#Corona: a negativity pandemic for urban dwellers?

Dick Vogel[1] *& Evert Meijers*[2]

[1] Esri Nederland (dvogel@esri.nl)
[2] Utrecht University, department of Human Geography and Spatial Planning, The Netherlands (e.j.meijers@uu.nl)

Working paper[i], Utrecht University

Abstract

Almost 2,5 million tweets were analysed to explore the variation of sentiment over space and time in relation to COVID-19 measures, infection rates and socio-geographical factors in the Netherlands. The sentiment analysis of geotagged and timestamped tweets shows that the mood of people was more negative during the first waves than in later stages. The period containing the strict lockdown and the curfew was not significantly more negative than open periods, which may be attributed to vaccines becoming available. Local COVID-19 infection rates do not explain the great spatial variance in sentiment across Dutch municipalities. However, we found clear evidence for higher urbanisation levels leading to higher negativity rates. Urban residents struggled more during the COVID-19 outbreak compared to people living in smaller towns or in the countryside. We discuss whether this could lead to counterurbanisation.

Keywords: Twitter | Sentiment Analysis | COVID-19 | Lockdown | Social Media | Netherlands

Acknowledgement: We would like to thank Daan Goslinga, Raoul Kalisvaart, Adriaan Pardoel, Cathrine Paulsen and Sarah de Wolf for developing the purpose-built Citysent database, and Caroline Newton for co-supervising the development of the database.

Universiteit Utrecht

Sharing science, *shaping tomorrow*

#Corona: a negativity pandemic for urban dwellers?

1. Introduction

Since early 2020 the world is battling the COVID-19 virus. At the time of writing, the Coronavirus Resource Center of Johns Hopkins University just announced the death toll had exceeded 5 million, while almost 250 million people had been diagnosed with the coronavirus SARS-CoV-2. These are confirmed numbers, real global deaths were estimated to be at least twice the official figures based on excess mortality rates (IHME, 2021, October 15 2021 update). Whether having personally experienced an infection or not, virtually all people have been at least indirectly affected in one way or the other by policy measures taken to control the pandemic – lockdowns, curfews, quarantines, social distancing and movement and gathering restrictions have been implemented in many countries. It is well known that these affect the well-being of people, and studies have reported changes in the mood of people, as the pandemic and related restrictions lead for instance to increased tension, depression, anger, anxiety, fatigue and confusion and less vigour (Terry et al., 2020)

This paper is concerned with the overall sentiment of the Dutch population and how it has been affected by the implementation of pandemic control measures and their liftings, but particularly considers the spatial variation in this sentiment. While most COVID-19 related geographic research has focused on the spread of the virus in relation to urbanization (e.g. Kuebart and Stabler, 2020; Boterman, 2020), we are concerned with the implications of the pandemic and the measures on the general mood of society and try to distinguish whether geographic factors related to the natural and built environment, as well as heterogeneity across places and regions in socio-economic terms, have an impact on this mood under different pandemic control policy regimes. This insight is essential to be included in larger debates about the impact of the pandemic on the future of the urban and the rural, as much speculation, whether based on theoretical arguments, historical parallels and/or anecdotical evidence in media outlets has risen on this topic. For instance, counterurbanisation is considered a likely outcome (Denham, 2021), fostered by considerably enhanced levels of working from home and telecommuting, which give new freedoms in location choices (De Vos et al., 2018). Others argue that in this experimental and seemingly temporary period that the pandemic is, patterns of urbanization are not likely to change drastically (Florida et al., 2021; Bandarin et al., 2021). We may assume however that places that are associated with higher levels of well-being are more attractive to live. The aim of this paper is therefore to inform this debate on post-pandemic urbanization trends by understanding where the mood of the population was more positive during different stages of the pandemic.

An increasingly popular resource to gauge public sentiment is user-generated content on social media, and millions of short online expressions such as tweets can be automatically gathered, processed, and identified, meaning that research can be conducted on a larger population, over a longer period and based on more observations (You, & Tunçer, 2016). Especially Twitter is important, as several studies showed the association between Twitter sentiment and levels of self-reported subjective well-being (Dodds et al., 2011; Volkova et al., 2016; Yang & Srinivasan, 2016), and it seems plausible that the effects of COVID-19 are visible in Twitter data. This study uses a sentiment-analysis service called 'CITYSENT' to analyse geo-tagged tweets in the Netherlands.

The existing literature relating social media sentiment to the COVID-19 pandemic primarily focuses on opinion mining as opposed to sentiment analysis, as they merely analyse the sentiment specifically *towards* COVID-19 (e.g. Manguri et al., 2020; Wang et al., 2020) instead of the overall sentiment *during* the COVID-19 period, which we focus on. One particularly interesting dimension of the pandemic is that traditional patterns in twitter sentiment across space (urban versus rural) and across different demographic variables (e.g. gender, age) seem to be reversed, as different areas and groups are affected more by restrictions that were imposed (e.g. Van Leeuwen and Bourdeau-Lepage, 2020), which makes a general sentiment analysis of the entire population of twitter users more opportune than just focusing on those that mention COVID-19. Moreover, whereas initial research into sentiments during the pandemic (necessarily) focused on relatively short time periods at the start of the pandemic (Boon-Itt and Skunkan, 2020; Zhou et al., 2020), or focused on specific groups in Dutch society such as the elderly (Van Tilburg et al., 2020) or students (Taquet et al., 2021), we are able to distinguish different time periods according to policy restrictions in place (which are inherently linked to the spread of the virus) and for society as a whole, allowing to assess the spatial variation of these restrictions over time. That is, as far as Twitter is representative of that society, but more on that later.

This paper is structured as follows. First, in section 2, the theoretical framework addresses the literature shedding light on the question whether social well-being can be captured through tweets, before turning to the question how different COVID-19 related factors (infections, policy regimes), socio-economic factors and geographic factors influence twitter sentiment. Next, in section 3, we detail our data collection and research approach. In section 4 we turn to our results, and in the last section 5 we discuss the wider implications of our findings, with a particular focus on future of urbani.sation trends.

2. Theory

2.1 You are what you tweet?

To understand the effects of COVID-19 measures on regional and temporal variation in Twitter sentiment, one must first consider what tweet sentiment reveals about the user and how this gets affected by external events. Generally, online user-generated information is conceived to be a fine source of sentiments and public views on societal issues or political decision-making (Giachanou and Crestani, 2016) and in particular the around 500 million tweets every day, transmitted by roughly 330 million users worldwide (Statista, 2019), have been considered 'social sensors' as Twitter users can react relatively spontaneously to real-time events or circumstances on the spot through their mobile phones (Günnemann & Pfeffer, 2015).

A question is whether subjective well-being and happiness can be observed and analysed through social media like Twitter. Ed Diener (2009), a prominent expert on the field of happiness, states that subjective well-being is induced by three factors: (1) the presence of positive sentiment, (2) the absence of negative sentiment and (3) life satisfaction. The last mentioned is a longer-term cognitive evaluation of individuals' own lives, for which twitter sentiment is not accurate (Yang & Srinivasan, 2016). However, the presence of positive and negative sentiment (also called affect) is largely determined by daily events (e.g., weather, activities or food; Diener, 2009). As people have subjectively experienced the impact of COVID-19 related policy measures and infections on a daily basis, it would be logical to assume that tweets can capture some dimensions of subjective well-being. Studies have indeed shown that meaningful observations about happiness can be gathered from Twitter and that these observations are uniform to earlier findings in social science on this matter (Dodds et al., 2011; Volkova et al., 2016; Yang & Srinivasan, 2016). In other words, negative or positive sentiment in tweets is not disconnected from the level of social well-being of its writers (Curini et al., 2015; Durahim & Coşkun, 2015).

