

“Social media is
transforming how
to study society...”

VOX DIGITAS

Jamie Bartlett
Carl Miller
Jeremy Reffin
David Weir
Simon Wibberley

Demos is Britain's leading cross-party think tank. We produce original research, publish innovative thinkers and host thought-provoking events. We have spent 20 years at the centre of the policy debate, with an overarching mission to bring politics closer to people.

Demos is now exploring some of the most persistent frictions within modern politics, especially in those areas where there is a significant gap between the intuitions of the ordinary voter and political leaders. Can a liberal politics also be a popular politics? How can policy address widespread anxieties over social issues such as welfare, diversity and family life? How can a dynamic and open economy also produce good jobs, empower consumers and connect companies to the communities in which they operate?

Our worldview is reflected in the methods we employ: we recognise that the public often have insights that the experts do not. We pride ourselves in working together with the people who are the focus of our research. Alongside quantitative research, Demos pioneers new forms of deliberative work, from citizens' juries and ethnography to social media analysis.

Demos is an independent, educational charity. In keeping with our mission, all our work is available to download for free under an open access licence and all our funders are listed in our yearly accounts.

Find out more at www.demos.co.uk.

First published in 2014
© Demos. Some rights reserved
Magdalen House, 136 Tooley Street
London, SE1 2TU, UK

ISBN 978 1 909037 63 2
Series design by Modern Activity
Typeset by Modern Activity

Set in Gotham Rounded
and Baskerville 10



Mixed Sources

Product group from well-managed
forests, controlled sources and
recycled wood or fiber

www.fsc.org Cert no. CU-COC-804101
© 1996 Forest Stewardship Council

VOX DIGITAS

Jamie Bartlett

Carl Miller

Jeremy Reffin

David Weir

Simon Wibberley

Open access. Some rights reserved.

As the publisher of this work, Demos wants to encourage the circulation of our work as widely as possible while retaining the copyright. We therefore have an open access policy which enables anyone to access our content online without charge.

Anyone can download, save, perform or distribute this work in any format, including translation, without written permission. This is subject to the terms of the Demos licence found at the back of this publication. Its main conditions are:

- Demos and the author(s) are credited
- This summary and the address *www.demos.co.uk* are displayed
- The text is not altered and is used in full
- The work is not resold
- A copy of the work or link to its use online is sent to Demos

You are welcome to ask for permission to use this work for purposes other than those covered by the licence. Demos gratefully acknowledges the work of Creative Commons in inspiring our approach to copyright. To find out more go to *www.creativecommons.org*



Contents

Executive summary	9
1 Social media: a new political theatre for Europe	25
2 Research design and methodology	33
3 How do people use Twitter to talk about Europe?	39
4 Case studies of real world events	49
5 Digital observation	73
Annex: methodology	85
Notes	103
References	121

Executive summary

Trust, engagement and support for the European Union (EU) is on a downward path. It suffers from a democratic deficit: the perception and reality that there is a large distance in understanding and reality between the governors and the governed. Democratic representation must mean more than electoral success – it must also include ‘responsive listening’: listening to people, understanding their fears, priorities and aspirations, and reacting to what is heard and learnt. Listening is a vital link between people and institutions, which underlies the practical, everyday task of representing people, and discharging a mandate on their behalf.

There are now more opportunities to listen than ever before. Over the last decade European citizens have gained a digital voice. Close to 350 million people in Europe currently use social networking sites – three in four EU citizens. More of us sign into a social media platform at least once a day than voted in the last European elections. Facebook has 232 million users across the EU and 16 per cent of European internet users have a Twitter account. EU citizens have transferred many aspects of their lives onto these social media platforms, including politics and activism. They use social media to discuss news stories, join political movements, organize new political movements and broadly discuss and dissect those public issues that matter to them, across boundaries and at essentially no cost. Taken together, social media represent a new digital commons, central places of assembly and interchange where people join their social and political lives to those around them. It is a new, focal theatre for Europe’s daily political life.

We have never before had access to the millions of voices that together form society’s constant political debate, nor the

possibility of understanding them. Capturing and understanding these citizen voices potentially offers a new way of listening to people, a transformative opportunity to understand what they think, and a crucial opportunity to close the democratic deficit.

However, making sense of digital voices requires a new kind of research.¹ Traditional attitudinal research relies on tried and trusted methods and techniques: the focus group, the interview, the national poll. But turning the cacophony of sometimes millions of social media conversations into meaning and insight requires the use of powerful new technologies that are capable of automatically collecting, storing, analysing and visualising information. This throws up questions of trust and rigour at every stage of the research cycle: the role of technology and automation, how to sample the data, how to make sense from the noise, how to interpret the information appropriately and use it, and how to do this all ethically.

This paper examines the potential of listening to digital voices on Twitter, and how far it might be an opportunity to close the democratic deficit. It looks at how European citizens use Twitter to discuss issues related to the EU and how their digital attitudes and views about the EU are evolving in response to political and economic crises faced by the EU. We ask whether social media analysis can provide a new way for the EU's leaders to apprehend, respond and thereby represent its citizens. It addresses the many formidable challenges that this new method faces: how far it can be trusted, when it can be used, the value such use could bring, and how its use can be publicly acceptable and ethical.

Listening to digital voices

The potential of social media as a source of attitudinal insight was tested using the practical case of the EU. The period between March and June 2013 was an extremely difficult time for the EU and related institutions. There were a series of economic bailouts, landmark and controversial European Court of Human Rights rulings, and the opening of the European Commission. We investigated two key themes:

- What kind of digital voices exist? How do EU citizens use Twitter to discuss issues related to the EU? What kind of data does Twitter therefore produce?
- How do we listen to these voices? To what extent can we produce meaningful insight about EU citizens' attitudes by listening to Twitter? How does this relate to other kinds of attitudes, and other ways of researching them?

Over this period around 3.26 million publicly available tweets were collected directly and automatically from Twitter in English, French and German, which contained a keyword considered relevant to one of six themes selected. These represent the many identities the EU has for the people who talk about it: an institution that drafts laws and enacts and enforces them, a collection of institutions which define and shape their economic lives, and a body of politicians and civil servants.

The volume of data collected was too large to be manually analysed or understood in its totality. We therefore trialled a number of different methods – automated and manual, some highly technological and others straightforward – to understand it. These included:

- *data overview*: examining the general characteristics of the Twitter data for each data stream such as volume of hash tags, retweets, linkshares, user mentions and traffic analysis
- *testing natural language processing*: which allows researchers to build algorithms that detect patterns in language use that can be used to undertake automatic meaning-based analysis of large data sets; these were built and applied in different contexts to see where it worked, and where it did not; these algorithms are called 'classifiers' – the research team built over 70 such classifiers, and tested how well they performed against human analyst decisions
- *manual and qualitative analysis*: using techniques from content analysis and quantitative sociology to allow analysts to manually discern meaning from tweets
- *five case studies*: examining how Twitter users responded to events as they happened in the real world, and whether they could be reliably researched

It was unclear at the outset what combination, and in what context, of these kinds of analysis would be effective or reliable. Different frameworks of use were therefore flexibly and iteratively applied throughout the course of the project.

Europe's new digital voice

This is a summary of what we found:

There are millions of digital voices talking about EU-related themes in real time; it is a new venue for politics

In four months, we collected 1.91 million tweets in English across the six English data streams; 1.04 million across the six French data streams; and 328,800 across the six German data streams. We considered 1.45 million tweets across all three languages to be 'relevant' to one of the six EU-related themes. This included almost 400,000 tweets about the euro currency in English, and 430,000 about the EU.²

These voices are event-driven and reactive, not steady and general

Most of the data collected are of people reacting to events, such as a major speech, ruling, or news story. These offline events provoke groundswells of online reaction that shadow events that have occurred offline – each a collectively authored digital annotation of the event, containing questions, interpretations, condemnations, jokes, rumours and insults. These 'twitcidents' will become a routine aftermath, a usual way that society reacts to and annotates the events it experiences.

These voices share information about events and express attitudes about them

Tweets were often used to keep up with recent developments in a rapidly changing world. Over half of every data set was tweets that shared a link to a site beyond Twitter, primarily to media stories, often containing no additional comment by the tweeter themselves. Where attitudes were expressed, it was often in the form of non-neutral reportage of a specific event.

Making sense of the noise: digital observation

These voices cannot be listened to in conventional ways. Twitter data sets are ‘social big data’. Conventional methods to gather and understand attitudes – polls, surveys and interviews – are overwhelmed by how much or how quickly data are created. Twitter offers a novel way of understanding citizens’ attitudes and reactions to events as they unfold, in a way that can be extremely powerful and useful for academics, researchers, advocacy groups, policy makers and others. Twitter is a new type of reactive, short-form expression, produced in large volume, and above all driven by events.

Current ways of researching society cannot handle these kinds of data in the volumes that are now produced. While there is a burgeoning industry in applying new computational techniques to try to analyse social media data, it can be misleading, and often hides sociologically invalid modes of collection and analysis. This is most important for the most popular way of analysing social media content ‘sentiment analysis’, which breaks conversations into ‘positive’, ‘negative’ and ‘neutral’ categories. This kind of analysis often uses natural language processing (NLP) in ways that our pilot found unlikely to be successful – generic, standardised, operating over a long period of time, and not related or trained to a particular event or conversation.

We found that it is possible to create new ways of combining new technology and traditional methodologies to understand the groundswells of digital voices that rise in reaction to important events as they occur. Through trial and error and case studies, we developed an approach to analysing these data sets, which we call ‘digital observation’. This includes:

- collecting tweets directly from Twitter on a given theme as they are posted in real time
- identifying groundswells of tweeted reaction when they occur on a particular theme and identifying the event(s) that are driving it. Our case studies and classifier tests revealed that people do not in general express generic sentiment on Twitter

about the EU, instead, Twitter was found to be fundamentally a reactive medium; a tweet is overwhelmingly a reaction to an event that the tweeter has otherwise encountered – either online or offline, whether through reading mainstream media or being told by a friend. Therefore, it is best used as a way of gaining insight into how people respond to events, rather than as a continuous ‘poll’ of opinion. The closest analogy to the value of insight from Twitter is perhaps not the population level opinion poll, but rather the noise of a throng of energised citizenry talking about a particular event

- using automatic NLP to build algorithmic classifiers, which can filter out tweets that are irrelevant to the theme in question
- flexibly and reactively building bespoke technology around these specific events to listen to the digital voices – what they are saying and the attitudes, hopes, fears and priorities that they carry with them at scale and speed
- situating these attitudes within the background of the events that were occurring, the media reportage that covered them, and the public discussions that were being carried out

What are digital voices saying?

Using this method, we found a number of specific features about tweets relating to the EU:

The silent majority ‘reaction’ phenomenon

While the general consensus in the UK is that the population is broadly hostile to the European Court of Human Rights (ECHR), it is of note that the response following the Cameron suggestion about leaving the ECHR led to a groundswell of hostile criticism. Even when it came to a very unpopular ruling – preventing the UK from deporting Abu Qatada – most Twitter users rallied around the *principle* of the ECHR (of 1,344 attitudinal tweets about this decision, 1,181 were classified as pro-ECHR and 163 negative).

Commission events are a good opportunity to gauge general views

There is clearly a significant surge in activity surrounding major

European events, such as summits – they are news stories in themselves. Rather than being based on a single news story (such as the other data sets) there was a significant number of tweets about the summit, which was an occasion for people to bring their own, related topics of concern to the table.

Variation across countries

By listening to how people on Twitter reacted to certain events rather than as a continuous whole, they tell a story of users responding to each EU-related case separately and on broad national distinctions. French tweeters thought the European Central Bank was ‘strangling’ Cyprus, while German tweeters continued to worry about Germany’s place in the Eurozone. Both British and French tweeters broadly applauded the ECHR over their own national governments, but French tweeters did not like Barroso’s incendiary admonishments of Hollande and France.

