelmingly supportive, they tend to show support for Lebanese institutions or individuals that are supporting the refugee community, and often verge on pity, or on highlighting Lebanese superiority. 'For example, the below post announcing the transportation of Syrian refugees to safe areas away from the storm.



The post describing the Lebanese man who took several Syrian children into his home during storm Norma is also evidence of this trend. The comments were considered overwhelmingly supportive, but the narrative was one of praise for the Lebanese man, rather than support for the Syrian refugee presence. For example,

يعني عنجد زلمي بكل المعنى كل الحكي ما رح يوصف اديش انت انسان محترم لازم يحطوك انت رئيس بلبنان بدل هودي لي عنا كل الاحترام والتقدير لصاحب هل عمل الخير انت لي بتفكر صح انسانيتك هيدي بتكفي لنقتخر فيها يا ريت لو الكل يصير هيك متلك والله يكون بعون الاجئين حرام مش ذنبون هل شي الله يطول بعمرك مشان ضل هيك فاتح ابواب المساعدةي عنجد كل الاحترام الك

'Really, a man in all sense of the word. No words to describe how respectable you are. You should be the president of Lebanon instead of the ones we have now. All my respect and appreciation for this do-gooder..This is enough to make us proud. I wish everyone could be like you and may God be with you and with refugees it's not their fault. May God give you a long life so you could continue to open doors for those in need. All my respect.'

Returns

The topic of returns was widely seen in response to events or incidents involving Syrian refugees, demonstrating its importance as an underlying narrative. Closely linked with the topic of assistance, calls for return appeared in the aftermath of Storm Norma, as well as in the aftermath of Tannourine and Aarsal incidents.

For example, this comment appeared on a post in response to Storm Norma:



Comment: "Send them to Syria it's warmer and better than here. But they like to stay in Lebanon to earn money from organizations and beg. [They are] a people who like to shame themselves."

Returns were not only discussed in response to events. Both on Twitter and on Facebook, polls were created to ask users whether they supported the return of refugees, and received a large number of engagements demonstrating the importance of this topic to users.

On Facebook, a post asked 'Are you with the return of the displaced Syrians today, or after the solution in Syria?', with a picture of the green buses that have become synonymous with the return of Syrian refugees. The post garnered 3,800 reactions and 1,300 comments with a majority of users commenting simply 'today' or 'definitely today'. Many comments received numerous replies, demonstrating the ability of this topic to provoke engagement. On Twitter, the below post by the same group asked 'are you with Syrians' safe and voluntary return to Syria'? The tweet was retweeted 31 times and received 87 replies. This is an unusually high number of engagements for a Twitter post in our dataset, reinforcing the nature of this topic as a trigger for high levels of engagement.



D) Are these issues discussed to an extent whereby meaningful sentiment analysis can be conducted to identify trends and spikes?

In order to analyse sentiment, the parameters of a particular sentiment need to be defined. The terms 'positive' and 'negative' are broad, and could be interpreted in multiple ways. For the purposes of this study we define positive and negative as 'supportive' (ie. denoting solidarity or feelings of support towards Syrian refugees or Lebanese as a host community)

and 'antagonistic' (ie. denoting anger, hatred, dislike or distaste for Syrian refugees or Lebanese as a host community). Working from this basis, we developed a set of terms that would denote support or antagonism for the general dataset. A categorisation of the two sentiments is provided in the table below.

Sentiment	Key findings (qualitative)
Supportive	<p>Key narratives:</p> <ul style="list-style-type: none"> Towards Syrians: help innocent refugee population, dignity, rights, victims of a political conflict out of their control, support for children. Towards Lebanese: solidarity, considered welcoming hosts, Largely sparked by natural disaster eg Storm Norma, as well as by events affecting Syrian children eg Ahmad Zoabi.
Antagonistic	<p>Key narratives:</p> <ul style="list-style-type: none"> Towards Syrians: return to Syria, burden on Lebanon, danger to Lebanese society, squeezing resources, benefitting from assistance, Syrian occupation, women as reproducers of burden on Lebanon, nationalism Towards Lebanese: unwelcoming hosts, racist hosts, lack of understanding for Syrian role in Lebanon, restrictive conditions. Sparked by incidents of crime or violence, or by events relating to underlying narratives eg return.

Although these two categories can refer to support and antagonism from both Syrian and Lebanese communities, **the majority of data gathered for this report referred to attitudes towards the Syrian community.** Whilst there are spaces where Syrians express supportive or antagonistic sentiment towards the Lebanese, these appear fewer in number in this research and are perhaps less accessible to public viewers. This can be explained by several factors, including sensitivities and security concerns of the Syrian community online, as well as the context of the research as Lebanon-centric and focused on the online public sphere. Future research may wish to further explore the private spaces where conversations expressed from a Syrian perspective are aired, as well as seek to better understand the Syrian influencer space, in order to ensure that any intervention work addresses both communities equally.

Sentiment is a fluid concept and is denoted in subtle and constantly changing ways, making quantitative tracking of these terms challenging. In addition, sentiment is not expressed in a vacuum—it is embedded in responses to specific events and incidents. As a result, words and phrases that can signal a certain sentiment at one moment may later signal a different sentiment in response to a specific event. To mitigate for these changes, we developed event-specific sentiment indicators. Due to this fluidity, sentiment analysis requires a combination of both qualitative and quantitative analysis.