Although the correlations are undeniably not strong enough to make reliable conclusions on an individual Twitter user's social well-being based on his or her tweet content, big datasets containing thousands or even millions of tweets could offer understanding on the overall level of well-being (among Twitter users) in a region or country and compare these. For instance, Schwartz et al. (2013) studied word use in tweets from different US counties and found it to be predictive for survey based social well-being levels of inhabitants of those counties. More specifically, they found that words in tweets associated with exercise (e.g., 'training', 'gym', 'fitness'), outdoor activities (e.g., 'camping', 'trip', 'sea') and engagement and social activities (e.g., 'meeting', 'experience', 'we') correlates with a higher level of social well-being as found for counties. Conversely, tweets containing words indicating withdrawal or disengagement from society or activities (e.g., 'bored' and 'tired'), as well as the use of slang and swear words, are found to show a negative correlation with social well-being measured as more traditionally on US county level (Schwartz et al., 2013). Taking these patterns further, it could be expected that the COVID-19 measures negatively affected subjective well-being, especially in those time periods that social and physical activities were more restrictive (e.g. closing of gyms and disallowing team sports) probably leading to less engagement, exercise and social activities and more boredom. Some studies do shed light on this. Zhou et al. (2020) exploited Twitter data over a five-month time span in Australia during the COVID-19 pandemic and compared it with other data (e.g., infection rates) and government policies (e.g., lockdown). They found an overall positive polarity, however the overall sentiment significantly decreased in positivity when the number of confirmed cases increased in the same period. Furthermore, the lockdown policy imposed by the Australian government had a further decreasing effect on the positivity rate in tweet sentiment (Zhou et al., 2020). Specifically for the Netherlands, Wang et al. (2020) compared the sentiment polarity of Dutch tweets in general with the sentiment polarity of Dutch tweets related to COVID-19 and found the latter to always have a more negative sentiment score, meaning that the public is significantly more negative towards COVID-19 compared to generic topics.

2.2 Regional variation in twitter sentiment

Some studies have examined twitter sentiment particularly also to detect geographical patterns, and these generally find that tweets are far from randomly distributed across space, largely reflecting regional variety in demographic, geographic and health characteristics (Mitchell et al., 2013). Cheng et al. (2010) even managed to create a model that can predict a Twitter user's location by solely using the content of the user's tweets, even when other geographical indications were absent. These studies evidently signify that tweet content is not only dependent on the characteristics of the user, but also on the geolocation of a particular user, which is an important assumption for studying spatial variation in tweet sentiment.

Before studying possible regional variation in tweet sentiment due to COVID-19 and pandemic-related measures, it is useful to acknowledge the existence of such regional differences in Twitter sentiment, autonomous from the pandemic. Positive tweets are less recurrent in counties with greater economic disadvantages, a younger population and a higher percentage of ethnic minorities (Nguyen et al,. 2016). On an individual level, it is found that Twitter users being female, having children, being in a relationship and having a higher annual income and education, significantly express more positively in their Tweets, whereas users being younger and liberal tend to express more negative sentiment while tweeting (Volkova & Bachrach, 2015). Consequently, substantial differences between regions in terms of such socio-economic and demographic differences need to be accounted for when comparing regions.

The pandemic may however also have a differential impact according to these characteristics. Compared to elderly people, young people's mental well-being is more affected by COVID-19 restrictions, and women experience higher levels of stress compared to men during quarantine (Kowal et al., 2020). This latter finding was however in contrast to the studies by Limcaoco et al., (2020) and Li et al. (2020). Kowal et al. (2020) also found that married or cohabiting people seem to be affected less by COVID-19, as indicated by lower levels of stress than single people during the pandemic. And income may well play a role as more affluent households can afford more spacious housing, and more private outdoor space.

The impact of the pandemic on twitter sentiment is likely to also vary over space. In terms of local breaking news or external events, spatial distance has been found to be associated with a decrease in tweeted words related to that particular news or event, as well as overall sentimental expression (Doré et al., 2015). Based on that study, one could predict that individuals and regions that are affected less by COVID-19 and the restrictions exhibit less disturbance in Twitter sentiment data throughout the pandemic. This assumption is supported by an Australian study, in which regional variation in (exclusively COVID-19 related) tweet sentiment between local government areas was investigated (Zhou et al., 2020). They found that the negative sentiment in some local areas significantly deviated from other areas, which the authors attributed to local COVID-19 outbreaks. We may therefore expect that local exposure to the virus matters for twitter sentiment.

The built environment may also be of influence, not only in terms of risk of contagion, but also in terms of the impact of pandemic-control measures. People that reside in urban areas might get confronted more intensely with the measures because there are less possibilities to recreate in nature and they face more closed amenities compared to people that reside in rural areas. Moreover, social distancing is simply more easy in less dense environments. The study by Van Leeuwen and Bourdeau-Lepage (2020), who measured spatial variation in the impact of the Dutch lockdown on well-being by means of questionnaires, partly confirms this presumption. Results showed that on average, the overall well-being decreased in the Netherlands during lockdown, however this decrease was greater in dense urban regions than in the least urban regions. What makes these results more remarkable is that significant differences in well-being between rural and urban regions were observed during the lockdown, while before the lockdown this was not the case. According to the authors, this relatively great decline in well-being for urban residents is due to being bored more frequently and a more significant reduction of exercise compared to rural residents (Van Leeuwen & Bourdeau-Lepage, 2020). Furthermore, people residing in an apartment without a balcony are found to be the most unhappy during the lockdown period.

Also the natural environment plays a role, and having access to green space has in particular been important. Studies already showed that tweets in green spaces (e.g., parks or gardens) contain more positive and less negative sentiment (Lim et al., 2018). During the first COVID-19 outbreak in the UK, Poortinga et al. (2021) showed that access to public green spaces is related to higher levels of social well-being, and that green spaces are especially critical for individuals without a private garden.

3. Research approach

3.1 Data

To analyse public sentiment changes across space and time during the COVID-19 crisis, a purpose-built dataset called CITYSENT is utilized, created by IT-specialists from Delft University of Technology (Goslinga et al., 2020). The system collects the Twitter data stream from the Netherlands (and Flanders in Belgium, but this region is not our focus here) to analyse sentiments combined with geographical data. The processing of tweets is visualised in Figure 1. First, a 1% random sample of tweets originating from the Netherlands is retrieved and pre-processed. The latter involves that Dutch-language tweets (about 80 % of the total) are automatically translated into English and cleaned by removing emoticons and images, videos, hyperlinks, Twitter-specific words (e.g. "RT") and user-mentions since our sentiment analysis tool only handles text and, except for emoticons, do not hold any sentimental value and such cleaning increases accuracy of the sentiment analysis (Angiani et al., 2016). Next, punctuation and numbers are deleted and words in tweets are stemmed, meaning that variations of words are replaced by their basic form or stem (e.g., 'saddest' and 'sadly' become 'sad').

Subsequently, the sentiment analysis processes the content of the tweets and categorizes their sentiment into three classes: 'positive', 'negative' or 'neutral'. The classification of tweets occurs in two stages. The first stage is based on the sentiment scores (also known as polarity scores) that have been given to each individual word of the tweet; this method is called lexicon-based analysis (Taboada et al., 2011). The sum of the score of all the words then determines the overall sentiment score of the piece of text. Such lexicon-based analysis has the advantage that its functioning is uniform to different types of texts, but it cannot autonomously co-evolve with the dynamic, constantly changing language on Twitter (Giachanou & Crestani, 2016). In addition, polarity scores are calculated on the base of individual words, without taking the context of the complete sentence into consideration, and especially for tweets that contain sarcasm, irony or cynicism this can affect the validity of the sentiment score. Therefore, the only purpose of the Lexicon-based analysis in CITYSENT is to establish whether a tweet contains enough sentimental words to be classified as either positive or negative. Tweets that do not apply to this condition, are being labelled as ‘neutral’. In a second stage, only those tweets made up of sufficient sentiment-containing words are analysed in more detail. The second stage consists of the usage of a machine learning model, a learning algorithm that is trained to categorize fragments of texts, to support the sentiment analysis process. The specific model that has been operationalised is called ‘Support Vector Machine’ (SVM), since this model appears to be most accurate compared to other options (Goslinga et al. 2020). Before the SVM can perform annotation tasks, the model needs to be trained by ‘feeding’ it existing data, with the purpose that it learns to predict independently when it is being confronted with new data. The developers of CITYSENT therefore supplied the learning model with a pre-developed dataset created by Stanford University called ‘sentiment 140’, containing 1.6 million tweets that have been classified by their sentiment class (Go et al., 2009). Based on this input, the machine learning model is now able to recognize patterns and can independently classify sentiments of short texts as positive or negative. This determination is the sentiment classification that eventually ends up in the final dataset.