Discussion: Twitter as a source of attitudinal data

Digital observation has considerable strengths and weaknesses compared with conventional approaches of studying attitudes. It is able to leverage more data about people than ever before, with hardly any delay and at very little cost. On the other hand, it uses new, unfamiliar technologies to measure new digital worlds, all of which are not well understood, producing event-specific, ungeneralisable insights that are very different from what has up to this point been produced by attitudinal research in the social sciences. Based on our research, we consider the following strengths and weaknesses to be most significant.

Strengths

Very large data sets available

Twitter data sets are ‘social big data’. The size of the data set gathered even for this pilot is far larger than comparative data sets gathered through conventional polling, interviewing and surveying techniques. Digital observation radically widens the number of voices that can routinely be listened to.

Real-time insight

Relevant tweets are collected almost immediately after they are posted. Digital observation, using automated technologies, draws meaning from these data very quickly after collection. It is therefore possible to understand attitudes about an event as the event happens, and as the public debate evolves. This is perhaps the most important distinction between digital observation and other ways of researching attitudes. Discerning real-time attitudes is a valuable power for institutions to have. It allows them to be agile, and react quickly to groundswells of anger, support or criticism quickly enough to influence the underlying developments and events that drive these attitudes.

'In conversation':

listening rather than asking removes observation bias

A well-known weakness in most attitudinal research is that data are collected in 'non-real world' settings. Most ways of gathering attitudes require a researcher intervening in someone's life – asking them questions, and recording what they say. This introduces 'observation effects', which change the attitudes expressed and views offered in a number of ways.³ Digital observation avoids these unwanted distortions by listening to digital voices as they rise, naturally, on social media platforms.

Cheap

Attitudinal research is often expensive. It is expensive to employ interviewers and to manage and incentivise panels of participants, to mail surveys to thousands of people and to hire rooms, technology and people to conduct focus groups. Digital observation is very economic in comparison. Acquiring tweets (in certain contexts and quantities) is free and the technology, once in place, can be trained and purposed in a matter of minutes. This lowers the threshold for attitudinal research – many more organisations will be able to listen more often to more conversation that they care about.

Weaknesses

There is no accepted 'good practice' for digital observation

Established ways of researching attitudes have long histories of use. This experience has consolidated into a body of good practice – dos and don'ts – which, when followed, ensures the quality of the research. Digital observation does not have a long history of use, or an established collective memory of what works and what does not. It uses new technologies in new ways that are unfamiliar with the social sciences, often with new and important implications for research.

The performance of the technology varies considerably

The technology sometimes performed very successfully, and at other times very poorly. In the research, the best performing classifiers were almost always correct, and the worst performing classifiers performed no better than chance. The performance of classifiers depends on the context of the task. We found that generic, long-term classifiers performed inaccurately. Language use – the kinds of words used and the meanings these words have – changes quickly on Twitter. Language is specific to a particular conversation at a particular time. Automated algorithms struggled to find generic meaning accurately independent of a particular event or discussion, and became drastically less accurate over a long period of time. However, bespoke short-term classifiers did well and proved to be able to reliably discern meaning, when trained on a specific event, at a specific time, and in a way that reflects the data. Classifiers performed best when making distinctions that reflected the data at a particular point. There are also other difficulties for classifiers. Non-literal language use, such as sarcasm, pastiche, slang and spoofs, are found to be common on social media. The 'real' rather than ironic meaning of these kinds of uses of language are inherently contextual and difficult to deduce via shallow automated analyses.

Sampling: the tweets that are gathered may systemically differ from those that are not

Data are acquired through Twitter by being matched to keywords. The pilots showed that these keywords can produce different kinds of problems – sometimes they are over-inclusive (and collect tweets on other, irrelevant topics), and sometimes they are under-inclusive (and miss relevant tweets). In both these ways, keyword matching is inherently prone to systemic bias – so data collected, and therefore the conclusions drawn, are affected in a non-random way by the search terms employed.

Insights from digital observation can be difficult to generalise

The power of much attitudinal research is that it creates representative data sets that allow for generalisations beyond the group that was actually asked – to age group, area, the country or even the world. Making these generalisations when using Twitter as a source of attitudinal data is difficult because of a problem of representativeness. Twitter users do not demographically represent wider populations: data are collected based on conversations rather than demographic details of a participant. Moreover, collected tweets often do not represent Twitter users. Tweets, in general, are produced by a small number of high-volume ‘power-users’. Compounding this problem, ‘power-users’ are sometimes automated, ‘bot’, fake, official or campaign accounts. Moreover, Twitter is a new social space, allowing the growth of a number of digital cultures and sub-cultures with distinct norms, ways of transacting and speaking and also a new communications medium whose format powerfully influences what is said and meant. The pace with which this context evolves and changes makes the meaning of tweets often unclear or ambiguous.

Recommendations and ways forward

Twitter has become an extremely significant venue for public debate and discussion. Increasingly, it is an important way for citizens to express their attitudes on a range of subjects, including the European project. We recommend that

representative organisations examine ways to listen and respond to these digital voices:

Investigate digital observatories

Organisations, especially representative institutions, now have the opportunity to listen cheaply to attitudes expressed on Twitter that matter to them. They should consider establishing digital observatories that are able to identify, collect and listen to digital voices, and establish ways for them be reflected appropriately in how the organisation behaves, the decisions it makes and the priorities it has. Digital observatories, constantly producing real-time information on how people are receiving and talking about events that are happening, could be transformative in demonstrating how organisations relate to wider societies. Just because it is possible to collect social media data does not mean it should be done. Digital observatories should be predicated on public understanding and openness about how they work; and conducted according to strict ethical principles for the collection analysis and use of data. This type of research should not replace existing methods of research, but supplement it.

The EU must adopt a leadership role on how to listen to citizens ethically and robustly

Increasingly, politics is moving online, enabling people to express opinions, politicians to mobilise voters, and anyone to form parties and movements. This opens new roles and opportunities for research to be powerful and useful: to rework communication campaigns that are misunderstood; to delay or halt policy roll-outs that have unintended and unforeseen consequences; and more broadly to allow those in democratic institutions to perceive, react to and represent views during the time when they most matter – as they are expressed. However, as a new field, this also creates ethical risks and dangers of poor research methods. To be a leader in the democratic governance and representation of an increasingly digital world, the EU must stake out leadership in the ethical and effective exploitation of these new technologies, grasping the opportunities they now offer.

Nine principles for social media research

Our ability to understand Twitter as a source of attitudes is nascent. Many of the tools that can handle large numbers of tweets have come from the computer science departments of academia, and the public relations and advertising industries. Their up-take within the sociological, psychological and anthropological disciplines has been slower, and new technologies have often not been reconciled with the values and principles of conventional attitudinal research.

It is necessary to arrive at a new discipline capable of turning social media into social meaning. This pilot demonstrated the strength of combining human and technological analysis, built around a specific event as it happens.

For this to be realised, we recommend the following nine principles for undertaking high quality social media research. They are designed for researchers, advocacy groups and others interested in understanding society, as a set of techniques approaches and methods for how to make the best use of these techniques, and turn the potential of listening to the digital voices into something useful and valuable:

1 Beware the numbers game and 'sentiment analysis software' – this will not always deliver the best results and can be misleading

Size is not everything. While there is a burgeoning industry in analysing social media, very large amounts of data often hide sociologically invalid modes of collection and analysis. This is most important for the most popular way of analysing social media content 'sentiment analysis', which breaks conversations into 'positive', 'negative' and 'neutral' categories. This kind of analysis often uses NLP in ways that our pilot found unlikely to be successful – generic, standardised, operating over a long period of time and not related to a particular event or conversation.

2 Digital observation can complement existing polling data, but not replace them

It is therefore necessary to use a new approach to 'attitudes' that reacts to events in real time. Traditional, representative

polling data still remain an extremely powerful way to ascertain attitudinal data, especially across large populations. It is based on tried and tested methods of randomised sampling and questionnaire design. Twitter data are of a different nature – dynamic, unstructured and event-driven. They should be viewed as a complement to, rather than replacement for, traditional polling.

3 Look for ways to mix qualitative and quantitative, automated and manual methods

Automated techniques are only able to classify social media data into one of a small number of categories at a certain (limited) level of accuracy for each message. They are a good first way to tackle scales of data that would otherwise be overwhelming. Manual analysis is therefore almost always a useful and important component; in this report it is used to look more closely at a small number of randomly selected pieces of data drawn from a number of these categories. In scenarios when a deeper and subtler view of the social media data is required, the random selection of social media information can be drawn from a data pool, and sorted manually by an analyst into different categories of meaning.

4 Involve human analyst and subject matter expertise at every step

It is vital that attempts to collect and analyse attitudinal ‘big data’ is guided by an understanding of what is to be studied: how people express themselves, the languages that are used, the contexts – social and political – that attitudes are expressed in, and the issues that they are expressed about. Analysts who understand the issues and controversies that surround the EU are therefore vital in order to contextualise and explain the attitudes that are found on Twitter, and to help build the methods that are used to find and collect these attitudes.

5 Beyond the ‘black box’ – new big data technologies must be presented in a way that non-specialists can understand

Non-technical specialists are often the end-users of the research, and it is vital that the technology, however sophisticated,

is explained in a way that clearly lays out how it was used and what the implications of its use are. This means clarity and detail must be provided about how the search terms were constructed and why, what type of data access terms were used, how well the classifier performed against a human analyst, and what the likely biases in the data were.

6 Use new technologies in contexts where they work.

NLP classifiers should be bespoke, not generic and driven by the data rather than predetermined

Overall, NLP classifiers seem to perform best when they are bespoke and event-driven rather than generic. When categories to sort and organise data are applied *a priori*, there is a danger that they reflect the preconceptions of the analyst rather than the evidence. It is important that classifiers should be constructed to organise data along lines that reflect the data rather than the researcher's expectations. This is consistent with a well-known sociological method called grounded theory.⁴

7 New roving, changeable sampling techniques

The collection of systemically biased data from Twitter is far from easy. The search terms that are used are vulnerable to Twitter's viral, short-term surging variations in the way that language is used to describe any particular topic, so keyword searches are liable to result in bias and/or incomplete data sets. Therefore, development is needed to improve ways of sampling in a more coherent and repeatable way.

8 From metrics to meaning

Numbers and measurements alone cannot talk for themselves, and do not represent meaningful insight that can be acted on. It is here, in the ability to translate measurements into insight and understanding that can be acted on, that most work is required. Findings from digital observation must be intensively contextualised within broader bodies of work in order to draw out causalities and more general insights.

9 Apply a strict ethical approach at every step

Researching people entails moral hazard. Research can harm the individual participants involved or more broadly the society from which they are drawn. Ethical codes of conduct are used by researchers to minimise these harms, and balance them against the social benefits of the research. In the UK, the standard best practice for research ethics is the ethical framework of the Economic and Social Research Council (ESRC), which is made up of six principles.⁵ It is unclear, however, how these can be applied for the mass collection of social media data. At the time of this writing, no official frameworks on internet research ethics have been adopted at any national or international level.⁶ Social media research of this kind is a new field, and the extent to which (and how) these ethical guidelines apply practically to research taking place on social media is unclear. We consider that the two most important principles to consider for this type of work are whether informed consent is necessary to reuse the Twitter data that we collected, and whether there are any possible harms to participants in republishing their tweets that must be measured, managed and minimised. Researchers must bear these considerations in mind at all times, and not assume that because data are available it is necessarily ethical to access and use them. We therefore suggest that all academic research work that involves collecting social media data relating to individuals should be subject to ethical review boards.