Of particular relevance to these categories is the nature of the term ‘supportive’. A broadly positive category, reflecting support for Syrian refugees from the Lebanese host community, it is more nuanced than a straightforward sentiment analysis can allow. Comments considered ‘supportive’ range from expressions of sympathy (genuine solidarity) to pity (some level of positive sentiment), through to superiority (expressing support in a way that reinforces the existing power dynamic between Syrian and Lebanese communities). At the far end of this spectrum, the link between ‘support’ and ‘nationalism’ can be blurred. The examples in Part C demonstrate these blurred lines in response to Storm Norma, through praise for the Lebanese man who hosted Syrian children and for the Lebanese Army, both

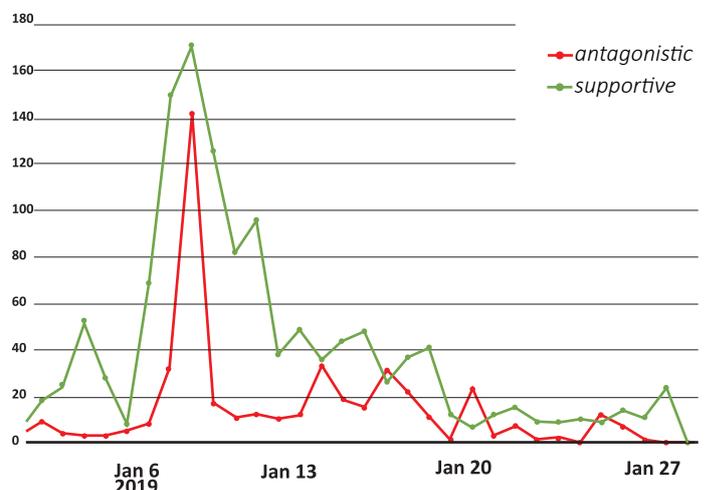
categorised as heroes. As a result of these nuances, we recommend breaking sentiment categories down into sub-categories in future research (section 2). It could also be worth exploring the offline social norms behind these sentiments, such as the perceptions of treason when Lebanese voices openly support Syrian refugees in a climate of general hostility (eg on a page where hostile sentiment prevails), and the impact this has on people’s willingness to express supportive sentiment online.

With these challenges in mind, it is possible to analyse overall levels of sentiment over time and in response to specific events or issues.

Sentiment over time:

Sentiment appears to spike in response to 1) overall political discourse (eg a Tweet or public statement by a major political figure or event); 2) macro-level events that affect the refugee community (eg Storm Norma); and 3) smaller scale incidents involving individuals (eg a small scale crime). Often, an event will trigger responses not only directly related to the event, but relating to underlying sentiment and broader narratives around the refugee presence (such as the economic crisis or the burden faced by Lebanon). Spikes in sentiment are exacerbated by relevant events, but they are not dependent on them.

The graph below shows number of tweets denoting sentiment on Twitter over time, incorporating all event and issue categorisations.



Sentiments by 24-h interval for twitter

A clear spike in supportive sentiment is evident following Storm Norma, as Tweets circulated expressing solidarity and calling for funding. In the same time period, a significant spike in antagonistic sentiment is observed. This spike was influenced by two separate Tweets by an important political figure that were widely retweeted (266 times) - both of them referring not directly to the Storm, but to the burden refugees place on Lebanon and its systems. The Tweet, which expresses fear of the continued refugee presence in Lebanon and the demand it places on the economic and security situation in Lebanon, demonstrates this link between events and underlying sentiment. The tweet, released as Twitter traffic about Storm Norma was rising rapidly, refers directly to the issue of returns.

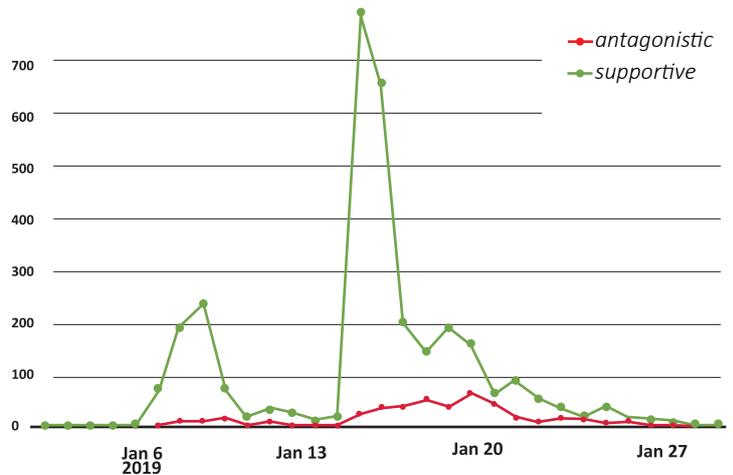
The fact that narratives are underlying, provoked by individual tweets, political events or incidents, shows that this analysis is not proportionally reflective of changes in overall sentiment. **These spikes do not represent the absolute level of sentiment among Twitter users. Rather, they are triggered by the awakening of individual users' pre-existing sentiments.** Spikes instead demonstrate how existing sentiment is magnified in the context of certain events, such as Storm Norma or the death of Ahmad Zoabi. The scale of the event (or the person leading the conversation) determines the scale of the spike in sentiment, not the actual level of sentiment. Whilst Norma - an event that reached an international audience - led to a large spike, Tweets about the overspill of sewage from a Syrian IS in the South into the Litani river irrigation channels led to a much smaller but still visible spike on January 22nd.

Regardless of the scale of the event, spikes are very rapid, reflecting the speed with which content is retweeted and then discarded on Twitter. For example, a Tweet that contributed to the spike in antagonistic sentiment on the 15th, referring to the issue of return, was not retweeted again after the 15th.

Whilst these trends also stand on Facebook, our methodology does not allow for a representative mapping of spikes. Due to the smaller quantity of data gathered from Facebook, single posts that have large numbers of reactions have a major effect on the trend line. For example, a single post on AJ+ Arabic reporting on a Lebanese man who took Syrian

children into his home during storm Norma garnered 2,800 comments, the majority of which were coded supportive. This accounts for the large spike on January 16th. As a result, we do not believe that the graph below is useful for analytic purposes, but include it to illustrate what would be possible if we had access to more data via an API.

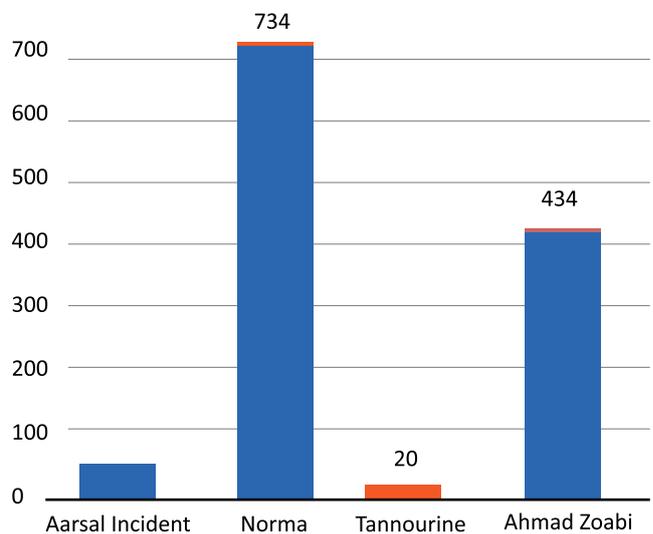
The graph below shows number of tweets denoting sentiment on Facebook over time, incorporating all event and issue categorisations.