Figure 1. Schematic model of CITYSENT's dataflow.

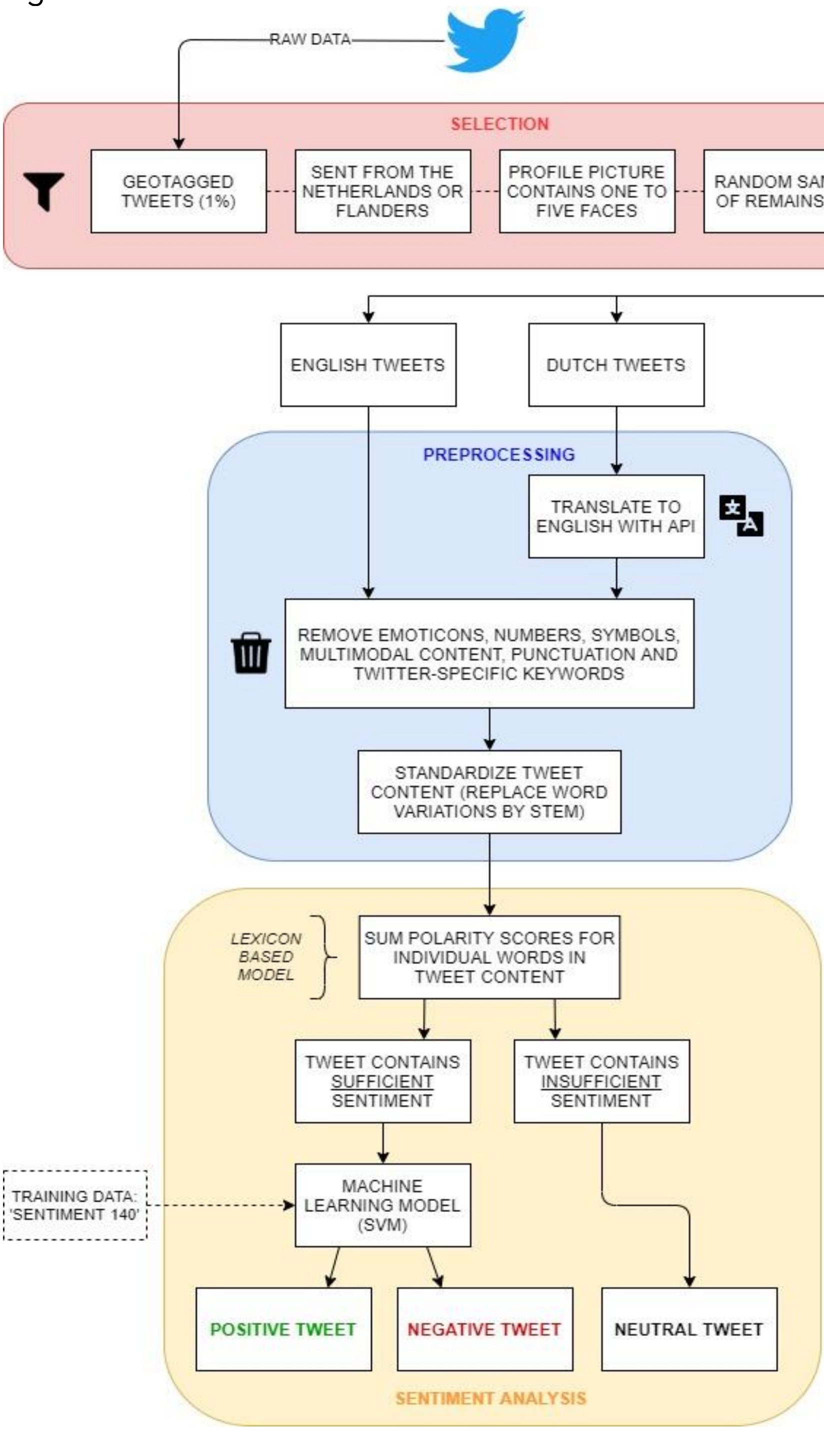

Since the sentiment variable is a crucial factor in this study, it is important to dilate upon the accuracy of sentiment classification by CITYSENT. Their creators found that 82,4% of the tweets were categorized in the right sentiment group (either positive, negative, or neutral) when it was provided with unseen data (Goslinga et al., 2020). We took another sample of 104 tweets, looked up the original and two different people classified them according to sentiment. We found an accuracy of 70,19%, but only 6 out of the 31 tweets (5,77% of 104 tweets) that we labelled differently were problematic in the sense that they were classified as positive in the dataset but classified as negative in the sample, and vice versa. The dataset is therefore still adequately usable for the intended research, especially given the immensity of the dataset, which contains millions of tweets. This suggests that the margin of error will remain roughly consistent when different time periods are being compared. However, it is important that the reader realises that 100% accurate classification is unattainable and that for instance irony and sarcasm are hard to classify. What is important for this paper, however, is the question whether there is spatial variation in incorrect labelling that would bias a geographic analysis. We think there is no theoretical or methodological reason to assume this is the case, also because the Netherlands is a relatively small country with limited variation in how language is used. There is a second language, Frysian, but only tweets in Dutch were analysed.

Measuring the location of tweets is a crucial element for our analysis, and we focus on the location from where a tweet was sent. Tweet location is collected by reading the 'geotag' of a tweet, which is the precise location from where a tweet is posted from. This feature is however disabled by default, meaning that it is required to be manually switched on by Twitter users before geotags are included in tweets (Sloan & Morgan, 2015). Because of the importance of the geographical component, the creators of CITYSENT decided to only incorporate geo-tagged tweets in the dataset. This appears to be a filter of large significance, as the developers observed that around 1% of all tweets are geotagged (Goslinga et al., 2020). This observation is in line with Sloan et al. (2013), who found a similar percentage (0,85%). Earlier studies did not find a big difference in demographics of Twitter users who do enable location services and the ones who do not (Sloan et al., 2013).

A large share (around 53%) of tweets in our dataset were not sent from personal accounts, but from accounts that represent a business or an organisation (Goslinga et al., 2020). These tweets can contain sentiment-carrying words (e.g., weather or news updates), however they do not depict the sentiment of the overall population (they could even be automated tweets sent by bots) and could therefore significantly disturb the accuracy of the data. To solve this problem, a pre-trained face recognition model analyses the profile picture of the sender of every tweet. Tweets of users that have no face shown on their profile picture thereafter get discarded. This intervention ensures that most tweets from businesses and organisations are excluded from the dataset, since those users generally use logos or texts as profile pictures. Additionally, tweets from users with profile pictures containing more than five people are also excluded to ensure that tweets are statements originating from individuals and not groups. The face recognition model has been found to verify faces with an accuracy of 99,38% (Goslinga et al., 2020). This model also estimates users' age and gender, based on their profile pictures, albeit with insufficient accuracy to be included in our analysis. Finally, the date and timestamp of every tweet is stored.

Our dataset starts with the CITYSENT tool becoming fully operational on June 11, 2020 and ends almost one year later on June 3, 2021, a moment when the pandemic seemed under control, the vaccination rate was rapidly increasing and society had practically returned to normal in the Netherlands. Note that this means that our data collection started during the pandemic, which prevents comparisons with pre-pandemic periods. CITYSENT is an autonomous system that is supposed to operate continuously throughout its activation period. Unfortunately, several days and even entire weeks (early December 2020, March/April 2021) spread over the study period are missing, presumably due to server defects or upgrades. In the end, 2,4 million tweets were analysed.