1 Social media: a new political theatre for Europe

A crisis of confidence

Throughout the European continent, there is a profound disaffection with politics and the political system, both towards national governments and the EU. Scepticism and uncertainty about the EU's future has grown. Anti-EU populist parties have garnered attention and momentum across EU member states, and performed well in the 2014 European elections. The future of the EU depends on the response to these critical events, and more broadly on bridging the real and perceived distances that now divide representative institutions, and those they represent.

Representation through formal democratic participation is trapped in a downward spiral. EU elections have consistently failed to attract the number of voters that participate in national elections. Political parties, sitting at the heart of both national and European elections, are highly distrusted almost everywhere. In Germany 73 per cent distrust them, as do 89 per cent of French citizens and 85 per cent of British citizens. Only around 2 per cent of voters in these countries are now members of a mainstream political party.⁷

In the wake of the economic recession and Eurozone crisis, distrust in EU institutions has increased in many countries. Between 1999 and 2009, trust in EU institutions was around 45–50 per cent. Since autumn 2009, trust levels dropped substantially from 48 per cent down to 33 per cent in autumn 2012.⁸ The latest Eurobarometer report shows a significant increase in the percentage of respondents who have a 'negative' image of the EU.⁹ Indeed, in May 2013, Pew Global proposed that the EU was 'the new sick man of Europe'.¹⁰

These figures – already worrying reading – may obscure a significant generational divide. Young people are the least

likely age group to have voted in the 2009 European Parliament election.¹¹ Those who did vote, tended to vote against the EU: 74 per cent of the Dutch 18–24-year-olds who voted, voted against the Constitutional Treaty, and 65 per cent of young Irish voters said ‘no’ to the Lisbon Treaty – both significantly higher than other age groups.¹²

However, research shows young people are interested in politics, but perhaps not the way it is done at present. A recent UK survey of 18-year-olds found that many respondents expressed an interest in political affairs when broadly defined, and many said they were keen to play a more active role in the political process.¹³ Even though young people have a fairly strong aversion to formal politics and professional politicians, they are relatively active in alternative modes of political participation.¹⁴ One venue for this new type of participation is social media.

Social media

The way people live their social lives in Europe is changing radically. While trust, engagement and support for our representative institutions continue to fall, there has been a democratisation in how our society produces, shares and consumes information. The explosion of a new, ‘social’ media – those platforms, internet sites, apps, blogs and forums that allow for user-generated content to be published and shared – have created a new digital commons.¹⁵ Around the world, 1.2 billion people use one of these platforms at least once a month.¹⁶ The most well known are Facebook (the largest, with over a billion users), YouTube and Twitter. They are only the most famous members of a much more linguistically, culturally and functionally diverse family of platforms and communities that span social bookmarking, micromedia, niche networks, video aggregation and social curation.¹⁷

Around three out of four Europeans use at least one social media platform, and 60 per cent of Europeans log into their social accounts every day,¹⁸ 62 per cent use Facebook, and 16 per cent use Twitter. What platforms people use, how often,

when and for what reason, and the value they get from them, differ greatly according to background, where they live, how old they are and how rich they are.¹⁹

Twitter

Twitter – the platform used for this study – is a social media platform that allows users to create accounts and post ‘micro-blogs’ to the site of no more than 140 characters in length. Since it began operating in 2009, its 250 million active users have posted over 170 billion micro-blogs, ‘tweets’. As a platform experiencing extremely rapid growth, the demography – geography, language, age and wealth – of these users is constantly changing. While struggling to keep pace with this changing reality, major studies have found that over 100 languages are regularly used on Twitter. English accounts for around half of all tweets, with other popular languages being Mandarin Chinese, Japanese, Portuguese, Indonesian and Spanish (accounting together for around 40 per cent of tweets).²⁰

In 2012, Twitter ranked as the third most popular social media site in France, the fourth in the UK, and the fifth in Germany.²¹ Approximately 6.6 million people regularly use Twitter in the UK, while in Germany and France the number of active users is estimated to be around 2.4 million and 2.2 million respectively.²² Other reports present higher figures.²³ In the UK, 55 per cent of Twitter users are female and 45 per cent male. In France, users are 40 per cent female and 60 per cent male.²⁴

A new venue for political activism

The role of social media in people’s lives continues to evolve and change. While it was once primarily a social tool for forming friendships and sharing content, it is increasingly a way to consume news, pursue niche interests, form new groups, identities and affiliations, and even coordinate offline activity. People increasingly use social media to engage in politics and political activism.²⁵ It is also beginning to affect formal politics in the way parties form, organise and communicate, the way

in which politicians can get their message out to the electorate, and indeed listen to potential voters.²⁶

'Clicktivism' has emerged as a new, distinct and exclusively online kind of political activism. In 2011, for the first time, people were more likely to contact a politician or a political party online (8 per cent) than offline (7 per cent). In 2011, 9 per cent of people sent an electronic message supporting a political cause, and the same number commented on politics in social media.²⁷ Individuals now increasingly participate in online consultations and voting: within a three-month period 6 per cent of people in Britain, 7 per cent of people in France and 11 per cent of people in Germany took part in an online consultation or vote about civil or political issues.²⁸

New forms of political affiliation based on social media are also growing quickly. According to recent research by Demos, in the UK there are now more unique Twitter users who follow MPs belonging to a party than there are formal party members.²⁹ In France, the Union for a Popular Movement (UMP) has about 205,000 formal members respectively, while President Hollande has 557,741 Twitter followers.

Perhaps more significantly, new kinds of social movements are emerging using social media, and challenging existing parties in a way that was unthinkable a decade ago. The English Defence League in the UK, Beppe Grillo's Movimento 5 Stelle in Italy, and Jobbik in Hungary are very different movements, but they all use social media effectively and are opposed to the EU, which they see as being distant, out of touch, and unrepresentative of national interests. For example, Beppe Grillo used his popular blog, Facebook page, Twitter feed and meet-up group to coordinate a huge number of supporters, becoming the leader of the single largest party at the latest Italian general election.³⁰ Other parties have looked for even more innovative ways to reconnect. The Swedish and German Pirate parties have combined an extensive use of social media with a commitment to values such as openness, dialogue and transparency.³¹

The growth of several anti-elitist, populist parties may at least partly be explained by the combination of these two trends. More people are looking for alternatives to the status quo and by offering new, non-hierarchical ways to communicate and organise, social media presents new avenues for political expression and mobilisation. It facilitates collective action on single issues across borders, with low barriers to entry and very few costs.³² Street-based movements across the continent have also used social media to connect and coordinate disparate groups effectively across the continent.

The Spanish *Los Indignados* movement is an early example of this new potency.³³ As the demonstrations progressed, participants systematically turned to such platforms to discuss relevant issues and improve the movement's coordination. In particular, a series of Twitter hashtags and accounts became a reference point not only in providing tactical information about the protests but also in promoting the movement's message and narrative.³⁴

Listening to the vox digitas

The way these two trends – rising levels of distrust and new ways of coordinating, organising and being part of politics – interact will be crucial for understanding the future of European politics. A whole new space for listening to and engaging with European citizens has opened up.

The legitimacy of democratic governments rests on more than just electoral victory. The challenge continues to secure and sustain representative government day by day. Representivity is vitally sustained by finding ways to understand people's attitudes accurately, and reflecting them in what the institution does. The Harrisburg Pennsylvania opened the era of political polling in 1824; readers preferred Andrew Jackson for president over John Quincy Adams.³⁵ Just over a hundred years later, George Gallup's first national scientific poll opened the way for a method that, evolving from postcards to the telephone to the internet, remains with us today.³⁶

Today, European citizens' opinions are measured by the Eurobarometer, a cross-national longitudinal survey conducted by the European Commission, which has been running since 1973, with all results available on an online database. It is run twice a year and consists of a number of standard questions that are asked in every 'wave' (such as life satisfaction questions) plus a number of thematic one-off or episodic questions. Eurobarometer is powerful and useful, and many of the questions it sets are explicitly written to inform or support particular policy decisions. However, as with any research method, it has limitations. It suffers from considerable lag with events. For example, it cannot tell us about immediate reaction and responses to quickly changing events across the continent – such as how citizens respond to major announcements, events, or crises (such as the Cyprus bailout in early 2013).

Europe now has a digital voice that is loud and passionate, and will continue to increase in importance. Taken together, social media is simply the largest body of information about people and society we have ever had – huge, unmediated and constantly refreshing bodies of behavioural evidence that are, in digital form, inherently amenable to collection and analysis.³⁷ Listening to this digital voice is a new way for European institutions to understand Europe in motion: to gauge public opinion, attitudes and beliefs in a way that can help reconnect people to politics. It can expose relationships, dynamics, processes, tipping points, information on causes and consequences that were previously unseen.³⁸

Turning this potential opportunity into something useful and useable is difficult. Research that produces trustworthy insight – evidence – into attitudes is based on the use of methods that are accepted and widely used by people who practise and use research. The attitudinal research methods used and trusted today to inform important and difficult decisions – from large scientific polling to in-depth qualitative ethnographies – have a long tradition of methodological development behind them. These form defined and codified bodies of good practice that identify the many threats to the accuracy or validity of the research.

Social media research – especially monitoring Twitter – is young. It is composed of a scattering of isolated islands of practice, rather than consolidated bodies of common experience. Private-sector companies, academic institutions and the third sector use it, applying very different research techniques from the computer sciences to ethnography, and with aims ranging from understanding networks of millions to the deep, textured knowledge of an individual. Consequently, there is no accepted or recognised body of best practice capable of satisfying the evidential standards of decision-makers.³⁹

To be powerful and useful, methods to listen to the digital voice need to demonstrate what new and different kinds of insight can be gained through these approaches, and how their strengths and weaknesses compare to other ways of learning about people's opinions and views.

Twitter is often used to share information rather than express opinions

Over half of every stream, and in many cases substantially more, were tweets that shared a link to a site beyond Twitter. A substantial number of these links were to media stories, and a substantial number of tweets linking to media stories contained no additional comment by the tweeter themselves.

2 Research design and methodology

Research aim

This research paper set out to determine the potential of researching Twitter to understand how European citizens' attitudes and views about the EU are evolving in response to the current political and economic crises, and explore methods and approaches that can provide useful, valid insight to these questions. It specifically aims to answer two core questions:

- What kind of digital voices exist? How do EU citizens use Twitter to discuss issues related to the EU? What kind of data does Twitter therefore produce?
- How do we listen to these voices? To what extent can we produce meaningful insight about EU citizens' attitudes by listening to Twitter? How does this relate to other kinds of attitudes, and other ways of researching them?

18 streams on Twitter

Six themes were selected as case studies for these questions: the EU, the euro, Barroso, the Commission, the European Parliament and the ECHR. There are many social media platforms that conceivably host these conversations. We selected Twitter because of the volume and availability of relevant data, and the (relatively) uncontroversial ethical considerations of collecting them.⁴⁰ There are 24 official languages of the EU; for reasons of time and resource we chose the three most commonly used – English, French and German. Together, they are used by about 43 per cent of native speakers in the EU.⁴¹

There was a separate data collection and analysis system for each theme, in each language. This resulted in 18 different, discrete flows of data, which we term 'streams'.