Sentiments by 24-h interval for facebook

Sentiment in response to major events:

On Twitter, sentiment tends to be one-sided when considering only Tweets that refer directly to an event. Responses to Storm Norma or Ahmad Zoabi were predominantly supportive, whilst those in response to the Tannourine incident were entirely antagonistic.



Volume & sentiment of Issues for twitter

However, this reflects only part of the broader conversation, seen in the overall sentiment graph above. It shows that whilst Tweets referring directly to an event such as Norma are overwhelmingly positive, it is only when zooming out to the non-event specific conversation that we get a full picture. The spike in antagonistic sentiment shown in the line graph above demonstrates that negative sentiments expressed in a time period in which conversations about Norma could not be ignored did not directly refer to the storm. This finding is key - **it is not enough to understand only conversations responding directly to events. To gain a full picture, it is necessary to see the full range of conversations happening in a given time period**, including broader political events and underlying narratives, particularly of the major political influencers within Lebanon.

Furthermore, a large proportion of Tweets about individual events are not coded as either supportive or antagonistic sentiment. This reflects the use of Twitter to publicise more than as a platform to display sentiment. Of 737 tweets about the Aarsal incident, for example, only 154 were coded supportive, and none antagonistic.

Qualitative analysis of events:

To dig deeper into sentiment in response to events it is necessary to run a qualitative analysis of conversations, particularly on Facebook, where comments on posts are often more subtle and therefore harder to code according to predefined sentiment categories.

For example, whilst Storm Norma provoked an overwhelmingly positive response on both Facebook and Twitter, subtle evidence of antagonistic sentiment was found even among largely supportive conversations.

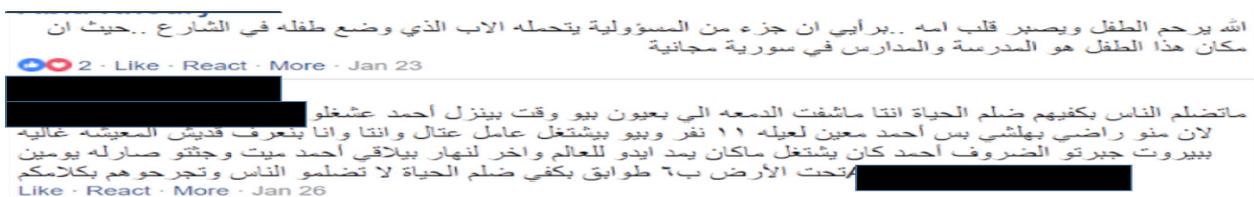


Translation

Main post: “If we, in houses, are flooded with water, what about people in tents? May God have mercy on those living in camps.”

All comments say variations of, “Amen”, and/or “May God help them” (widely considered expressions of sympathy), except one which says “put a bucket out”, sarcastically addressing the rhetorical question about flooding.

In the case of Ahmad Zoabi, an event which also provoked a seemingly uniform positive response, negative responses could be found encouraging return. Below, the notion expressed by the commenter that the father is to blame and that Syrians get free schooling can be considered antagonistic although at face value the expression ‘May God rest this child’s soul and calm his mother’s heart’ could be considered supportive. Such edge cases make event-specific sentiment analysis challenging on Facebook.



Reply 1: “May God rest this child’s soul and calm his mother’s heart. I think part of the responsibility is on the father who put him on the street. This child belongs in school and school is free for Syrians.”

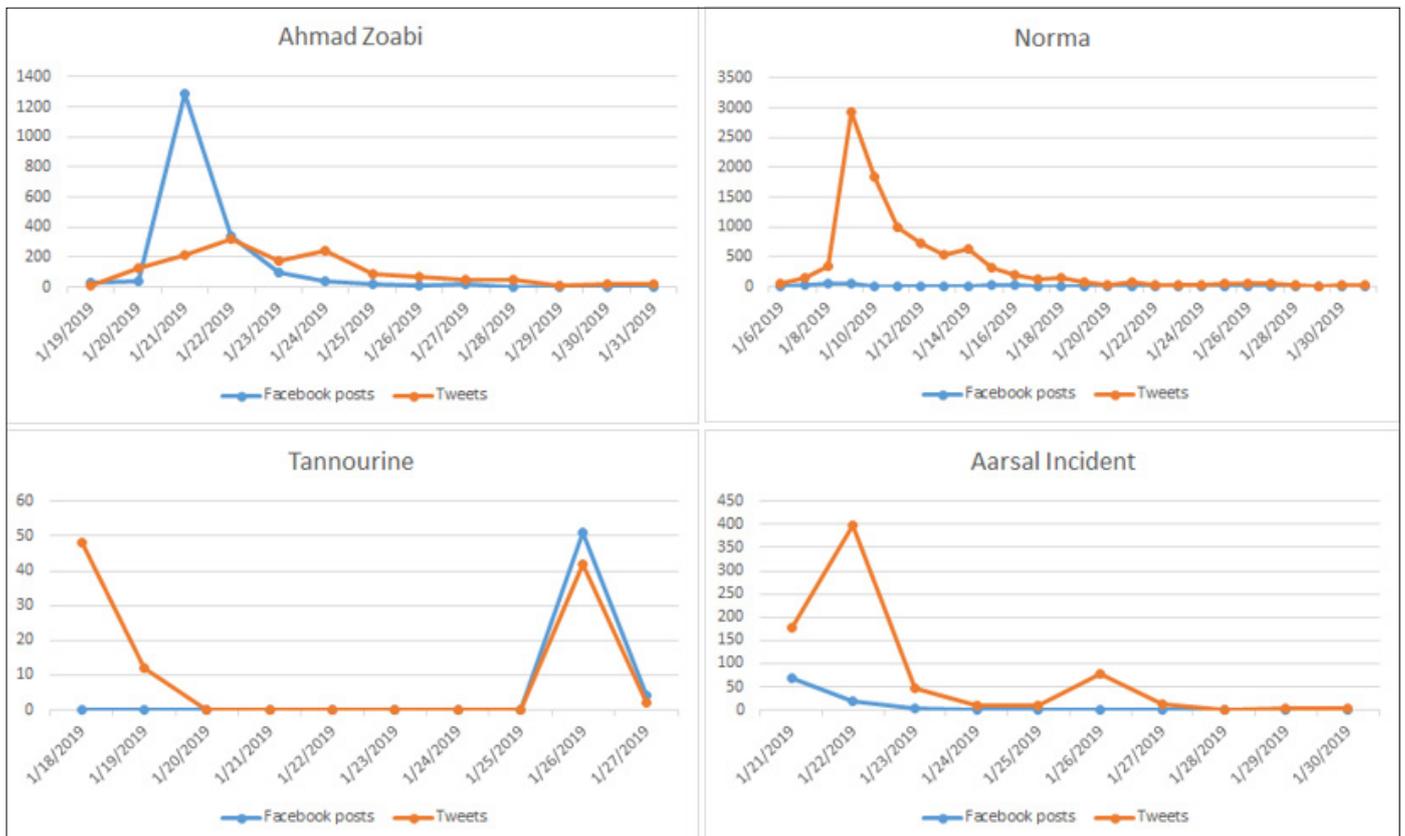
E) Are there instances when these spikes coincide or lead to violent incidents, or shape public discourse after violent incidents?