3.2 Model: explaining negativity rate

The tweets in the CITYSENT dataset provide the basis for the creation of a new dataset, that aggregated those tweets according to time and space. This includes the calculation of our dependent variable that captures the sentiment of residents in the Netherlands, which is the 'negativity rate'. This is measured by the ratio of negative tweets as a percentage of the total number of negative and positive tweets, as the following formula encapsulates:

$$Negativity\ rate\ = \frac{total\ number\ of\ negative\ tweets}{(total\ number\ of\ negative\ tweets\ +\ positive\ tweets)} \times 100$$

This negativity rate was calculated for all Dutch municipalities (#343)[1] for five different time periods that we can discern during the year of observations we collected. Since these periods are distinguished on the basis of the level of strictness of the pandemic-control measures in place (which are inherently linked to the spread of the pandemic), we discuss these in more detail below under 'COVID-related factors'. In addition, we present the other sets of independent factors, namely socio-economic factors and geographical factors that will be incorporated in our regression models explaining the negativity rate.

3.3 COVID-19 related factors

We employ two factors related to COVID-19 when explaining the negativity rate. The first is policy-related, as we split the time period of our study into five subperiods that are characterised by different policy regimes in terms of the level of strictness of pandemic-control measures. Figure 2 presents the COVID-19 timeline in the Netherlands for our study period. This was preceded by a period in which some policy measures had been in place, somewhat pretentiously framed as an 'intelligent lockdown' by Dutch Prime Minister Rutte.

<u>Period 1 A relatively 'normal' summer (11 June 2020 - 27 September 2020)</u>: This first period of the data collection is characterised by relative freedom and openness compared to previous months. High schools, restaurants and bars are (partially) re-opened and everyone had the possibility to get tested. However, the economic impact started to become clear as reports predicted a decrease of 6% in GDP and a doubling of unemployment (CPB, 2020). In July infection rates were still decreasing making the government allow gatherings of maximum 100 people. The relatively 'normal' summer proceeded in August, however concerns increased as the virus started to quickly spread again, especially among young people.

<u>Period 2 Infections increasing and restrictions returning (28 September 2020 - 13 December 2020)</u>: On 29 September the first serious restricting measures (i.e., limited number of visitors per household, bars and restaurants closing earlier, no audience at sport events) were

[1] Ten municipalities were not included in the analysis due to lacking or inaccurate data from the CITYSENT system, or because place name disambiguity made it impossible to allocate tweets to similarly-named municipalities (Bergen).

announced since the first wave of infections in March 2020. It would be the start of a series of restricting numbers to control the reproduction number of the virus that fuelled the second wave of infections. In October, the Netherlands went into a 'partial lockdown', which implied that bars and restaurants were fully closed, events were prohibited and team sports restricted. Although this seemed to have a positive effect in November as infection rates decreased, this trend stabilised later in the month.

Period 3 Lockdown and curfew (14 December 2020 - 22 February 2021): On December 14 a lockdown was announced, which meant that non-essential shops, gyms, schools, day-care and public facilities like museums and swimming pools fully closed down. These measures would initially last until midway January, however due to fears for mutations of the virus, the lockdown was extended till March. Additionally, the Dutch Government followed the example of most of its European neighbouring countries and implemented a curfew on the 23rd of January. Between 21:00 till 4:30, people were not allowed to go outside (except for workers and dog owners). Both the lockdown and curfew remained in place longer than initially expected and announced.

Period 4 Slowly reopening after lockdown (23 February 2021 - 27 April 2021): As the average infection rates dropped during the first months of the new year and the first groups were being vaccinated, the government could start to contemplate about the timing of lifting restrictions. However, labelling a clear ending mark to the lockdown is difficult because rather than withdrawing all the restrictions concurrently, the government gradually re-opened the country over a few months. For example, already on the 8th of February 2021, the first relaxations of the lockdown were announced as the government decided to open up primary school and day-care facilities, as well as non-essential shops to pick up purchases. However, most drastic changes occurred from the end of February onwards when high schools partly re-opened, contact professions could start again, non-essential shops could welcome some visitors and team sports were allowed under certain conditions.

Period 5 The beginning of the end? (28 April 2021 - 3 June 2021): In April colleges and universities reopened, as well as terraces at bars and restaurants. Non-essential shops could accept more visitors, people could invite more visitors at home and the curfew was abolished (28th of April). This day marks the start of the last study period and is similar to the first period in the sense that it is characterized by relative openness and push for public freedom. Gradually easing restrictions was possible for the government due to the decreasing infection rates and the growing share of the population that was vaccinated. Throughout May, the country continues to gradually return to a 'normal' society with the reopening of gyms, libraries, zoos, amusement parks and less restrictions in terms of sport and cultural events. The 3rd of June is the last day of the data collection period, and not much changed in June.

Figure 2: Timeline of the COVID-19 periods including the infection rates and restrictions (infection data from RIVM, 2021)