Data collection

It is possible to collect social media data manually in a number of ways – copying, screen grabbing, note-taking and saving web pages. Where large volumes of data are involved, the most appropriate (and sometimes the only possible) method is to collect the data automatically through connection to a platform's application programming interface (API).⁴² The API is a portal that acts as a technical gatekeeper of the data held by the social media platform. APIs allow an external computer system, such as the researcher's, to communicate with, and acquire information from, the social media platform. APIs set rules for this access that often differ in the type of data they allow researchers to access, and the format and quantity they produce it in.⁴³ We collected data via Twitter's API,⁴⁴ which returns tweets with up to 33 pieces of metadata – data about the data – attached, such as location, text and author name.⁴⁵

Only tweets that matched the keywords for each topic and in each language were collected. Choosing keywords is an extremely important component of sampling. Some keywords return very specific samples, others very general ones.⁴⁶ 'Euro' cuts across many different types of issues that are often discussed in high volumes, from football competition to foreign exchange speculation. Others, like 'Barroso', are often used much more specifically in the context of discussing José Manuel Barroso. Generally speaking, the more expansive the cluster of search terms used the more likely it is possible to collect a comprehensive sample, but there will be more irrelevant data included within it.⁴⁷

Data were collected between 5 March and 6 June 2013. During the early stages of the study, a search strategy was developed through a number of formal steps. Both very specific and expansive clusters of keywords were trialled, and the returns were monitored by analysts for relevance to the specific topic. Through a process of incremental improvements, a final cluster of keywords was finalised for each topic, in each language.⁴⁸

The finalised search terms and the numbers that each produced are provided in the annex. In total, approximately

1.91 million tweets were collected in English, 1.04 million in French, and 328,800 in German across the data streams.

Data analysis

The volume of Twitter data collected was too large to be analysed manually or understood in their totality. This sort of natural language, as it occurs on social media, can be analysed automatically at great scale and speed using NLP. A long-established sub-field of artificial intelligence research, NLP combines approaches developed in the fields of computer science, applied mathematics and linguistics. It is increasingly used as an analytical ‘window’ into ‘big’ data sets, such as ours.

A core component of the value of NLP is its ability to create ‘classifiers’, which are trained to place tweets automatically in one of a number of predefined categories of meaning. This process – machine learning – is achieved through mark-up. A machine learning approach that involves semi-supervised learning and active learning significantly reduces the time taken to build classifiers. Carefully selected messages are presented to the analyst via an interface, which the analyst reads, and then decides which of a number of preassigned categories they should belong to. The NLP algorithm looks for statistical correlations between the language used and the meaning expressed to arrive at a series of rules-based criteria. Having learned these associations, the computer applies this criteria to additional (and unseen) tweets and categorises them along the same, inferred, lines as the examples it has been given. The statistical nature of this approach renders it notionally applicable to any language where there is a statistical correlation between language use and meaning. Further details about this method and how we used it are available in the annex.

For each of our data streams, we built a series of separate classifiers with their own discrete jobs. For example, an analyst would build a classifier to recognise whether tweets were relevant or irrelevant. A second classifier would be built to recognise if the relevant data expressed an attitude or not,

and so on. Because of the exploratory nature of the project, while we started with a fixed idea about what classifiers might be built, by the end of the project we would build classifiers based on what any set of conversations appeared to look like.

Interpretation

The outputs of each stream's analytical architecture were subject to four broad modes of interpretation:

- We determined the *type* of data that existed for the stream, for example, looking at the volume, the general use of hashtags, and the popular links shared within the data set.
- We determined how people talk about each of the themes over time, examining the data changes over time, especially as they fluctuated in response to real world events.
- We then attempted to use Twitter to learn about people's attitudes. How this was approached was a key point of evolution over the lifetime of the project. We began by attempting to measure generic attitudes as they related to each of the selected themes. It became increasingly clear that this was not feasible. We then changed our approach to measure attitudes as reactions to events related to the project's themes, by looking in detail at seven real-world case studies that touched on one of our related themes.
- We undertook a constant evaluation of how well the research method itself was working. Of especial importance was the assessment of the technology: whether or not the classifiers performed well, and under what circumstances, by testing them against a human analyst and drawing lessons about where they work, and where they do not work well. It was also important to develop a mix of technological and manual methods to measure attitudes, and to assess how the eventual product of this process, digital observation, relates to wider social science.

Ethics

Research that involves people possibly entails difficult moral questions. In general, it must be conducted in a way that is consistent with a body of fundamental principles – human dignity, respect for individuals and the maximisation of social value, which are codified in documents such as the UN Declaration of Human Rights, the European Convention on Human Rights, and the Declaration of Helsinki. At the time of this writing, no official frameworks regarding internet research ethics have been adopted at any national or international level.⁴⁹ Social media research of this kind is a new field, and the practical guidelines for applying these principles to social media research is often unclear, and remains an issue of debate and disagreement between institutions and individuals.

We consider that the two most important principles to consider for this work are whether informed consent is necessary to collect, store, analyse and interpret public tweets, and whether there are any possible harms to participants in including and possibly republishing tweets, as part of a research project, which must be measured, managed and minimised. We carried out a series of measures to respect these principles, which are set out in the annex, including:

- carefully reviewing Twitter’s terms and conditions and determining whether API-based research for this project was compatible with informed consent
- generally treating data collected as non-individual: they are anonymous and aggregated wherever possible
- carefully reviewing all tweets selected for quotation in this report and considering whether the publication of the tweet and the links, pictures and quotations contained within might result in any harm or distress to the originator or other parties involved; for example, if any possibly invasive personal information were revealed in the body of the tweet, this was not used, and as a further measure we removed any user names, and in a small number of cases ‘cloaked’ the text so its originator could not be identified

3 How do people use Twitter to talk about Europe?

We deployed the method described in chapter 2 as a constantly operating technological system, one that, for the three months of the study, continuously collected tweets into one of 18 specific and discrete streams of data, and then for each stream applied NLP classifiers to reduce these collected tweets successively into those that were first relevant, then attitudinal, and then either positive or negative, towards the theme of the stream.

This chapter examines how European citizens discussed issues related to the EU on Twitter, whether the conversations that do exist can be listened to, and whether this listening informs us about people's attitudes.

What kinds of data exist on Twitter?

The outputs of our research system were first examined very broadly to establish the overall contours and attributes of the data that were collected. Tables 1–3 show the features of 'relevant' English, French and German language tweets by theme.

There are a number of different ways for users to use Twitter, for example uploading a linkshare (a url that links to a story on another site), or responding directly to another user's tweet. This information can be captured as metadata, information about each tweet itself.

These are some of the features of Twitter:

Linkshares

Tweets can contain one or more linkshares, url links to other online material. The proportion of tweets that share a link often denotes the role and influence of other material – including news, commentary and analysis – in any given conversation.

Retweets (RTs)

These are tweets that relay or repost the content of another tweet. A conversation with a high proportion of retweets often implies that the conversation is dominated by a smaller number of influential, heavily retweeted online personae.

Replies

Replies (often described as an @reply) are tweets that are directly replying or addressing another Twitter user. They are often used to sustain a conversation between users, and a high proportion of replies can indicate that a given conversation on Twitter is more sustained and conversational than others are. Importantly, they are not private ‘whispers’ to other users: these are facilitated on Twitter by another function – the ‘direct message’.

User mentions

Tweets can contain one or more user mentions, explicit mentions of other Twitter accounts somewhere in the tweet. The presence of a high proportion of tweets containing user mentions, similar to replies, implies that the tweets on a given theme are more conversational. Note that all tweets that are replies are also classified as tweets with user mentions.

Hashtags

Tweets can contain one or more hashtags (or #tags). These are used to ‘tag’ a tweet as belonging to a particular topic or conversational thread. These tags are decided by the users themselves, and include a rapidly changing landscape of annotations that locate a tweet as a member of a wider conversation. A theme that collects a large proportion of #tags implies it is the subject of a broader-reaching discussion on Twitter. Tweets without #tags are often intended for the Twitter user’s own followers. Tweets can include several #tags.

Table 1 **The features of ‘relevant’ English language tweets, by theme**

Theme	With linkshare (%)	Retweets (%)	Replies (%)	With #tag (%)	User mentions (%)
Euro	64	27	6	35	23
ECHR	67	50	4	92	37
Barroso	59	41	5	52	31
European Parliament	67	40	4	54	35
European Commission	76	34	2	38	23
EU	65	35	4	59	27

Table 2 **Features of ‘relevant’ French language tweets, by theme**

Theme	With linkshare (%)	Retweets (%)	Replies (%)	With #tag (%)	User mentions (%)
Euro	66	32	7	37	23
ECHR	54	47	10	56	32
Barroso	53	45	6	47	27
European Parliament	60	51	3	66	37
European Commission	62	43	5	44	23
EU	59	39	8	43	34

Table 3 **Features of ‘relevant’ German language tweets, by theme**

Theme	With linkshare (%)	Retweets (%)	Replies (%)	With #tag (%)	User mentions (%)
Euro	72	29	4	77	14
ECHR	72	30	14	61	27
Barroso	—	—	—	—	—
European Parliament	81	33	3	47	15
European Commission	81	31	3	63	12
EU	69	29	6	62	18

Findings: data types

Overall, there is a large volume of ‘relevant’ tweets available for every subject, and in every language

A total of 1.45 million tweets considered to be ‘relevant’ to one of the six EU-related themes were collected over three months. There were almost 400,000 tweets about the euro currency and 430,000 about the EU in English.⁵⁰ Perhaps more interesting is the large volume of tweets on the more niche and specific institutions. The ECHR was mentioned in over 30,000 English language tweets, over 12,000 tweets in French and 750 German language tweets. Predictably, there are more English language tweets (almost a million) that are relevant to one of the study’s themes than German (176,000) or French (286,000) tweets. This is because there are roughly three times more British users of Twitter than either French or German users of Twitter, and more English language users on Twitter than French or German language users.⁵¹

Users frequently use hashtags to link to the wide discussion

A large proportion of tweets also contain a hashtag – often around half depending on the stream. Hashtags allow users

to join larger conversations, making it easier for other users to find their tweet. In a random sample of 500 English language tweets about the EU, around 45 per cent used hashtags to link their conversations to signal the topic(s) that they discussed (most used were: #Cyprus, #EU and #Eurozone). A random sample of tweets in French about the EU also found that around 40 per cent used hashtags, and also usually to link to current news stories. At other times, though, hashtags are used as a shorthand by users to express an opinion on the subject matter in hand, and 11 per cent used a hashtag to express an opinion in English. Examples included #betteroffout, #immorality, #no2eu, #Eurogeddon, #fail, #Merkelstan, #WakeTheFuckUp and #DayLightRobbery. In French, this was only 7.5 per cent, and included #TroikaGameOver, #maisouimaisoui, #oceantwentyseven, #Basta!, #danslcul, #volteface and #anticonstitutionnel. Common non-attitudinal hashtags include #UE, #Chypre, #austerite, #Melenchon, #Europe and #MotsCroises.

Twitter tends to be used to ‘broadcast’ rather than as ‘conversation’

The extent to which tweets are ‘broadcast’ (simply sharing a message) or ‘conversations’ (a dialogue between two or more users) can be partially hinted at by the number of tweets that contain another users’ Twitter name. Only a small proportion of tweets are direct replies to other tweets (typically under 10 per cent) but around one-third include a user name, although these are often in the context of users quoting other well-known accounts such as ‘can we the people beat @bobjarr at his lobbying?’

Twitter is often used to share information rather than express opinions

Over half of every stream, and in many cases substantially more, were tweets that shared a link to a site beyond Twitter. A substantial number of these links were to media stories, and a substantial number of tweets linking to media stories contained no additional comment by the tweeter themselves.

For example, of a random sample of 500 English language tweets about the EU, the majority (60 per cent) included a url link, which contained the headline of the article being shared. Similarly, a random sample of 500 tweets in French and 500 tweets in German revealed a similar prevalence of link sharing: 81 per cent of French and 62 per cent of tweets in German included a linkshare. This is echoed in the significant proportion of 'relevant' tweets that were also retweets. Overall, between a quarter and a half of all tweets are retweets, often with the functional aim of relaying a particular nugget of information to the Twitter user's followers.