Relevant conversations on social media do spike following incidents, providing an ‘instantaneous and immediate’ response to offline events.¹⁸ After every major incident that occurred in January, a notable uptick in the volume of data was recorded (see section D).

However, it is very difficult to draw predictive conclusions from such data. As some interviewees pointed out, this analysis can help provide ‘an accurate temperature gauge’ but would require a lot of additional data (which UNDP may have) to compliment the social media data¹⁹. We found no evidence of spikes occurring before events during our data collection period, supporting the view that social media tends to be a reactive space, very responsive to events but not always predictive of them.

There are some exceptions to this (eg the organisation of protests) where spikes may be seen on social media prior to an event²⁰. However, no such events occurred during the period of data collection. In order to better understand the frequency or nature of these rare predictive cases, data would need to be considered over a much longer term. In addition, by the time spikes occurred on social media, even if they were before an event, an offline, traditional monitoring system would likely have already picked up heightened social tensions. Social media can be understood as a mirror of offline sentiment. Even events like protests do not occur in a vacuum, but are a reaction to growing disenchantment and sentiments likely to be observed offline in the run up to the event.

Figure: number of Facebook and Twitter posts related to key events over time



18 Interview with UNHCR Lebanon, Communications with Communities Unit

19 ‘Interview with Nasser Yassin, Director of Research, Issam Fares Institute for Public Policy and International Affairs, American University of Beirut’

20 Interview with expert, Beirut

Whilst we can analyse the online discourse that emerges in the immediate aftermath of events, this is of limited analytical value. As shown in the graphs above, spikes following events are short-lived, with conversations fast moving on to other issues or incidents. This makes any in-depth analysis of discourse in specific relation to an event challenging as reactions provide only a snapshot of the broader conversation. As explained above, to analyse sentiment of Tweets relating solely to Norma (as supportive) would have ignored the underlying antagonistic sentiment that peaked in the same period. Public discourse is shaped outside of events, making it difficult to distinguish discourse that is related to events from more general discourse.

Where spikes do occur they often reflect a combination of an offline event and underlying narratives that shape conversations. As a result, social media analysis serves to deepen our understanding of sentiment over time, helping us to understand the overall discourse on a topic rather than on individual events. It is only through social media analysis over the long term, coupled with offline analysis of perceptions and tensions, that a full picture of discourse around a particular event can emerge.