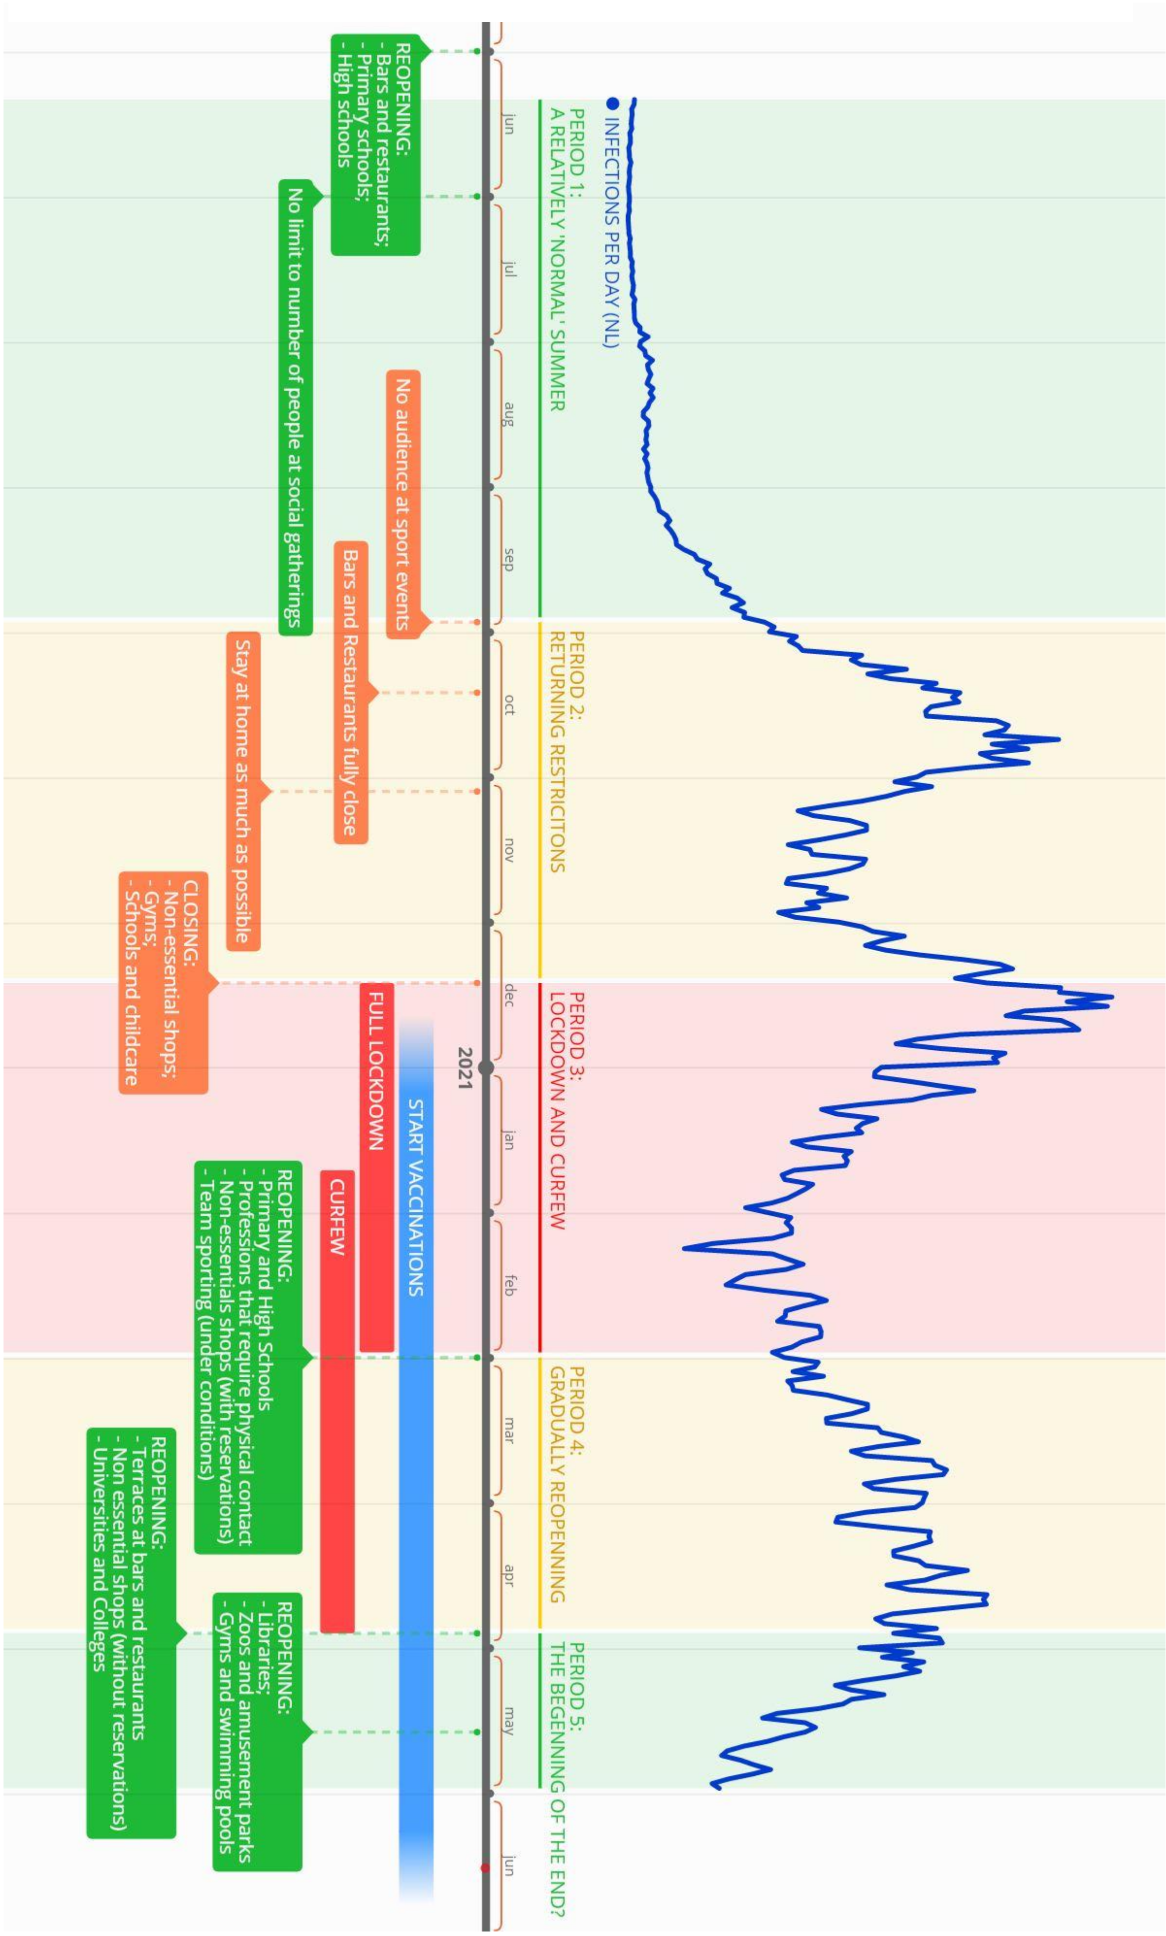

Infection rates

A second variable that may explain the negativity rate in a municipality obviously is the extent to which people are affected by, and confronted with the virus. Zhou et al. (2020) showed that overall sentiment significantly decreases in positivity when the number of confirmed cases increases. To capture such an effect in the Netherlands, we calculated the *local infection rate*, measured as the number of infections per 1.000 inhabitants in each of the five time periods. This data was provided by the National Institute for Public Health and the Environment (RIVM, 2021).

3.4 Socio-economic factors

Although the ability to collect enormous amounts of information on a particular population is a massive advantage of Twitter data, it should be reckoned that Twitter users do not inherently represent the general population. Social media studies essentially rely on people who are willing to supply the needed data. In 2020, less than twenty percent of the Dutch population used Twitter, even a smaller percentage is an active user (Van der Veer et al., 2020). Moreover, it is plausible that the demographics and other characteristics in social media samples do not align with the overall population. For example, Twitter users in the UK are wealthier, younger and better educated than non-Twitter users, and similar results have been found in the US for the factors age and wealthiness (Blank & Lutz, 2017). Gender does not seem to be disbalanced among Twitter users compared to Census data (Blank & Lutz, 2017; Sloan et al., 2013). This all likely applies to Dutch twitter users too, and Van der Veer et al. (2020) found for instance that the 20-39 age group is better represented than the 65+ age group. What is more, such socio-economic characteristics have a possible impact on the sentiment that is expressed, as was discussed in our theoretical section. As was explained before, the results of an exercise to obtain information on the micro-level of Twitter users based on their profile picture were unsatisfactory, which is why we try to control for a variety of socio-economic characteristics using data on the municipal level that may influence sentiment during the pandemic:

- Gender, measured by the *percentage men* in a municipality (2020).
- Single households, measured as share of all households (2021).
- Income, measured by *average disposable income* of private households (excl. students) (2019) per municipality.

3.5 Geographic factors

As was discussed in the preceding theoretical section, the role of geographical factors influencing sentiment in times of a pandemic is likely to be large. Even so large that traditionally much more negative farmland areas (Cao et al., 2018), now seem to have a higher level of well-being, being less affected by lockdowns than urban dwellers (Van Leeuwen & Bourdeau-Lepage, 2020). We capture these factors by a common measure in the Netherlands called 'urbanity level' provided by Statistics Netherlands (2021), which is based on the density of addresses within a 1km radius, calculated for the centre of every 500m by 500m grid cell, and then takes the average density for the municipality. Given that Dutch municipalities are large and often contain multiple places, this is to be preferred over general population density or population size. The urbanisation level is representative for other factors, including the presence of amenities and jobs. We reversed the order of the original variable so that higher values represent higher levels of urbanisation.

Table 1 presents the descriptives of all non-dummy variables in our model. The missing values for our dependent variable, Negativity rate, are due to some low-density municipalities not containing enough tweets in every period, which does not guarantee a representative negativity rate for their particular place and time. Specifically, municipalities that contain less than 100 tweets per period have been removed from the dataset. Some of the independents have missing values due to recent mergers of municipalities. Table 2 presents the correlation matrix.

Table 1. Descriptive statistics of the variables

	N	Minimum	Maximum	Mean	Std. Deviation
Negativity rate per period per municipality	1546	3.17	55.17	24.82	6.62
COVID-related factors					
Average infections per 1.000 inhabitants per period	1715	0	68.37	18.35	12.12
Socio-economic factors					
Percentage men	1710	47.07	52.83	49.78	0.76
Percentage Single Households	1715	6.19	34.95	15.07	4.21
Income	1685	35.2	102.6	47.74	6.44
Geographical factors					
Urbanisation level	1715	1	5	2.69	1.15

Table 2. Correlation matrix.

	Negativity rate per period per municipality	Average infections per 1.000 inhabitants per period	Percentage men	Percentage Single Households	Income	Urbanisation level
Negativity rate per period per municipality	1					
Average infections per 1.000 inhabitants per period	.031	1				
Percentage men	.006	.053	1			
Percentage Single Households	.065	-.083	-.297	1		
Income	-.028	.045	-.246	-.434	1	
Urbanisation level	.099	.049	-.480	.585	-.138	1

4. Results

Figure 3 maps our dependent variable, the negativity rate, over the entire study period (right) as well as the spatial distribution of COVID-19 infections per 1000 inhabitants over the study period (left). Starting with the left map, there is a clear pattern visible in the distribution of the virus as the south of the country (mainly the provinces of Noord-Brabant and Limburg) manifested itself as the epicentrum of the pandemic in the Netherlands (so-far) while the Northern provinces coped with much lower infection rates. In contrast, the map of Twitter sentiment appears to be an incoherent mosaic. Negativity rates range from 5,38% (Woudenberg) all the way up till 42,56% (Mill and Sint Hubert). Remarkably, municipalities that are characterised by extreme negativity share borders with municipalities that are distinguished by extreme positivity, which makes it difficult to spot consistent patterns with the naked eye. For example, at first sight, the north-south contrast that is visible in the infection rates does not get converted to a similar pattern in the negativity rate.