The dominance of information-sharing tweets has profound implications for the kind of attitudinal insight that can be drawn from Twitter. Although some users certainly share stories they agree with, others appear to share stories out of general interest and even sometimes because they strongly disagree with them. Of those where there was a linkshare and news headline, around 45 cent were deemed by an analyst to be some expression of an attitude (whether by the poster or whoever had written the headline). Similarly, there is no clear relationship between a retweet, and an endorsement or condemnation of the message being retweeted.

Findings: traffic

There is one graph for each of the streams, with all three languages (only relevant data) included at weekly intervals on the y axis and a volume on the x axis.

Findings: technology

Finally, we systematically tested the ability of the technology to analyse reliably the data that we collected.

Automated classifiers are useful for research purposes and for policy makers when they make meaningful distinctions that contribute to useful insight, and they make these distinctions with sufficient individual accuracy such that aggregated measures are reliable. The performance of all the

classifiers used in the project was tested by comparing the decisions they made against a human analyst making the same decisions about the same tweets. A full description of how the classifiers were evaluated, and the results, are included at the end of this report in the methodology annex. The key findings from this process of evaluation are discussed below.

The performance of classifiers varied considerably

In these tests the best performing classifiers were almost always correct, and the worst performing classifiers were almost always wrong (and, indeed, worse than chance). Definite patterns emerged about the contexts where the classifiers were successful, and where they were not.

Relevancy classifiers — filtering the correct data sets — are valuable tools

These results suggest that classifiers trained to decide whether tweets are relevant are extremely valuable. Typically, they performed well, correctly classifying the tweets over 60 per cent of the time (F1 score of over .60), which suggests they are a useful way to categorise large data sets.

Classifiers to identify ‘attitudinal’ tweets performed less well

The idea behind training a classifier was to be able to determine how many of the tweets were from EU citizens expressing a clearly defined opinion on a relevant subject. The training data made this quite difficult to create a model for. We believe this was because a large proportion of tweets did not contain a clearly or obviously expressed opinion. Many were linkshares, where an attitude may be inferred but not clearly expressed. This creates difficulties for classifiers, because they are trained on the data they are provided. For example, if a classifier was trained to place shared links to a headline from an article into ‘no attitude’, because the structural and linguistic features of that text do not necessarily bear any relationship to the category ‘no attitude’, it is not able to extract clear rules by which to make decisions. In general, classifiers work well when the data are more uniform and human beings can decide clearly what the

meaning of the tweet is. This is something that has not happened for attitudinal tweets, and part of the problem there has been an inability to define clearly the difference between each class, and very changeable and inconsistent mark-up as a result. For instance, the English Barroso data have very high performance – the data set contains a large proportions of tweets that directly praise or criticise Barroso himself

Fast moving, event-specific language hinders the performance of long-term classifiers

Many of the conversations taking place on these platforms were responding to very specific rather than generic events. This may also make replicable language use patterns less likely. There are other difficulties for classifiers: non-literal language use such as sarcasm, pastiche, slang and spoofs are found to be common on social media. The ‘real’ rather than ironic meaning of these kinds of mobilisations of language are inherently contextual and difficult to deduce via automated analyses.

Classifier performance improves when it becomes more specific to a particular conversation

Classifier performance varied according to the task assigned, and in some cases (see the annex) it performed poorly. However, when trained closely against an event-specific data set, performance was vastly improved, which reflects the event-specific nature of language on Twitter.

Classifier performance improves when it is trained to make distinctions that are naturally present in the data

When the categories of meaning are clearly present in the data, rather than applied from above by an analyst, it is likely that humans will agree more often on what tweets fit these categories, the training data will present clear patterns and correlations to the NLP algorithm, and the decisions the algorithm thereafter makes will be consequently more accurate.

4 Case studies of real world events

A number of key insights implied that continuously running, broad and top-level research did not reflect either how people use Twitter, or the best ways to analyse it. People did not in general express generic sentiment on Twitter about the EU. Our analytical method was also not well suited to produce an accurate picture of constant, rolling sentiment: specific topics changed and the way that people spoke about those topics also changed.

Instead, Twitter was found to be fundamentally a reactive medium. A tweet is overwhelmingly a reaction to an event that the tweeter has otherwise encountered – either online or offline, whether through reading mainstream media or being told about it by a friend. Twitter use fits into how a person engages with the world as they learn about it from a much wider ecology of different currents of information. In this chapter we examine how Twitter responded to real-time events through a series of case studies.

An attitude expressed on Twitter is usually a social gloss, a non-neutral piece of commentary about a specific event. A body of tweets is really a snapshot mosaic of opinions from people who have been spurred to react to something they have read about, either in the news or on Twitter, and almost always something that has happened, either online or offline. Sentiment-bearing tweets therefore are almost always anchored in the context of important events that prompt discussion, and the mainstream media environment that reports on them.

What we are witnessing is the reaction of ordinary people to events as they unfold – so-called ‘twitcidents’ – a digital annotation of an important event. A complex, varied and evolving storm of reaction on Twitter is a new kind of aftermath to events of significance – an online shadow of interpretations, condemnations, jokes, rumours and insults.

The opportunity to learn about attitudes from these kinds of data is not in any sense to learn about them as generic or general. The opportunity is to identify and analyse these twitcidents as bodies of a specific kind of reactive sentiment expressed within a specific context. Only from this very event-specific context would it be possible carefully to begin to infer the more general or fundamental attitudes that this specific reaction implies. This is in line with the optimal performance of the technology we developed – specific conversations in specific contexts with specific ways of using language to express meaning.

An event-specific method of analysis was developed along the following lines:

- *Identify surge in relevant Twitter traffic.*
- *Describe the contours of the surge:* Identify when it began, how it evolved and when it ended.
- *Determine the cause and topic of the surge:* Understand what the tweets included in the surge are talking about, and the broader context within which they are made. Qualitative dips were taken into the tweets located at one or a number of points during the surge, analysing information about the tweets – such as what links they were sharing, and what #tags they contained, and also to build a picture of the backdrop against which the surge occurred – what offline events were occurring at the time – a relevant speech by a EU politician – and whether the media was reporting an important and related news story.
- *Determine event-driven attitudes:* Only now, with a developed understanding of the context within which people use Twitter, is it possible to infer people's attitudes towards EU institutions. This was done through moving carefully (where possible) through three stages: unstructured, qualitative analysis of randomly selected tweets to suggest broad distinctions present in the data; manual structured coding of randomly selected tweets to formalise and measure these distinctions; and

the training of a bespoke automated classifier to make this distinction for all the tweets that were part of the surge.


- *Draw out wider, more general insights:* From this very specific, contextualised, event-driven analysis, wider insights can be drawn, including how different twitcidents relate to each other, and how different language groups react to a common story, event or controversy.

To make, test and demonstrate this way of understanding Twitter, we undertook a number of case studies using digital observation.⁵² Surges of Twitter traffic in one or more of our streams were identified. For each, we applied a mix of qualitative and quantitative methods to provide an idea of the context and content of the twitcident. Overall, this method was tested to see how far listening to the tweeted reaction surrounding an event can provide a useful and meaningful insight into citizens' attitudes about these events, and then more broadly the themes these events relate to.

Case study 1: Cyprus bailout

The 2012–13 Cypriot financial crisis involved the exposure of Cypriot banks to overleveraged local property companies, the Greek debt crisis (Cypriot banks had made loans to Greek borrowers that were worth 160 per cent of the island's GDP), the downgrading of the Cypriot Government to junk status by international rating agencies, the consequential inability to refund its state expenses from the international markets, and the reluctance of the Cypriot Government to restructure Cyprus' troubled financial sector. On 16 March 2013, Cyprus became the fifth nation (after Greece, Ireland, Portugal and Spain) to get a Eurozone bailout as the Eurogroup, European Commission, European Central Bank and International Monetary Fund agreed on a €10 billion bailout with Cyprus to recapitalise its ailing banking system in return for a series of drastic measures which would hit the country's depositors.

Figure 1 **Tweets in French and German about the euro, 1 March–9 May 2013**



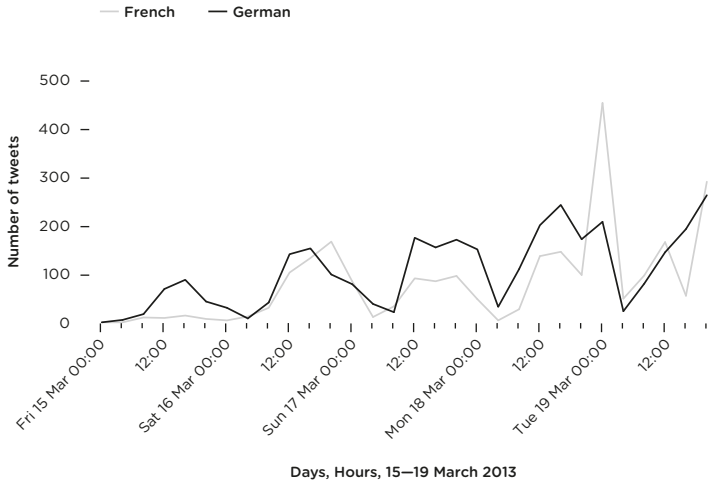
From 16 March to 19 March 2013, the number of tweets in French and German about the ‘euro’ spiked dramatically (figure 1).

16–17 March 2013: bailout announcement

On the morning of 16 March 2013 the French newspaper *Libération* ran with the headline ‘Dix milliards d’euros pour sauver Chypre’ (‘Ten billion euros to save Cyprus’); *Le Figaro* had ‘Chypre: un sauvetage inédit à 10 milliards d’euros’ (‘Cyprus: an unusual 10 billion euros rescue’) and *Le Monde*, ‘À Chypre, la population sous le choc, le président justifie les sacrifices’ (‘In Cyprus, people in shock, the president justifies the sacrifices’).⁵³ In Germany, *Der Spiegel* announced, ‘Hitting the savers: Eurozone reaches deal on Cyprus bailout’.⁵⁴

Immediately following the news, the number of French and German euro streams surged. The greatest number of conversations happened between midday and 6pm every day, building up to higher and higher peaks of traffic. German conversations consistently built up peaks of conversation

Figure 2 Tweets in French and German about the euro, 15–19 March 2013



between 12pm and 6pm every day: 116 on 15 March, 214 on 16 March, 257 on 17 March, 329 on 18 March, dipping on 19 March to 299 and climbing on 20 March to 380. French conversations were less consistent, with a high peak of 516 between 6pm on 18 March and midnight on 19 March (figure 2).

Using randomly drawn qualitative dips of 100 tweets on these days, we found the French and German conversations about the euro were dominated by conversations about the Eurozone, especially Cyprus. Around 90 per cent of tweets in French were about Cyprus – 87 per cent actually contained the word ‘*Chypre*’ – Cyprus – and 30 per cent used the #cyprus hashtag; 70 per cent of tweets in German were specifically about Cyprus, and an additional 26 per cent of tweets were about the Eurozone crisis as a whole. In contrast, the day before – 15 March 2013 – just 12 of 100 randomly selected relevant tweets about the euro referred to the Cypriot banking crisis.

At this early stage of the twitcident, tweets in German and French overwhelmingly just shared the information that the bailout had been agreed.⁵⁵ However, attitudes related to the Cypriot bailout were soon expressed more frequently. In French tweets at around noon on 16 March, people started sharing more negative attitudinal headlines ('Cyprus in shock after bailout plan', 'Cyprus: does the Eurozone still exist?'), though without any explicit comments. Later that day, and the next, people started posting attitudinal remarks of their own about savers' bank accounts being 'plundered' by an 'arbitrary tax', indicating that this detail of the bailout agreement had started to become more widely known.