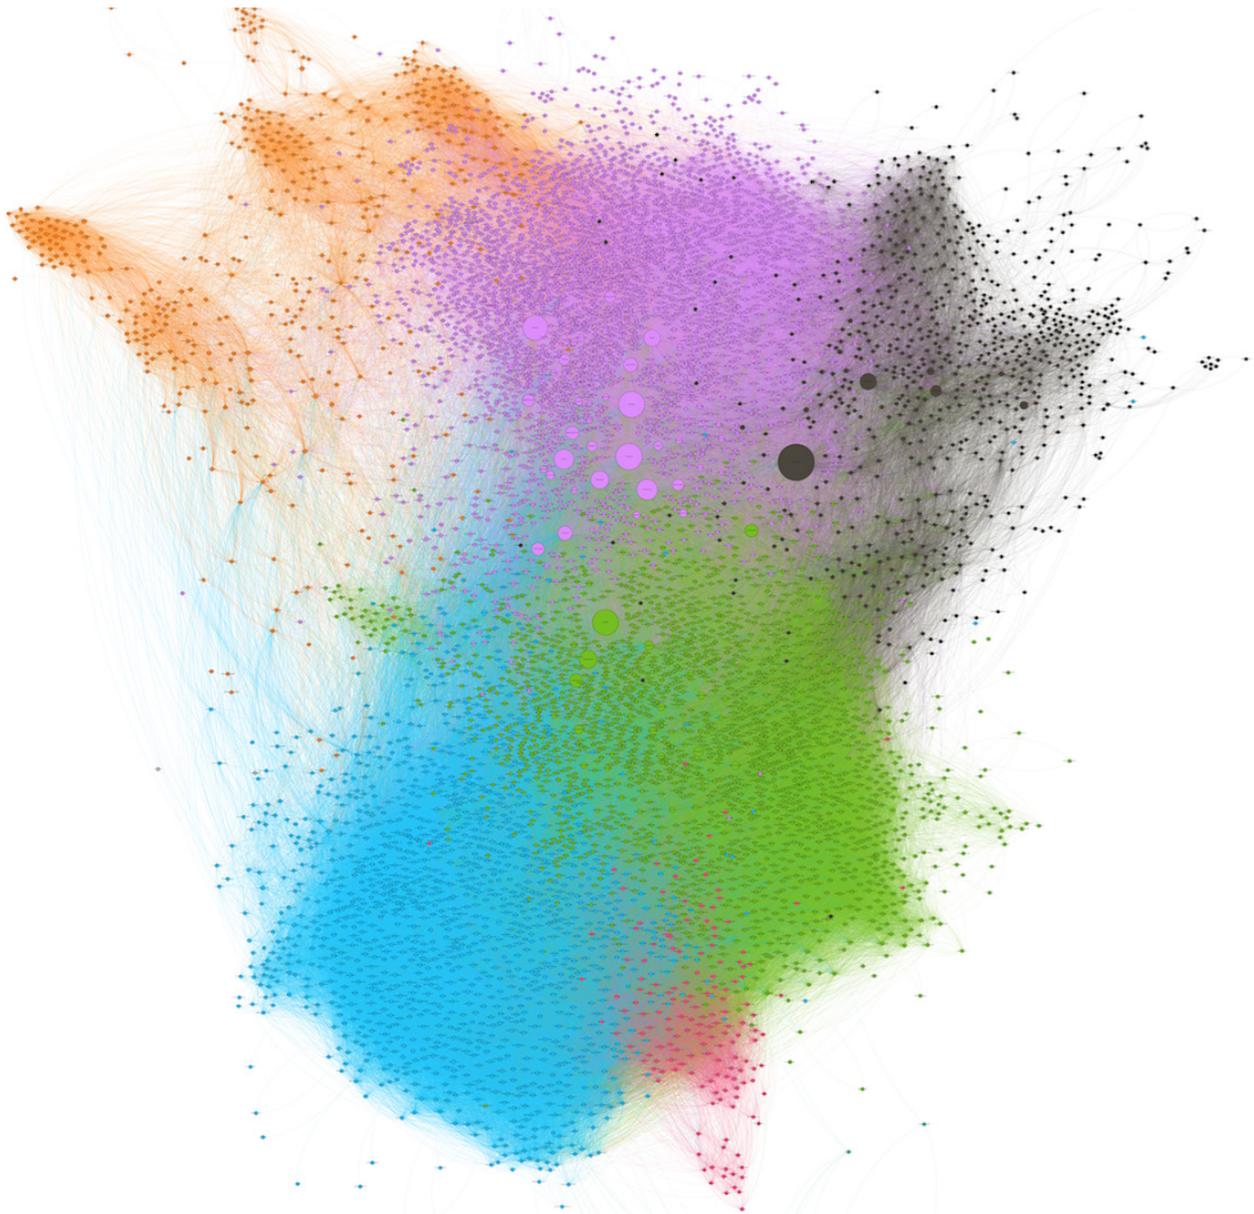
Even where discourse can be analysed in response to an event, it is difficult to disentangle whether online discourse is shaping offline discourse or vice-versa. Does expressing supportive sentiment online translate into supportive discourse offline and from that to supportive action? Does expressing antagonistic sentiment online result in a greater propensity to engage in antagonistic actions? More research is needed to fully understand how social media conversations shape real felt perceptions in this context.

F) Who are the people / organizations discussing these issues online and to what extent are they 'influencers' of others?

This research question focuses on the people / organisations discussing these issues on Twitter. Whilst technically possible to view a network map of Facebook users, our sample of Facebook posts was a snapshot and therefore not representative of the full conversation and its users. We did not feel the network map of Facebook users would provide reliable insights. As a result, the below analysis focuses entirely on Twitter.

The following network graph represents the communities active within this discussion. It displays all Twitter users that tweeted about one of the key events, as well as the seed users, with nodes representing a Twitter user, the size of the node representing how many followers it has, lines connecting nodes (Twitter users) if one user follows another, and colours representing distinct communities (clusters of accounts with many follows between them, technically referred to as a 'modularity class').





Five distinct communities emerged from these discussions on Twitter. Analysing their profiles and the extent to which they discuss events reveals the following characteristics of each community:

Colour / community	User description
Green- Lebanon-centric influencers	These users are a combination of Lebanese news sites and public figures, and appear to have the most influence on Lebanese opinions online. Public figures here are either politicians like Saad Hariri, Walid Joumblatt, and Gebran Bassil, popular journalists with regular shows/platforms for engagement with the public like Dima Sadek and Marcel Ghanem, or celebrities such as Elissa and Ragheb Alama. Within this community we also find local news sites and institutions, such as Al Jadeed TV, LBCI, and Internal Security Forces. The majority of content from these users is in Arabic.

Purple- Arabic speaking influencers	Users represent news sites that are Arabic-speaking, but not based in Lebanon, and their followers. Most of the sites are based in Syria, Egypt, and the Gulf. They have significant influence among Lebanese users, but also overlap with clusters of non-Lebanese Twitter users that have discussed Syrian refugee-related topics online. Although these users do not target Lebanon as a specific audience, they still shape opinions and discourse around current events, particularly for Arabic speakers based in Lebanon. Examples include Al Arabiya, BBC Arabic, and Al Watan News 24.
Grey- English speaking influencers	This category brings together English-speaking, international news sources (e.g. Reuters and Al Jazeera English), INGOs (e.g. World Vision and NRC), and UN agencies (e.g. UNHCR, UNICEF, and UNDP). These institutions are closely tied to individuals that have relatively progressive political opinions and/or are linked to an international institution like the American University of Beirut.
Blue- pro Syrian regime users	Here both individuals and institutions are apparently pro-Syrian regime. These users are mostly individuals and smaller-scale news sites or institutions that have the same political affiliation. They are predominantly Arabic speaking.
Orange- outsider groups	This category groups clusters of allied users located outside of Lebanon. Three distinctions can be made here. One (top-left) is suspended accounts and spam that could have been temporary accounts created to influence opinion around certain topics. The next cluster (centre) is pro-Palestine, and appears to express solidarity towards refugee populations in Lebanon. The third cluster is location-specific to Egypt. These final users are on a tangent and engaged in the conversation solely on the Aarsal incident.
Red- pro-FPM users	More likely to consume international and English-speaking media than those in the blue cluster. Many appear as supporters of the Free Patriotic Movement in Lebanon. They are predominantly Arabic speaking.