Figure 3. The COVID-19 infections per municipality over the study period (left) and the negativity rate per municipality (right)

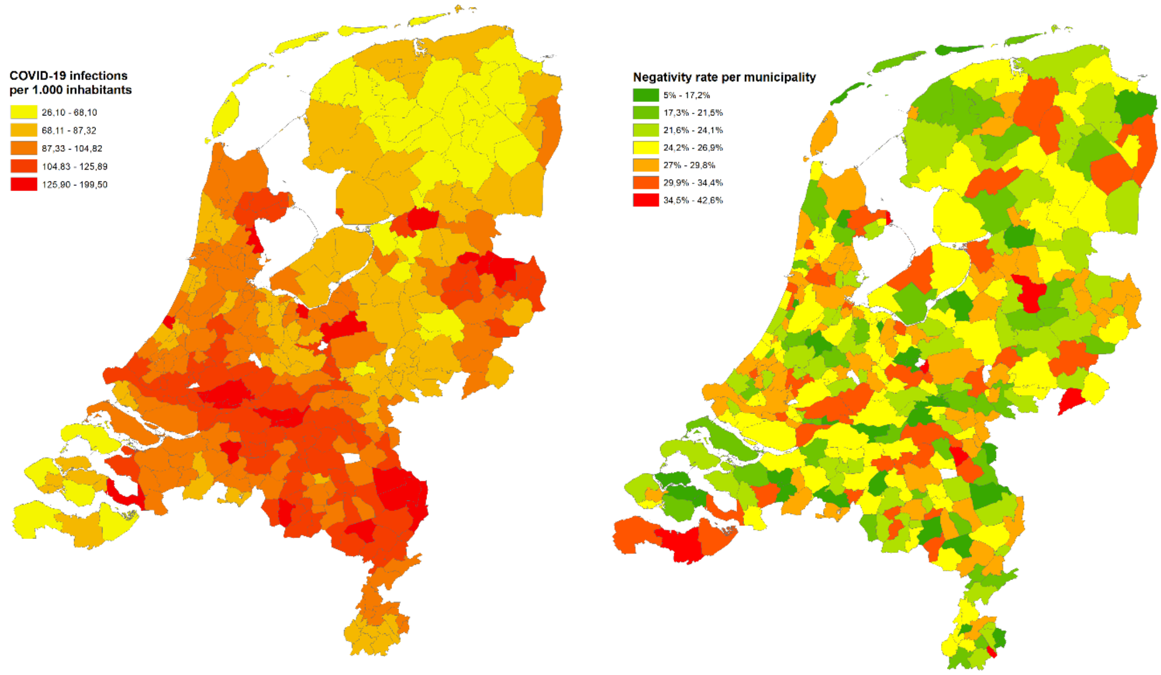

Our panel data includes both COVID-19 related variables that vary over time, and time-invariant explanatory variables that capture socio-economic and geographical factors. Breusch-Pagan Lagrange multiplier (LM) tests reveal that random effects is to be preferred over OLS regression (Table 3). Hausman tests confirm the choice of a random effects models over fixed effects models (Table 3). Since we have the availability of repeated observations on municipalities during the COVID-19 period, the random effects model allows inserting an additional term in the regression, capturing individual-specific, time-invariant factors affecting the dependent variable but that remain unobserved. The random effects models are presented in Table 3. We defined different models according to the sets of COVID-related, socio-economic and geographical factors that we use. As can be told from the model statistics, it is rather hard to explain twitter sentiment, and our variables of interest are only able to explain a fraction of the variance.

Table 3. Random effects GLS models explaining negativity rate of Twitter sentiment across Dutch municipalities during the pandemic.

	Model 1 COVID-19		Model 2 COVID-19 + Socio-Economic		Model 3 COVID-19 + geographic		Model 4 COVID-19 + urbanisation categories		Model 5 – All	
Variables	*b*	*SE(b)*	*b*	*SE(b)*	*b*	*SE(b)*	*b*	*SE(b)*	*b*	*SE(b)*
Constant	.2349**	.0046	.1711	.2283	.2179**	.0078	.2250**	.0080	-.0057	.2442
COVID-19 related factors										
P1 'Relatively normal summer'	.0152**	.0046	.0152**	.0046	.0153**	.0046	.0153**	.0046	.0147**	.0046
P2 'Infections increasing and restrictions returning'	.0137*	.0059	.0135*	.0059	.0144*	.0059	.0144*	.0059	.0147*	.0059
P3 'Lockdown and curfew'	.0008	.0059	.0006	.0060	.0015	.0059	.0014	.0059	.0019	.0060
P4 'Slowly reopening'	-.0022	.0055	-.0026	.0055	-.0027	.0055	-.0027	.0055	-.0017	.0055
(Reference period: P5 'Beginning of the end?'										
Infection rate	.0004	.0002	.0004	.0002	.0003	.0002	.0004	.0002	.0003	.0002
Socio-economic factors										
Gender (% men)			.0009	.0042					.0044	.0046
Disposable income			-.0000	.0005					-.0000	.0005
Single household			.0013	.0008					.0004	.0009
Geographic factors										
Urbanity level					.0063**	.0024			.0063*	.0032
Very Strongly Urbanised (dummy)							.0324*	.0127		
Strongly urbanised (dummy)							.0148	.0092		
Moderately urbanised (dummy)							.0119	.0093		
Somewhat urbanised (dummy)							.0057	.0085		
(Reference category: rural)										
N Observations	1546 (343 municipalities)		1524 (337 municipalities)		1546 (343 municipalities)		1546 (343 municipalities)		1524 (337 municipalities)	
Wald-Chi^2	31.88		32.71		39.21		40.09		36.65	
Prob > Chi^2	<.001		<.001		<.001		<.001		<.001	
R^2 Within	.0261		.0248		.0261		.0261		.0248	
R^2 Between	.0000		.0070		.0190		.0215		.0165	
R^2 Overall	.0127		.0166		.0247		.0263		.0244	
Hausman	1.61 (p=0.9)		1.27 (p=0.94)		0.74 (p=0.98)		0.76 (p=0.98)		1.42 (p=0.92)	
Breusch-Pagan LM test	321.70 (p=0.00)		305.09 (p=0.00)		306.04 (p=0.00)		303.90 (p=0.00)		293.34 (p=0.00)	

** p<0.01 * p<0.05

COVID-19 related factors

All models in Table 3 present a consistent pattern regarding the impact of pandemic control measures. With period 5 as a reference category, both period 1 (the relatively normal summer in 2020) and period 2 (when infections increased again and restrictions returned) are significantly more negative. Somewhat surprisingly, the period with the most drastic and strict measures, period 3 with the curfew and lockdown, is not significantly more negative than period 5. Another interesting find is also that the municipal infection rate per period does not significantly influence the negativity rate. Earlier studies that did find this relation (Zhou et al., 2020) can therefore not be confirmed for the Dutch context.

Socio-economic factors

Socio-economic factors (Models 2 and 5) seem to have a marginal influence on the negativity rate per municipality per period, as the variables income, gender and single households do not have a significant explanatory effect. Signals that women experience more stress during the COVID-19 pandemic, as described by earlier studies (Kowal et al., 2020), and theories that suggest that Twitter users with a higher annual income tweet more positively (Volkova & Bachrach, 2015) are therefore not confirmed. However, the finding that gender does not play a role may actually be suggesting that the pandemic has a bigger negative impact on women as we also could not confirm the widespread idea that men express themselves more negatively on social media (Volkova & Bachrach, 2015). Also the common assumption that living on your own during lockdown is more stressful could not be confirmed. However, it needs to be stated again that we do not know the gender, income or residential situation of the actual senders of Tweets in our database; we are relying on the assumption that general socio-economic profiles of municipalities align with those of Twitter users in those municipalities which is however less accurate. If the ambition is to focus on just these socio-economic factors, we can imagine that a research approach targeting individuals is more appropriate.