On 17 March, *Die Welt* took the firm editorial line: 'Zypern schröpft die Sparer' ('Cyprus fleeces savers'). Echoing this, many German speakers seemed to be in solidarity with Cyprus savers, implying that the use of individuals' savings is unfair and anti-democratic. However, plenty of users were not in favour of a bailout, and criticised the amount paid by other citizens of the Eurozone. A very significant number predicted or called for the end of the euro. Some people mentioned the comparative strength of the deutschmark.

19 March: demonstrations

Against the background of large demonstrations outside the House of Representatives in Nicosia by Cypriots protesting at the bank deposit levy, the number of German (927) and French (853) tweets about the euro reached a peak. Again, nearly all were to do with the bailout and Cyprus: 94 per cent of tweets in French were about the bailout, 56 per cent of tweets in German were about the story, with the remainder talking about the euro crisis more broadly. Many German tweeters expressed concern that Cyprus had refused the terms of the bailout, and said that other alternatives were unconvincing. Some called for Cyprus to leave the Eurozone rather than be bailed out. French tweeters were more neutral, sharing news stories as the situation developed, especially after the 5pm announcement that Cyprus had rejected the terms of the bailout. Those French tweeters who did express attitudes usually did so by expressing solidarity with Cyprus' rejection of the ECB-imposed levy.

21 March: plan B

In the wake of the rejection of the terms of the bailout, Cypriot politicians tabled seven ‘plan B’ bills to parliament. As protestors clashed with riot police outside parliament, the European Central Bank piled the pressure on Cyprus by warning it would cut off its emergency liquidity assistance after 25 March unless an EU–IMF programme was in place.

The volume of French and German conversations about the euro remained high on 21 March – 1,000 in French, and 669 in German. Nearly all (98 per cent) of the tweets in French were about the situation in Cyprus (82 per cent still used the word ‘Chypre’), although the tweets in German became more general – 36 per cent were about the Cypriot banking crisis, but discussion was turning more generally to the euro and the Eurozone.

The tweets in French focused on the ultimatum put to Cyprus by the European Central Bank. Most tweets simply shared news stories and headlines. Of the few that expressed attitudes, all were negative, saying that the European Central Bank had ‘declared war on a European country’, and that the move was ‘an act of war under international law’. By mid-afternoon, the story that it was now a real possibility that Cyprus might leave the Eurozone broke, and a negative headline stating that the ECB was proposing to ‘strangle’ Cyprus was much shared. Later in the evening it was reported that Cyprus was working on a plan B, and that the EU was ready to discuss it. Similar to the tweets in French, German tweeters principally used Twitter to keep track of the fast changing events. Most simply shared links to articles reporting neutrally on the changing situation. Several tweeted the statement, ‘Euro-Retter meinen, die Situation auf Zypern sei emotionsgeladen! Vielleicht merken sie auch noch das die ganze Euro Zone geladen ist!’ (‘Ministers proposing a bailout think that emotions are running high in Cyprus. Perhaps they will notice that the entire Eurozone is hopping mad!’) The anti-euro sentiment in the German data set continued: some tweets linked to an article about Germans’ wish to return to the deutschmark.

24 March: bailout terms agreed

By 24 March, discussions of the 'euro' in both languages had begun to decline (there were only 466 in German and 249 in German that day). However, they spiked again on 25 March (722 tweets in German and 645 in French) with the news that an eleventh-hour deal for a 10 billion euro bailout was agreed between the Cyprus Government and the Troika, which safeguarded small savers while inflicting heavy losses on uninsured depositors (including many wealthy Russians using Cyprus as a tax haven). That day it had been decided that deposits up to €100,000 would be protected, but that any holdings larger than this would suffer a 'haircut' of up to 40 per cent. The revised agreement, expected to raise €4.2 billion in return for the €10 billion bailout, did not require any further approval of the Cypriot parliament.

Almost all (99 per cent of) tweets in French and over half of tweets in German were about Cyprus (with the other half referring to the Eurozone crisis more generally). In France, nearly all tweets shared headlines, quite a number of them attitudinal but with no express endorsement or comment by the poster. The deal would have 'heavy social consequences' according to one much-shared article, but another claimed that the deal 'brings an end to the uncertainties facing Cyprus and the Eurozone'. Some noted how the markets had reacted with relief to the agreement. Another much-shared article later in the day announced that Cyprus staying in the Eurozone was still not guaranteed. The last-minute bailout of Cypriot banks was the main topic of the news for German language tweeters, in particular the consequence of the bailout on Germany's relationship with the rest of the Eurozone; 54 out of 100 randomly selected tweets referred specifically to the Cypriot bank bailout, with the other 46 referring to the Eurozone crisis more generally. Anger at the bailout was still in evidence. General dissatisfaction with the single currency began to replace the previous focus on Cyprus – many tweets mentioned the European crisis in more general terms without specifically referencing Cyprus. There were 'harsh but fair' calls for Germany to leave all the other countries. There was a call for

the politicians responsible for the crisis to be jailed, and an ‘us against them’ sentiment: ‘So now Cyprus is saved. Who will save us from the saviours? #euro #merkel’.

29–30 March: the end of the crisis

By 29 March, volumes in relevant conversations about the euro in both languages had declined from their 25 March high. On 29 March there were 386 relevant conversations on the euro in French and 302 in German. The tweets in French continued to discuss Cyprus (half of the sampled tweets shared an article from *Le Monde* reporting that the Cypriot president assured people that Cyprus would stay in the Eurozone). From the morning onwards, another much repeated tweet was the headline ‘Let’s make the fiscal paradises jump’. In German, there was significantly less discussion about the Cyprus issue. High-levels of anti-euro sentiment continued throughout a more topically varied and general discussion, with one of the most shared stories an article looking at the facts behind a claim that Germans have less money than Italians or Spaniards, and another by the Federation of German Wholesale and Foreign Trade (BGA), which feared the collapse of the Eurozone. By 30 March, volumes of tweets in French and German had returned to pre-Cyprus levels. A high proportion of the remaining French conversations still discussed Cyprus (92 per cent on 30 March) – while tweets in German had become more wide-ranging.

Discussion

The issue of Cyprus caused a long-running twitcident, following the many twists and turns of the story as it evolved over a number of days. It showed the two important functions of Twitter that underlie many twitcidents: a way of sharing information to announce and learn about important events and keep apace with them as they rapidly develop, and a way of exchanging opinion about those events.

News of the Cyprus bailout had different implications for the Germans and the French. While it caused similar immediate surges of conversations on Twitter in both languages, in France, the conversation remained tightly focused on Cyprus, and

conversations declined as the prominence of the issue of Cyprus itself declined. In Germany, it awakened a broader conversation about the Eurozone, its future, and Germany's place within it, which continued on Twitter beyond the issue of Cyprus itself.


The overall attitudes expressed throughout the twitcident were critical towards the Cyprus bailout (often in solidarity with Cypriot depositors), sceptical that it would stabilise Cyprus, and in broader terms increasingly concerned about the stability of the Eurozone, and the implications of the instability for the individual and their own national economy. This is consistent with what else we know about attitudes on this subject. A *Guardian* poll in March 2013 found 91 per cent of people thought 'this is just the beginning of the island's problems', and only 9 per cent that 'the agreement means things can only get better'.⁵⁶ This concern for the economic future of neighbours, at the same time as concern for what it means for themselves, has divided the French and German electorates – slightly more than half of Germans generally support helping others, while 60 per cent of French people are against it.⁵⁷ On the specific issue of Cyprus, Germans supported the bailout, and French people opposed it.⁵⁸

Case study 2: European institution events

European Commission Summit: 13-16 March 2013

On 14 March 2013, the leaders of EU member states met in Brussels. From 13 to 16 March, there was a clear spike in the volume of tweets in English about José Manuel Barroso, reaching around ten times the background average on the day of the summit.⁵⁹ The increase in traffic lasted for three days, with smaller spikes in volume on the afternoons of 13 March and 15 March surrounding the most significant spike on the afternoon of 14 March (figure 3).

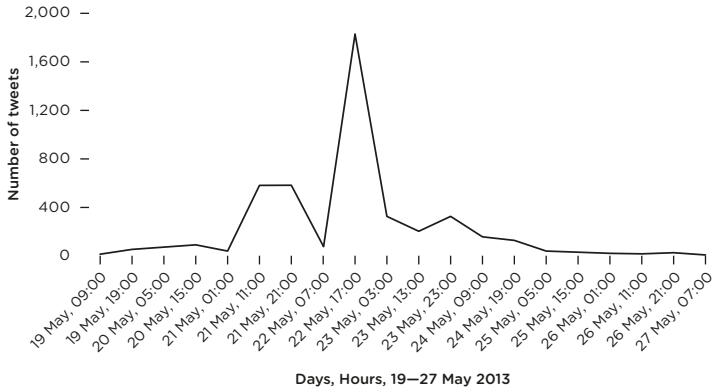
Figure 3 **Tweets in English about José Manuel Barroso, 7–18 March 2013**



Of the 3,518 English language tweets ‘relevant’ to José Manuel Barroso from the beginning of 13 March to the end of 15 March, over 70 per cent of tweets were about the summit. Yet an unusually low number (1,394, or around 40 per cent) shared a link, and no single story was notably dominant. The top ten most shared stories constituted only 299 (21 per cent) linkshares (seven of these ten were official EU websites, primarily covering a speech by Barroso in anticipation of the summit). There was a wide variety of coverage about the summit the day before (13 March), including a number of press releases and speeches by European politicians, and during the summit itself (on 14 and 15 March).

Instead of a single issue dominating discussion, users took the occasion of the summit to talk about the issues related to the EU that affected them. The summit therefore acted as a sounding board for a range of different concerns, fears and hopes that people felt about the EU.

Figure 4 **Tweets in English about the European Commission, 19–27 May 2013**



This analysis found that 45 per cent of tweets were about the EU generally, which includes tweets concerning the EU budget, the European Commission, the European Parliament, the EU’s relationship with Russia and a series of other issues related to single member states. One in three of the tweets voiced economic concerns, discussed healthcare policies or demanded new initiatives in the social sphere.

European Commission Opening: 22 May

After a series of anticipatory press releases and briefings on 21 May, The European Commission opened on 22 May 2013. Also beginning on 21 May and continuing until 22 May, there was a sharp spike in the number of tweets in English ‘relevant’ to José Manuel Barroso (figure 4).

The first surge of tweets (between 8am and 5pm on 21 May) was primarily reportage of the upcoming Commission launch, sharing links to the publications released in anticipation of it. The overwhelming majority of tweets from 21 May referred to the publication of the European Commission’s (and José

Manuel Barroso's) contribution to the European Council meeting the following day, which called for measures on tax evasion (including full tax data exchange) and, to a lesser extent, progress on energy policy.

The second surge, predictably, was a body of reaction to the opening of the European Commission itself. As above, no single issue was dominant – the opening of the Commission spurred people to talk about their own, specific concerns and interests in relation to the EU. The topics raised were diverse, from tax evasion proposals to youth unemployment, the debate about arming Syrian rebels, and David Cameron's remarks about lower taxes for business to increase growth and employment. A significant minority were 'live tweets' – people commenting directly about statements as they were made.

We therefore trained a classifier to distinguish between tweets that were in general optimistic, pessimistic or neither (irrelevant) about the European Commission's ability to enact positive influence on their lives.⁶⁰ Of 1,684 tweets that were posted during the duration of the summit itself, 667 were broadly optimistic, and a very large majority of the rest were non-attitudinal or irrelevant. The generally optimistic attitude of Twitter users towards the Commission's opening is surprising – and appears in direct contradiction to other data about attitudes.