These five communities are to a large extent distinct - each operates within their own sphere of influence. Conversations had by those in one sphere are not widely seen by those in other spheres. This is particularly relevant for those on the outskirts of the diagram, or in isolated clusters (eg. the Egypt cluster). Those in the centre of the diagram or on the border between two colours trend more towards interaction with other communities. **Notably, the English-speaking community, which includes the majority of UN agencies and INGOs, does not exert much influence over the domestic Lebanese audience represented by the green cluster.** As a result, attempts to influence the narrative on these issues should come from within the Lebanon-centric cluster wherever possible.

The grey, purple and green clusters are made up of most of the 'influencers', i.e. those with most followers. These can be loosely distinguished by their audience (international, Arabic speaking or Lebanon-centric). They are a mixture of political, media and entertainment personalities and organisations. These influencers tend to be closer to the centre of the diagram, demonstrating their wider reach across communities. The blue, red and orange clusters are made up of individuals with less relative influence, and can be broadly defined by their political affiliation. Whilst less influential, individuals within these communities are relevant to engage with as they may be more prone to express personal sentiment towards refugees.

The map is not easily broken down between pro and anti-refugee sentiment. Polarization between communities does not appear to be happening based on attitudes towards refugees - within the Lebanon-centric cluster alone there are both pro and anti-refugee sentiments being discussed. Those that have the most influence are not easily categorised based on their Twitter community. Instead, **polarization between communities appears to be occurring on the basis of political and confessional leaning** - people within the blue, orange or red clusters show greater tendency to follow those with the same political affiliation as themselves. This makes it more difficult to assign sentiment towards refugees across the Twitter-sphere.

Communities respond differently to different events. Whilst the majority of those on the map tweeted about Storm Norma (79%), the Egypt cluster was notably absent from that conversation - in fact they only engaged in the Aarsal incident discussions demonstrating a single interest. The Ahmad Zoabi incident was predominantly discussed by the Arabic-speaking media category and to a lesser extent the Lebanese influencers, demonstrating the regional interest in the case. The Aarsal incident was discussed by 11% of users from across the spectrum. Across events, communities could be identified that retweeted a particular narrative. For example, the Egypt cluster within the outsider groups community almost uniformly put out a tweet absolving the Future Movement of any involvement in the incident. The Tannourine incident was discussed by a smattering of individuals (1%), largely in the Lebanon-centric cluster. It was not picked up by the most influential users within that category, explaining the limited reach of the event discussions.

The below graphs demonstrate the distinct responses from these communities to events. Colours represent users tweeting about an event, grey represents users who did not tweet about an event.



G) To what extent is it possible to geo-locate where more hostile sentiments are being expressed?

Through manual Facebook analysis, it is not possible to provide any quantitative location analysis, due to the need to anonymise the dataset entirely. Although qualitative analysis does reveal some information about the location of users and commenters, it requires individually visiting user profiles and is therefore challenging to do on a large scale. Furthermore, there is no straightforward quantitative measure of location on Facebook. We did not download any data from the 'check-in' function for this study, nor did we fully explore the possibility of using Facebook ads targeted to specific locations which, when coupled with filters for specific interests, could lead to greater understanding of location and size of audiences by showing us the size of the ad audience that Facebook has determined are interested in specific topics within specific geography. This could help roughly geo-locate specific sentiments.

On Twitter, the ability to analyse by location is restricted to self-reported location settings by users. The geo-location tagging is of very limited [or no] use as it is automatically disabled and thus requires an opt-in function. If an account lists its location, that location will be included in the dataset, if not the user location field will be left blank. As a result, it is likely that some users do not report their locations, resulting in location filters under-estimating real figures. Of 2.6 million Tweets, 56,900 listed their location as somewhere in Lebanon (many listed only cities and not the country, so this is a combined figure for all major cities in Lebanon).

Even where users do not report their location through Twitter, it is often possible to locate a user based on the content of their posts, including their keywords and range of content. For this reason, we considered users tweeting about our final keywords as relevant even when they did not explicitly state their location as being in Lebanon.

It is also possible to analyse the location about which users are talking, even without knowing where they

are, for example through keyword searches for 'Aarsal'. This could be a more useful analysis for UNDP than geo-location of users, particularly given that influence on social media goes beyond geographical boundaries. Users follow other users, not only on the basis of geographic location, but on the basis of content and mutual interest. As a result, influence on perceptions of those in Lebanon is not restricted to users located inside the country, limiting the value of analysis that focuses on location.

H) To what extent are hostile sentiments being expressed by users outside of groups and pages on Facebook?

Current terms of service for Facebook limit quantitative analysis through their API to public pages. As a result, we did not run any quantitative analysis on private posts or on closed groups. Qualitative analysis suggests that although this is a limitation, a lot of the debate and inter-communal interaction is seen on public pages, particularly on large news pages.

Individuals do appear to share information on their personal pages more quickly than public pages- both the Aarsal incident and the Ahmad Zoabi case first surfaced on personal pages. However, their posts were made public, enabling for a larger audience and suggesting a deliberate attempt to reach an audience beyond private followings. In fact, these individuals with a broad influence often have private pages rather than public pages- they appear as individuals but in fact have a large number of followers that are not necessarily friends.

I) Does the current regulatory framework (both national legislation and social media platform rules and regulations) enable the identification of and engagement with polarized users?

The apparent lack of regulatory framework in Lebanon makes this uncharted territory at the national level. Few organisations appear to have explored the identification of and engagement with polarized issues in this context. However, concern has recently been expressed by activists and INGOs over the ISF's 'cybercrime and intellectual property rights

bureau' formed in 2006 and apparently responsible for recent crackdowns on activists on social media.²¹ Between June and August 2018, several individuals were summoned for questioning in response to content criticising political figures on social media. These included an individual detained for changing his WhatsApp profile picture to a picture critical of President Aoun and an individual questioned for criticising Foreign Minister Gebran Bassil on Facebook. It is notable that none of those detained or questioned were accused of spreading hate speech towards Syrian refugees, although we have seen this to be a widespread issue. Whilst these incidents do not directly impact the space to identify and engage with polarized users, they demonstrate the blurred lines of the local legal framework and the prevailing atmosphere of censorship, particularly in response to political content. More should be done to understand this legal framework.