Geographical factors

Besides the temporal dimension, our main interest is in the spatial variation in Twitter sentiment. Whereas earlier findings by Cao et al. (2018) suggested that farmland areas accommodate the most negative sentiment on Twitter, it turns out that the most rural areas in the Netherlands witnessed the most positive twitter sentiment during the pandemic year studied here (Models 3-5). We found that the more urban a municipality is, the higher the negativity rate. This confirms that inhabitants of Dutch rural areas experience a higher level of well-being, especially during COVID-19 restrictions, compared to Dutch urban dwellers (Van Leeuwen and Bourdeau-Lepage, 2020). Considering the dummy variables for the level of urbanisation in Model 4, the difference between the most urban places and rural areas is significant. Municipalities belonging to this very strongly urbanised category predominantly include the largest cities in the Randstad, their immediate suburbs and most medium-sized cities in the Randstad, but also cities elsewhere like Groningen, Tilburg, Eindhoven and Maastricht. Note that this relation between urbanisation and negativity rate holds, also when controlling for different socio-economic profiles of municipalities (Model 5).

5. Conclusion

This study utilized Twitter data to analyse the variation of sentiment over space and time in relation to the COVID-19 measures, the infection rates, socio-economic profiles of places and the level of urbanisation. The data was obtained from a sentiment analysis system named CITYSENT, a tool that autonomously classifies the sentiment of a tweet into 'positive', 'negative' or 'neutral', which was subsequently used to determine a negativity rate per place and period.

These periods were determined by different pandemic control policy regimes during a year. We established that twitter sentiment during the first periods, characterised by a relatively 'normal' summer (period 1; 11 June 2020 - 27 September 2020) and the subsequent period when infections were increasing and restrictions returning (period 2; 28 September 2020 - 13 December 2020), was significantly less positive than the last three of our five periods. Surprisingly, the period containing the strict lockdown and the curfew was not significantly more negative than open periods. Several explanations come to mind. First, these results might reveal that transition periods, i.e., times in between open periods with relatively many freedoms and strict periods with drastic measures, provoke the most negative sentiments to people during the COVID-19 pandemic. Possibly, the weeks in which infection numbers are rising again and with more returning restrictions are the most distressing because of the unexpectedness and the contrast with previous (open) weeks. Especially days on which new COVID-19 restrictions were being announced (or extended) on national press conferences show explicit spikes in negativity on Twitter, rather than days on which the new restrictions were firstly into effect. Second, the holidays in December had a damping effect on the negative sentiment in period 3. However, the most negative day of the entire study period is the day on which the second lockdown was officially announced (14th of December 2020), marking the start of period 3. A third explanation for this temporal pattern is that vaccines were not yet registered and available in the first two periods, and the 'light at the end of the tunnel' that vaccines have said to provide may have led to generally more positive levels of twitter sentiment in later stages.

At first sight, it seems impossible to spot any patterns in maps of the spatial variation in twitter sentiment, but statistically we found a clear association as higher urbanisation levels are associated with higher negativity rates. Urban residents struggled more during the COVID-19 outbreak compared to rural residents. While our explanatory factor was a general measure of the level of urbanisation, it must be noted that it correlates with many more specific factors, like the availability of urban green and distance to national parks, the presence of all sorts of amenities and services and higher real estate prices. Further research could delve into the question what it is precisely that created a more negative mood among urban dwellers. It could be for instance higher levels of boredom and less physical exercise (Van Leeuwen & Bourdeau-Lepage, 2020). But one can also imagine that paying a premium for living in highly urbanised places while not being able to enjoy many of their agglomeration benefits, e.g. in terms of amenities and services, affects one's mood negatively. The same applies for having less personal space, inside and outside the home.

The question is whether this more negative mood among urban dwellers is a temporary difference or a kickstart to a more structural re-evaluation of urban and rural living. As Florida et al. (2021) recently stated, much depends on the longevity of the pandemic.

Glaeser (2021) draws the history lesson that pandemics generally do not affect the attractiveness of cities in the longer term, unless they come coupled with economic and political shocks that engender urban decline. He states that telecommuting and online shopping may pose such post-pandemic economic risks for cities, and to this we would like to add that housing affordability has now come under pressure also for middle incomes. At the same time, while highly localised agglomeration advantages remain (e.g. in terms of knowledge transfers), the idea is that many sorts of agglomeration benefits do spread over a larger territory, especially in small, polycentric and highly connected countries like the Netherlands. Here, agglomeration benefits are increasingly a zonal rather than nodal phenomenon as captured by the term 'agglomeration externality field' (Burger and Meijers, 2016). Locating somewhere in this field allows access to most agglomeration benefits, while avoiding some of the agglomeration costs that became more evident during the pandemic and that are likely to have affected the mood of the population. In a spatial context like the Netherlands, the mood of a population could more easily translate into changes in spatial location behaviour. Also Glaeser et al. (2016) draw attention to how such a spatial context may change the perspective on urbanisation. However, much of the evidence for a strengthened process of counterurbanisation is still anecdotical, at least in the Netherlands, and requires further empirical research.

While Twitter sentiment describes the general short-term mood of a population, with Twitter users fully fitting the description of 'social sensors on the ground' (Günnemann & Pfeffer, 2015), we need to mention some drawbacks. Most notably, is the lack of information on Twitter users. Attempts to derive gender and age from profile pictures were not yet satisfactory for our purposes. Also, the fine-grained geographical detail that we sought for was in some cases problematic, as some municipalities only had a few hundred tweets over the entire study period, a problem that became more apparent when municipalities were analysed per period, and forcing us to exclude too small samples from the analysis. Although results of this study were compared with findings from earlier Twitter data using studies, one should keep in mind that outcomes of such studies strongly depend on the geographic detail, the sentiment classification and measuring methods, which vary per study and can hence explain different findings. Unfortunately, our method disregarded emoticons used in tweets, but this would be an interesting extension. Finally, a limitation of this study was that only data during the COVID-19 pandemic was available, which made it impossible to compare it with pre-pandemic time periods. To better understand the implications of the COVID-19 related results, future studies could compare the Twitter sentiment used in this study with pre- or hopefully post-pandemic data.

References

Angiani, G., Ferrari, L., Fontanini, T., Fornacciari, P., Iotti, E., Magliani, F., & Manicardi, S. (2016). A Comparison between Preprocessing Techniques for Sentiment Analysis in Twitter. *KDWeb*, 1-10

Bandarin, F., Ciciotti, E., Cremaschi, M., Madera, G., Perulli, P., & Shendrikova, D. (2021). After covid-19: A survey on the prospects for cities. City, Culture and Society, 25 doi:10.1016/j.ccs.2021.100400

Blank, G., & Lutz, C. (2017). Representativeness of social media in great britain: investigating Facebook, Linkedin, Twitter, Pinterest, Google+, and Instagram. *American Behavioral Scientist*, *61*(7), 741-756.

Boon-Itt, S., & Skunkan, Y. (2020). Public perception of the COVID-19 pandemic on Twitter: sentiment analysis and topic modeling study. *JMIR Public Health and Surveillance*, *6*(4), doi: 10.2196/21978

Boterman, W. (2020). Urban-rural polarisation in times of the corona outbreak? the early demographic and geographic patterns of the SARS-CoV-2 epidemic in the netherlands. *Tijdschrift Voor Economische En Sociale Geografie, 111*(3), 513-529.

Burger, M. & Meijers, E. (2016) Agglomerations and the rise of urban network externalities, Papers in Regional Science, 95(1): 5 – 15.

Cao, X., MacNaughton, P., Deng, Z., Yin, J., Zhang, X., & Allen, J. G. (2018). Using Twitter to better understand the spatiotemporal patterns of public sentiment: a case study in Massachusetts, USA. *International journal of environmental research and public health*, *15*(2), 250-270.