Discussion

People understood and related to these European Commission events through the lens of their own specific grievances, concerns and priorities. They therefore provoked a different kind of twitcident – a heterogeneous collection of different volunteered statements that suggest people's underlying issues of interest and concern. The longer duration illustrates how Twitter can move beyond knee-jerk reaction to a sustained engagement with current affairs as they play out across our computer screens.

Case study 3: European Court of Human Rights controversies

Example 1: 'Casse toi pov'con'


In 2008 Hervé Eon was arrested for waving a small placard exclaiming 'Casse toi pov'con' ('Get lost you sad prick!') during a visit of then French President Nicolas Sarkozy to the French town of Laval and convicted under an old French law that forbids insulting the head of state. His initial conviction was appealed at the ECHR, on the basis that his freedom of expression had been infringed. On 14 March, it was reported that the ECHR had ruled in his favour, arguing that by repeating a phrase ('Eh ben casse toi alors, pauv' con!'), which Sarkozy himself had used during a visit to the Salon de l'Agriculture in 2008 (and which subsequently went viral) the individual was using political satire, which should be protected as legitimate political criticism under human rights law.

A sudden surge of tweets in French 'relevant' to the ECHR began at 9am on 14 March, which lasted for around 24 hours. This was a sharp and symmetrical twitcident, beginning at 6am, peaking around midday at just over 1,800 tweets, and declining over the afternoon and the evening of that day (figure 5).

Within that total surge of 2,710 tweets, 1,934 of the tweets (71 per cent) directly referred to the case by the 'casse toi' quote in the tweet text itself, and 865 used a relevant #tag. Three-quarters (2,025) shared a link, most prominently (451 shares in total) to the article in *Le Monde*⁶¹ that originally broke the news, while similar articles in other mainstream new outlets (*Le Figaro*, *Le Nouvel Observateur*, *Libération*, *Le Parisien*, *20 Minutes*, *France Info*) were also widely shared. However, the second most widely shared link (240 shares) was to the actual ruling itself made by the court.⁶²

Around two-thirds of these tweets did not express an attitude, but simply shared the story or the court's decision without further elaboration. The vast majority of the remaining attitudinal tweets were positive about the ruling. Most took a light-hearted tone: 'Let's not hide our pleasure: let's tweet it!' Many praised or thanked the ECHR explicitly. One tweet reported the court's decision then added, 'That's what Europe is for.'

Figure 5 Tweets in French about the ECHR, 10–23 March 2013

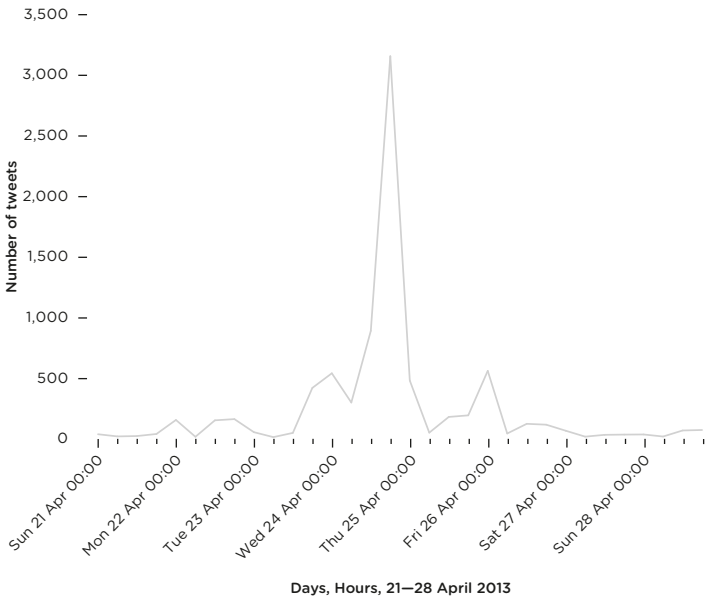


Example 2: The deportation of Abu Qatada

Abu Qatada al-Filistini, a Palestinian Muslim of Jordanian citizenship, had since 2002 been the subject of a long legal battle to deport him from the UK to Jordan, where he had been sentenced to life imprisonment for conspiracy to carry out terror attacks. In 2012 the ECHR – the last legal hurdle to deportation – had ruled that sending Qatada to Jordan would violate his right to a fair trial.

According to one poll, 61 per cent of Britons supported the view that Britain should ‘ignore the court ruling’ and ‘deport Abu Qatada anyway’.⁶³ Most people pointed to the ECHR, ahead of the home secretary or civil servants, as the reason for the delay.⁶⁴ On the morning of 24 April 2013, it was reported that David Cameron was exploring ‘every option’, widely understood to mean a temporary withdrawal from the European Convention of Human Rights, in order to deport Abu Qatada to Jordan. This temporary withdrawal followed by a reratification with certain reservations, it was announced, had been discussed between David Cameron and other cabinet-level ministers.

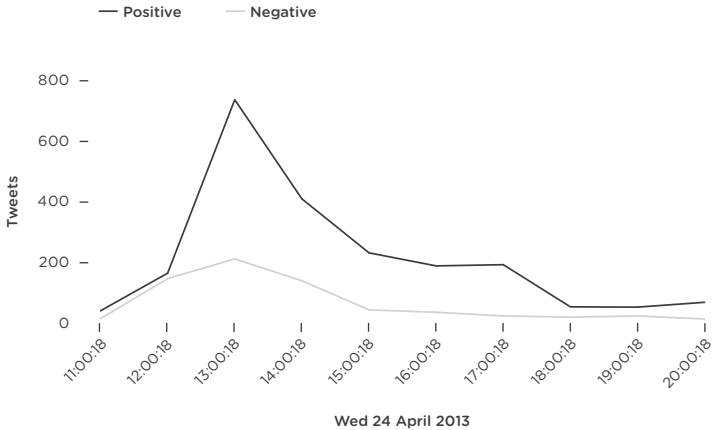
Figure 6 Tweets in English about the ECHR, 21–28 April 2013



Spurred on by this announcement, a passionate debate took place on Twitter about the relative merits of leaving the ECHR in order to deport Abu Qatada. On 23–24 April 2013 the number of English language tweets discussing the ECHR increased above the background level, and surged to a peak of over 3,000 around 5pm (figure 6).

Over the days of 24 and 25 April 2013, 5,834 tweets in English ‘relevant’ to the ECHR were posted. Cameron’s proposal dominated this discussion – 94 per cent of the tweets randomly selected were related to it and 1,785 of these shared a link. The most shared stories were mainstream media and blogs discussing the Government’s proposals. This wider body of shared commentary was primarily hostile to the proposal to leave the ECHR, including a widely circulated and (critical)

Figure 7 Positive and negative tweets about the ECHR, 24 April 2013



legal opinion from the campaign group Liberty about the lawfulness of the UK Government's proposal, and a recording of Thomas More's speech on the primacy of man's laws over God's laws (also deployed in apparent criticism of Cameron's proposal), which appeared on YouTube.

The majority of tweets were also strongly hostile to the idea of a temporary withdrawal from the ECHR: 'The rule of Law clearly means nothing to this government. It is absolutely shocking #ECHR.' Some fitted the suggested move into a wider narrative of recent illiberal government policies; some argued it was a slippery slope towards further abuses; some pointed out the absurdity of such a large change for one person; and many questioned whether it was legally possible: 'Actually quite worried about only having rights when it's convenient to my government. This is not how it's supposed to work. #echr.' However, a smaller group argued that Abu Qatada should be deported at all costs, and Britain did not need the European Convention on Human Rights to safeguard its liberties.

A classifier was trained to classify each tweet as ‘positive’ towards the ECHR (and therefore hostile to Cameron’s proposal to withdraw from it), ‘negative’ towards the ECHR, or ‘non-attitudinal/irrelevant’. Of 1,344 attitudinal tweets, 1,181 (88 per cent) were classified as positive, and 163 negative (12 per cent) (figure 7).

Discussion

The case of the French response to the story of Hervé Eon shows how Twitter is used not only to express disagreement and discontent at perceived injustices but also to thumb one’s nose at authority – as demonstrated by the repeated use of the very phrase that had landed Eon in trouble. Regarding Abu Qatada, we again see a strong reaction against perceived authoritarianism. Both incidences are examples where some domestic authority – a French court, the British prime minister – is seen to take or propose a drastic measure at odds with European legal institutions, and in both instances Twitter users sided with Europe (although this may also be a response to domestic political issues).

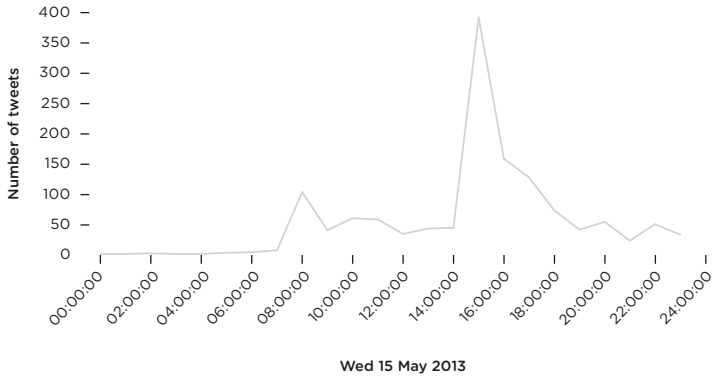
The very strong signal of hostility towards Cameron’s proposal and support for the court is consistent with evidence from opinion polls on British views towards the legitimacy of the ECHR. While, unlike many other European countries, the court was viewed in the UK as something that both improved and harmed democracy, twice as many people viewed its influence to be broadly positive as those who considered it negative.

Case study 4: José Manuel Barroso on the French economy

At 6.30 (GMT) on 15 May 2013, the French National Institute of Statistics and Economic Studies (INSEE) announced that the French economy was officially in recession (-0.2 per cent growth for the second consecutive quarter) while François Hollande was due to meet all 27 European Commissioners later in the day to request an extension for France’s budget reforms.

Barroso was interviewed at 7.20am about his reaction to the French recession and the upcoming meeting with Hollande.

Figure 8 Tweets in French about Barroso, 15 May 2013



Meanwhile, the number of French language tweets about Barroso began to surge. Volumes increased even more sharply that afternoon as in a joint French language press conference at 2pm Hollande and Barroso announced the outcome of that meeting – a grudging acceptance of a two-year extension.

From a background average of 86 per day, 1,419 French language tweets about Barroso arrived over 24 hours on 15 May 2013 – peaking during Hollande and Barroso’s joint press conference (figure 8).⁶⁵

Around 40 per cent of these tweets shared a link, and many of these shared a version of a media narrative that dominated the depiction of the relationship between Barroso and Hollande – that Barroso was admonishing Hollande for failure (table 4).

Barroso’s statements were remarkably incendiary – that France lacked ambition and France prioritises debt servicing over educating its children. Given this dynamic, a classifier was created to identify whether the tweets within the twitcident were supportive of Barroso’s statements, unsupportive, or neither.⁷¹ The ‘neither’ category included any non-attitudinal linkshares, and straight quotations or paraphrases, as well as

Table 4 **Headlines and types of link for tweets in French about Barroso, 15 May 2013**

Headline	Link type	Release time	Number
'Barroso: "the French should be more ambitious"'	Article from TV news channel's business-specific website	9.24	20 (3.7%)
'Recession: Hollande says "the situation is serious"', ⁶⁶ 'Barroso: "France must suggest credible reforms"' ⁶⁷	Newspaper article	9.51	63 (11.5%)
'José Manuel Barroso: "France spends more on its debt than it does educating its children!"'	Blog post — includes a link to Barroso's Europe 1 interview	12.22	56 (10.2%)
<i>Unavailable</i>	The live news stream from the French Government home page	—	32 (5.9%)
<i>Unavailable</i>	European Commission live streaming of the joint press conference	—	16 (2.9%)
Live press conference JM Barroso and F Hollande [<i>link no longer working</i>]	Page on Belgian equivalent of BBC news website	—	14 (2.6%)
<i>Unavailable</i>	European Commission's timetable for live event streaming	—	13 (2.4%)
'José Manuel Barroso: "To be against globalisation is spitting into the wind"' ⁶⁸	Article on French news website	—	12 (2.2%)
'Hollande will not have to pass an "exam" in Brussels, insists Barroso' ⁶⁹	Newspaper article	—	10 (1.8%)
'François Hollande before the European Commission: the President of the Republic met the 27 European commissioners in Brussels to discuss the structural reforms they demanded of France' ⁷⁰	Magazine article	13:40	9 (1.6%)

tweets that were not about Barroso (eg attitudes about or quotes from Hollande).