Regardless of local regulations, shifting regulations imposed by social media platforms' terms of service appear to be making it more challenging, although not impossible, to identify and engage with polarized users. A full review of platforms' regulatory framework is provided in Annex II, outlining our ethical guidelines.

J) What are the challenges with APIs for social media monitoring and user engagement and how can they be overcome?

As outlined in the full methodology document, gaining access to APIs for Facebook and Twitter is a significant challenge and a potential impediment to further research. For this research project, no access was granted despite multiple requests and the mobilisation of contacts within both organisations. This followed a recent shift in 2018 that restricted access to APIs by private organisations and individuals.

21 <https://civilsociety-centre.org/article/conflict-analysis-bulletin-issue-12-october-2018-focus-crackdown-social-media-lebanese>

Part 2

Opportunities and challenges (intervening)

A) Based on the findings for questions I-VII, what kind of analysis and tools would be useful in monitoring and analysing social media?

Our research indicates that **the kind of quantitative analysis we produced in this scoping project would be a useful complement to traditional monitoring of social tensions**. Several organisations in Lebanon are conducting social media analysis, but the majority of those are qualitative processes. Although some have attempted to develop quantitative analyses, these trials have been largely unsuccessful, reflecting the challenges of the task laid out in the methodology limitations section above. However, we believe the quantitative analysis we arrived at overcomes many of the limitations of previous similar attempts. The method could be further enhanced by gaining access to the Facebook API.

An extended and more sophisticated version of the tools and processes we used could be developed to provide relevant results on a rolling basis. Now that the infrastructure for this kind of analysis is set up (ie the methodology, data collection tools and graph formats), making adjustments to the inputs (such as keywords and seed users) is relatively straightforward. Once refined, the process could support the production of regular graphics to display overall volume of traffic, sentiment and networks on social media. Specific elements of the tools and process to be refined would include:

- Creating a further **refined list of keyword** combinations, including more accurate terms to denote specific issues and events. This would include a full list of Arabizi terms as well as alternative Arabic and English spellings to form a Lebanon-specific glossary of relevant terms.
- A **refined sentiment analysis** function. This could include:
 - ◇ Further qualitative research to explore how phrases are used in different contexts, as well as several rounds of testing different sets of words and phrases in the tool to explore how they translate into datasets. It is likely that sentiment indicators would also shift over time and in response to specific events, so this would require an ongoing manual review.
 - ◇ A division of sentiment categories into further subcategories (beyond simply supportive and antagonistic). For example, it would be interesting to distinguish sentiments of pity, allyship and solidarity within the 'supportive' category, or nationalism, fear and misogyny within the 'antagonistic' category. This would overcome the generalisations inherent in the supportive and antagonistic distinction, and could serve as the basis for more specific recommendations for intervention (see recommendations in Part B).
 - ◇ Exploring the inclusion of emojis as keywords, as these provide valuable insights into sentiment and cannot be captured through text searches.
- Securing API access. Given the significant workload involved in manual data collection, the expansion of this method would only be feasible for Facebook with API access (that would allow for the process to mirror that used for Twitter). Given Facebook's apparent push to engage in these conversations, it may be worth UNDP engaging specifically with senior management at the company on this issue. In addition, specific API access for this research (not through personal API keys) should be further sought from Twitter.

Short term analysis of narratives on social media provides only a snapshot of what is a rich ground for understanding refugee-host community tensions. In order to ensure sustainability, **such a tool would need to be integrated into the UNDP workflow**, with staff managing the back-end of the database, making relevant adjustments and reviewing analysis outputs. Given the nature of social media, and particularly of sentiment analysis, no useful tool for analysis will be fully automated, and therefore would need regular updates. Despite these capacity challenges, such a process could provide UNDP with a valuable complement to its current social stability monitoring process and would further set it apart from other tensions analysis in Lebanon.

Short of a major shift in terms of service (and ethical framework), **the ability to quantitatively monitor and analyse WhatsApp is unlikely to become a possibility**. A qualitative assessment of WhatsApp usage in relation specifically to tensions would however be of relevance. The voice note analysis tool used by UNDP in recent surveys may be worth expanding to explore the platform's role in social tensions. A further exploration of the Instagram platform would also be worthwhile as it is likely to grow in importance in these discussions. Currently owned by Facebook, quantitative analysis of content on Instagram would require similar permissions to that of Facebook, although permissions appear more challenging to access and would likely require either a strong relationship with the platform, or the development of an app that fits the permission requirements but is limited in its analytical capability (short of being deployed outside of its original use case).

Finally, interviews also revealed that any efforts to more systematically analyse social media narratives on refugee-host community relations would be **of interest to many organisations and researchers in Lebanon**. UNDP should consider sharing the findings of this report, and convening a discussion with other interested organisations to explore opportunities for continuing this work. In taking this work further, UNDP should also closely engage with local communities, from both refugee and host communities, to ensure that this research does not become extractive.

B) Challenges and opportunities of a tool to engage polarized online users on refugee-host community relations

Amid unanimous understanding that social media platforms are playing a role in influencing refugee-host community tensions, few organisations are working to intervene in this space (AUB, Banayet Box and Smex stand out as examples). Those that are intervening online agree that **more could and should be done to combat the increasingly toxic nature of online conversations**, although there is no clear view on who is both capable of and responsible for doing more. The difficulty of this work is clear: as one interviewee put it 'shifting public opinion is so challenging ...[because] divisions are so wide, even within the Lebanese community'.²²

Current interventions to combat polarization and hate speech on social media have a tendency, both at a global level and at a Lebanon-specific level, to target those individuals who are already to some extent concerned with the polarization or hate speech in their online environment. Initiatives such as '101 facts about refugees' are highly commendable, but they require users to "opt in", to be interested in identifying factual information. In other words, they are not designed to reach out to individuals who consume, spread and potentially act on negative and false sentiment online - with limited concern for its veracity.