Cheng, Z., Caverlee, J., & Lee, K. (2010). You are where you tweet: a content-based approach to geo-locating twitter users. *Proceedings of the 19th ACM international conference on Information and knowledge management,* 759-768.

Curini, L., Iacus, S., & Canova, L. (2015). Measuring idiosyncratic happiness through the analysis of Twitter: An application to the Italian case. *Social Indicators Research*, *121*(2), 525-542.

Denham, T. (2021). The limits of telecommuting: Policy challenges of counterurbanisation as a pandemic response. Geographical Research, doi:10.1111/1745-5871.12493

De Vos, D., Meijers, E., & van Ham, M. (2018). Working from home and the willingness to accept a longer commute. Annals of Regional Science, 61(2), 375-398.

Diener E. (2009) *Subjective Well-Being. In: Diener E. (eds) The Science of Well-Being. Social Indicators Research Series*, vol 37. Dordrecht: Springer

Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A., & Danforth, C. M. (2011). Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PloS one*, *6*(12).

Durahim, A. O., & Coşkun, M. (2015). # iamhappybecause: Gross National Happiness through Twitter analysis and big data. *Technological Forecasting and Social Change*, *99*, 92-105.

Florida R, Rodríguez-Pose A, Storper M. (2021) Cities in a post-COVID world. Urban Studies. June 2021. doi:10.1177/00420980211018072

Giachanou, A., & Crestani, F. (2016). Like it or not: A survey of twitter sentiment analysis methods. *ACM Computing Surveys (CSUR)*, *49*(2), 1-41.

Glaeser, E., Ponzetto, G. & Zou, Y. (2016) Urban networks: Connecting markets, people, and ideas, *Papers in Regional Science*, 95 (1): 17 – 59.

Go, A., Bhayani, R., & Huang, L. (2009). Twitter sentiment classification using distant supervision. *CS224N project report, Stanford*, *1*(12), 2009.

Goslinga, D., Kalisvaart, R., Pardoel, A., Paulsen, C. & De Wolf, S. (2020) CITYSENT - Sentiment analysis of the Netherlands and Flanders. Delft: Delft University of Technology.

Günnemann, N., & Pfeffer, J. (2015). Finding non-redundant multi-word events on twitter. In *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 520-525.

IHME (Institute for Health Metrics and Evaluation) (2021) https://covid19.healthdata.org/global (accessed November 1, 2021)

Kowal, M., Coll-Martín, T., Ikizer, G., Rasmussen, J., Eichel, K., Studzińska, A., & Ahmed, O. (2020). Who is the most stressed during the COVID-19 pandemic? Data from 26 countries and areas. *Applied Psychology: Health and Well-Being*, *12*(4), 946-966.

Kuebart, A., & Stabler, M. (2020). Infectious diseases as socio-spatial processes: The COVID-19 outbreak in Germany. Tijdschrift Voor Economische En Sociale Geografie, 111(3), 482-496.

Limcaoco, R. S. G., Mateos, E. M., Fernandez, J. M., & Roncero, C. (2020). Anxiety, worry and perceived stress in the world due to the COVID-19 pandemic, March 2020. Preliminary results. *MedRxiv*. 1-21

Lim, K. H., Lee, K. E., Kendal, D., Rashidi, L., Naghizade, E., Winter, S., & Vasardani, M. (2018). The grass is greener on the other side: Understanding the effects of green spaces on Twitter user sentiments. In *Companion Proceedings of the The Web Conference 2018*, 275-282.

Manguri, K. H., Ramadhan, R. N., & Amin, P. R. M. (2020). Twitter sentiment analysis on worldwide COVID-19 outbreaks. *Kurdistan Journal of Applied Research*, 54-65.

Mitchell, L., Frank, M. R., Harris, K. D., Dodds, P. S., & Danforth, C. M. (2013). The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PloS ONE*, *8*(5).

Nguyen, Q. C., Kath, S., Meng, H. W., Li, D., Smith, K. R., VanDerslice, J. A., ... & Li, F. (2016). Leveraging geotagged Twitter data to examine neighborhood happiness, diet, and physical activity. *Applied Geography*, *73*, 77-88.

Poortinga, W., Bird, N., Hallingberg, B., Phillips, R., & Williams, D. (2021). The role of perceived public and private green space in subjective health and wellbeing during and after the first peak of the COVID-19 outbreak. *Landscape and Urban Planning*, *211*, 104092.

Rijksinstituut voor Volksgezondheid en Milieu (2021, June 27). *Ontwikkeling COVID-19 in grafieken*. RIVM. Retrieved from: https://www.rivm.nl/coronavirus-covid-19/grafieken

Schwartz, H., Eichstaedt, J., Kern, M., Dziurzynski, L., Lucas, R., Agrawal, M., ... & Ungar, L. (2013). Characterizing geographic variation in well-being using tweets. In *Proceedings of the International AAAI Conference on Web and Social Media*, *7*(1), 583-587

Sloan, L., & Morgan, J. (2015). Who tweets with their location? Understanding the relationship between demographic characteristics and the use of geoservices and geotagging on Twitter. *PloS one*, *10*(11).

Sloan, L., Morgan, J., Housley, W., Williams, M., Edwards, A., Burnap, P., & Rana, O. (2013). Knowing the tweeters: Deriving sociologically relevant demographics from Twitter. *Sociological research online*, *18*(3), 74-84.

Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational linguistics*, *37*(2), 267-307.

Taquet, M., Quoidbach, J., Fried, E. I., & Goodwin, G. M. (2021). Mood homeostasis before and during the coronavirus disease 2019 (COVID-19) lockdown among students in the Netherlands. *JAMA psychiatry*, *78*(1), 110-112.

Terry, P. C., Parsons-Smith, R. L., & Terry, V. R. (2020). Mood responses associated with COVID-19 restrictions. *Frontiers in Psychology, 11* doi:10.3389/fpsyg.2020.589598

Van der Veer, N., Boekee, S., & Hoekstra, H. (2020). *Nationale social media onderzoek 2020*. Newcom Research & Consultancy.

Van Leeuwen, E., & Bourdeau-Lepage, L. (2020). Spatial differences and the impact of the lockdown on well-being in the Netherlands. *SSRN, 3597707*, 1-7.

Van Tilburg, T. G., Steinmetz, S., Stolte, E., van der Roest, H., & de Vries, D. H. (2020). Loneliness and mental health during the COVID-19 pandemic: A study among Dutch older adults. *The Journals of Gerontology: Series B*.

Volkova, S., & Bachrach, Y. (2015). On predicting sociodemographic traits and emotions from communications in social networks and their implications to online self-disclosure. *Cyberpsychology, Behavior, and Social Networking*, *18*(12), 726-736.

Volkova, S., Han, K., & Corley, C. (2016). Using social media to measure student wellbeing: a large-scale study of emotional response in academic discourse. *International Conference on Social Informatics,* 510-526.

Wang, S., Schraagen, M., Sang, E. T. K., & Dastani, M. (2020). Dutch general public reaction on governmental covid-19 measures and announcements in twitter data. *ArXiv,* 1-25.

Yang, C., & Srinivasan, P. (2016). Life satisfaction and the pursuit of happiness on twitter. *PloS one*, *11*(3), 1-30

You, L., & Tunçer, B. (2016). Exploring public sentiments for livable places based on a crowd-calibrated sentiment analysis mechanism. *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 693-700.

Zhou, J., Yang, S., Xiao, C., & Chen, F. (2020). Examination of community sentiment dynamics due to Covid-19 pandemic: A case study from Australia. *ArXiv,* 1-15

[i] This paper was written and submitted for publication to a scientific journal in early 2022. While positively and constructively reviewed, we did not manage to submit the revised version within the required time due to personal circumstances. Working on the paper again in summer 2023, we realized that it needed updating to include more recent phases of the pandemic, including the post pandemic phase, in order to be timely again. However, because of time constraints and our shared feeling that we were somewhat done with COVID19 (our own little negativity pandemic), we decided to publish this slightly revised paper as a working paper and hope others can build on this work.