Moreover, there is a growing recognition amongst donors that funding on the link between social media, polarization and conflict has gone towards analysis (with a view to predicting or detecting conflict trends) and not towards testing methods for intervention (prevention or early response). However, as demonstrated in our scoping, social media analysis is an effective way of identifying how trending narratives are being framed, what sentiments they are evoking, and the users influencing them, especially after an event has taken place. Social media analysis is less effective at providing signals or indications that a particular event is about to take place, or at any rate it is no more effective than other traditional forms of monitoring social tensions. As a result, **interventions on social media should shift attention away from early warning and more towards conflict prevention through narrative shifting**. Concretely:

²² Interview with Nasser Yassin, Director of Research, Issam Fares Institute for Public Policy and International Affairs, American University of Beirut

- ◇ Social media analysis could be used to design a **strategic communications campaign to influence attitudes and behaviours** towards refugee-host community sentiment. Through in-depth analysis of online audiences, influencers and segmentations of online communities, researchers can understand different narratives further and delve into the factors that affect discourse for or against certain issues, or specific communities. It would then be possible to counter the prevailing tension-heavy narratives and test for behavioural changes.
- ◇ UNDP should consider engaging with **social media influencers and equipping them with tools to better shape local narratives**. This would build on efforts by UNDP to equip conventional media actors through the 'Journalists Pact for Strengthening Civil Society'²³. The strategy of equipping social media influencers with a broad reach across and within communities recognises the limited ability of actors like UNDP to influence the online conversation (see Section F). Efforts should be made to identify online influencers that have strong reach at a local level to ensure messages have impact.
- ◇ **A strategic communications campaign could be designed to connect to offline dialogue**. A handful of actors (Moonshot CVE, ISD Global, Build Up, the ACLU) are working on interventions that identify social media users at risk of polarization or radicalization, and use a variety of semi-automated methods to target them with messaging and draw them into a dialogue. Most of these approaches are still highly experimental, and none focus specifically on refugee-host community tensions. **These strategic communication approaches could fit well with UNDP's current initiative under the Lebanon Host Community Support Program (LHSP)**, which uses committees for its Mechanisms of Social Stability (MSS)

programming. These committees, made up of moderate actors and diverse stakeholders in conflict-prone areas, could be a launchpad for changing discourse around certain topics.

UNDP could use a combination of social media listening and offline tension mapping to identify issues pertinent to the committees, and then bring together online influencers to design a locally-appropriate strategic communications campaign to shift narratives and promote dialogue. Working with local influencers both online (social media influencers or prominent actors in the online space) and offline (local leaders, community influencers etc) would be a compelling strategy.

23 http://www.lb.undp.org/content/lebanon/en/home/library/crisis_prevention_and_recovery/media-monitoring-studies.html

Annex I: Ethical guidelines

The ethics around social media research is a complicated and constantly shifting conversation, as lines between civic, public, and private are blurred or outright scrambled. Under the strictest interpretation of terms laid out by the companies in question and the most strenuous research ethics, this would be a nearly impossible task. Navigating this line is an ongoing conversation in academic circles, and many institutions have published social media research ethics and guidelines, some of which we have consulted for further guidance.

Our starting point for ethical considerations is that this is not academic research; it is part of a process to find actionable methods of monitoring and intervening in social media conversations. The following ethical guidelines formed the basis of this research:

1. All raw data should be treated as confidential material and stored with appropriate safety measures by Build Up.
2. All raw data will be anonymized for analysis, and reporting purposes. If PII is discovered within a data set, it will be removed, though there will be no active search to determine its presence.
3. Raw data should be used by Build Up solely for analysis and reporting. No raw data should be shared with UNDP or any other party, except in cases where a sample excerpt should be used to illustrate a point. Trackback links to individual conversations will be kept for the research team's internal use but will not be shared.
4. In the cases where user profiles are identifiable, Build Up and UNDP should make no attempts to connect user profiles and determine specific individual behavior either within a social network, between social networks, or between social networks and the real world. Models of individual behavior will not be built.
5. Should any user of a platform ask if their information is within the dataset, Build Up will make a best-faith effort to answer affirmatively or negatively, and will respond to requests to delete the data.
6. Build Up and UNDP should understand that security systems are fallible and anonymization is reversible, and that data exposure risks exist regardless of measures taken to protect against them.
7. Additionally, Build Up will follow the Risk Mitigation and Research Ethics sections laid out in the Technical Proposal.
8. Based on Build Up's prior experience with social media interventions, Build Up and UNDP should understand that these guidelines inform this research project only, and our recommendations for interventions to engage polarized users on refugee-host community relations will outline any additional or different risks and conflicts related to social media companies' Terms of Services.

Annex II: Additional analysis tools reviewed

IBM Watson

IBM Watson is a software that uses AI to determine, amongst other things, sentiment behind online content, general categories per post (such as business, war, or fashion), and the relationship between the entities in the post. In order to assess the validity of this software for this kind of research, we entered 698 Facebook posts and 1,445 Tweets that we had already analysed, deemed relevant and coded according to our own sentiment categories, into IBM Watson. The results showed that all three functions that were potentially of interest and offered for Arabic language (sentiment, category and relation) were deemed ineffective and not advanced enough for use in this research, at least for Arabic language posts.

For example, tweets that were clearly antagonistic were coded positive, whilst others referring directly to the conflict in Syria were categorized as related health and fitness/men's health. These inaccuracies were repeated, making IBM Watson an ineffective tool for this research.

Graphext :

Graphext is a promising data-analysis tool that retrieves information from a number of platforms including Google, Google News, Twitter, Instagram and Facebook and provides data manipulation, computational and statistical tools, and visualizations.

We explored Graphext as a possible tool available on the market that accomplishes similar tasks to our work on this report. Our work was unable to proceed when we discovered the rendering of Arabic text in the visualizations was broken, making some of our key analysis work impossible.

Ultimately, we believe Graphext could be a promising platform for targeted social media monitoring in and if the language implementation is fixed, it would be worthwhile to investigate it more thoroughly.