Over the entire twitcident, the classifier found 19 supportive tweets (1.4 per cent), 363 unsupportive tweets (27.1 per cent) and 956 tweets that were neither (71.4 per cent). Over time, the classifier suggests there was a large number of neutral, ‘reporting’ tweets, followed by a smaller number of ‘commentary’ tweets that were, on the whole, unsupportive. These criticisms of Barroso ranged from the polite – #Barroso ‘Nous attendons des réformes crédibles de la France’ Celles de FH [François Hollande] ne le seraient elles pas jusqu’à maintenant?’ (‘We await credible reforms from France’) – to outright attacks – #BarrosoOnTeMerde.

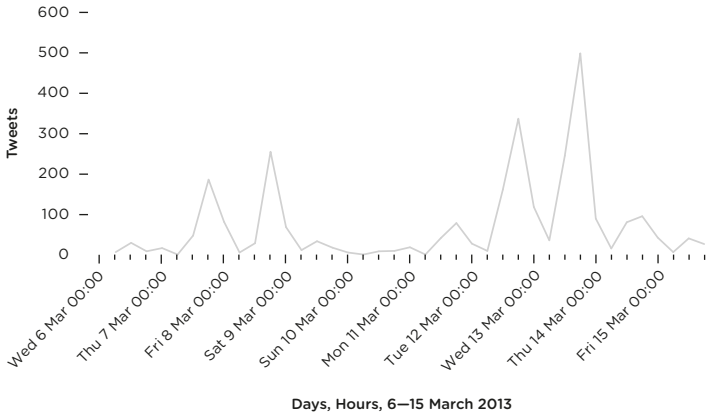
Case study 5: a possible ban on pornography

In the late evening of 6 March 2013 the news was circulated that the Women’s Rights and Gender Equality Committee of the European Parliament had proposed a vote to ban pornography from all forms of media, so that a ‘true culture of equality’ could be achieved in the digital world.

The next morning, the number of German conversations ‘relevant’ to the European Parliament surged from a very low background level to 318 tweets sent over the course of that day, 7 March 2013; 301 of these tweets shared a link and the two most shared – together comprising 196 of these linkshares – discussed the Committee’s recommendation.

There were two significant spikes in the volume of tweets in German ‘relevant’ to the European Parliament – one over 7 and 8 March, the second beginning of the morning of 12 March and continuing until 6pm on 13 March. Consistent between these two clusters of two days, the spike in volume was sharp and symmetrical; there was a rapid rise in the number of tweets from lunchtime to evening, and a rapid decline to very low levels from the evening to that night (figure 9).

Figure 9 **Tweets in German about the European Parliament, 6–15 March 2013**



On 7 March 2013, 89 per cent of tweets referred to the possibility of a ban on pornography, but 68.5 per cent did not record an attitudinal view. Of those that did, the majority (57 per cent) were relatively dismissive of the plan, seemingly viewing this proposal as unlikely to garner any substantial support with the European Parliament; the remainder were highly critical of the potential pornography ban. Indeed, these tweets revolved around the theme that the European Parliament simply did not, or should not, have the authority to enforce such a ban.

On 8 March, as the story reached a broader public, thousands of emails began to arrive from concerned voters to their members of the European Parliament (MEPs). However, at around 12 midnight, the flow of emails was suddenly interrupted. Christian Engström, MEP for the Swedish Pirate party, publicly announced that after receiving more than 350 protest emails, they had suddenly ceased. It was soon discovered (with around half of all tweets that day sharing a link to the story) that following complaints by a number of MEPs, the European Parliament's IT Department had started to filter out these emails as spam.

When the news broke that the European Parliament had enforced an email filter, the story was reflected on Twitter too: 83 per cent of tweets were about the EU Parliament blocking emails. Of these tweets, 76 per cent did not display a negative or positive attitude towards this measure. However, the 24 per cent of tweets that did convey an attitude were almost unanimous in tone – that this episode was a prime example of the European Parliament repressing the opinions of voters (EU citizens), and demonstrated how out of touch the EU institutions had become.

The vote on the proposal took place on 12 March 2013 and concluded at around 1pm. The European Parliament approved the overall report on gender inequalities with 368 out of 625 members voting in favour, but rejected the controversial section containing the ban on pornography. The four most shared stories that day reported on the vote. However, the following day, the largest daily spike in relevant tweets across all three months (early March to early June) (872) focused on a completely different story – the overwhelming rejection by the EU Parliament of the tabled 2014–2020 EU budget ‘in its current form’. The topics of the tweets over 12 March and 13 March therefore drastically changed in reaction to this important announcement.

Discussion

The proposed ban on pornography shows how Twitter is used first to share information about events, especially from mainstream media, and then to talk about them.

The proposed ban allowed people to think about wider European legal institutions and what should be the limit of their power. Broadly, in this instance, Twitter users were supportive of the existing way of things. The blocking of emails critical of the proposed measure shows precisely how sensitive Twitter is to what its users perceive as injustices. However, it seems that to be popular, the message also has to be self-contained, and to demonstrate what it needs to, whether through a photo or link, within the bounds of Twitter itself.

5 Digital observation

The chapters above have described and showcased a new kind of event-specific research method to understand attitudes on Twitter: digital observation.

It is essential to know whether, how far, and in what ways this method of analysis can actually tell us something about people's attitudes – their values, concerns, dispositions, fears and convictions. Finally, what is its future?

Our study found that these data are extremely valuable. We found millions of digital voices talking about EU-related themes, in real time. Many tweets expressed political attitudes about pressing events as they were happening. These tweets were surrounded by a cloud of metadata – everything from when the tweet was made, to how many followers the tweeter has, and sometimes where they are. Some of these metadata were leveraged in this project to aid analysis – but much more could be done (and is being done elsewhere). Overall, Twitter is a new venue for politics, and there exists an extremely valuable opportunity to understand it.

We found that such data sets are 'social big data'. They are often far larger than comparative data sets gathered through conventional polling, interviewing and surveying techniques. Social media data are also noisy, messy and chaotic. Twitter is prone to viral surges in topic, kinds of language used, theme and meme. Twitter data sets are also subject to 'power-laws': the most prolific tweeters tend to be much more prolific than others, those with the most followers tend to have many more followers than anyone else, the most shared links tend to be much more shared than any other. Taken together, any given data set will be profoundly influenced by a number of factors that are very difficult to anticipate beforehand.

Conventional polls, surveys and interviews are not designed to handle the speed and scale at which data are created on Twitter. We found that in order to understand Twitter data, we needed to deploy new technologies that are unfamiliar to sociologists and sociological methods.

Our solution – digital observation – attempted to reconcile and integrate new technologies with conventional techniques, and the long-standing values of social science, but as with any new method of analysis there is a pervasive concern for its quality and credibility.

Generalisability

A key challenge to digital observation is generalisability. When a smaller, representative group is studied, it allows us to extend the findings onto the wider group from which it is drawn. Digital observation does not study representative groups for various reasons:

The data gathered from Twitter may not represent Twitter
Strategies to gather data from Twitter, including our own, often return large bodies of data that are non-representative expressions of systemic non-random bias.⁷² As we described above, we used APIs to deliver tweets that match a series of search terms. The search terms that we used attempted (imperfectly) to gather as many tweets about a given topic as possible, and as few tweets about any other topic as possible. This is difficult to achieve: language use on Twitter is constantly changing, and subject to viral, short-term changes in the way that language is mobilised to describe any particular topic. Trending topics, #tags and memes change the landscape of language in ways that cannot be anticipated, but can crucially undermine the ability of any body of search terms to return a reasonably comprehensive and precise sample. It is therefore probable that tweets about the relevant issue were missed and these tweets, through virtue of using different words and expressions, may be systematically different in attitudes to the ones we did collect.

Tweets may not represent Twitter users

In general, tweets are produced by a small number of high-volume tweeters. Some research suggests that a small number, around 5 per cent, of ‘power-users’ on Twitter are responsible for 75 per cent of Twitter activity.⁷³ These include a small number of dedicated commentators or campaigners on a related issue.

Twitter users may not represent actual people

We found a number of prolific accounts in the data sets that we gathered that not only accounted for a large number of tweets, but were also not EU citizens – our target demographic. These included:

- ‘Twitterbots’ or ‘fake’ accounts programmed to produce automated posts on Twitter
- Official accounts, especially from the EU itself, including the accounts of EU politicians, communications and external affairs agencies and EU offices.⁷⁴

Twitter users may not be representative of EU citizens

Take-up and use of Twitter has not been consistent across EU member states or within them:

- *Geographically*: Around 16 per cent of Europeans use Twitter, and a higher proportion of the population use Twitter in Britain than in France or Germany. Most tweets cannot be accurately located to a particular area – and this study differentiated only on the basis of the language, not specific location, of the tweet.
- *Demographically*: The background of people who use Twitter continues to change, and is linked to the complex phenomenon of how people adopt technology and new habits of using technology. The demographic of the EU’s Twitter users is unlikely to reflect the overall demographic of the EU. The most detailed demographic studies of Twitter use, from the USA, have identified that Twitter users there tend to be young, affluent, educated and non-white.⁷⁵

Truly getting hold of attitudes is a fraught process. Attitudes are complex constructs, labels for those myriad ‘inclinations and feelings, prejudice and bias, preconceived notions, ideas, fears, threats, and convictions’, which we can only infer from what people say.⁷⁶ Does digital observation really uncover attitudes? Can it reliably measure what people say, and does what people say relate to the attitudes that they have?⁷⁷

We have drawn the following conclusions:

Attitudes on Twitter are mixed with a lot of ‘noise’

A significant proportion of our data did not appear to include any discernible attitude at all: the general broadcasting of information, in tweets and through the sharing of links.⁷⁸ Practically, therefore, the mixture of attitudinal and non-attitudinal data drawn from Twitter are not always readily distinguishable. Why precisely people decide to share certain stories is not well understood – and has, to our knowledge, not been studied in detail.

The use of natural language processing is necessary

Faced with far too much data of differing quality and relevance to read and sort manually, the use of new, automated technologies was necessary. The ability of digital observation to measure accurately what millions of people are saying depends on the success or failure of a vital new technology – NLP. Assessing whether and when it can work is vital to understanding when digital observation can add insight, and when it cannot.

To be successful, natural language processing must be used on events, not generically

We showed in chapter 2 that the success of NLP technology overwhelmingly depends on the context in which it is used. Natural language processing tends to succeed when built bespoke to understand a specific event, at a specific time. It tends to fail when it is used in attempts to understand non-specific data over a long period of time.

When used correctly, natural language processing is highly accurate

Where NLP was used appropriately, it was very accurate. As it continues to improve, it is clear that NLP has great potential as part of a reliable and valid way of researching a large number of conversations.

Digital observation will always misinterpret some data

The meaning of language – its intent, motivation, social signification, denotation and connotation – is wrapped up in the context where it was used. When tweets are aggregated as large data sets, they lose this context. Because of this, neither the manual nor automated analysis of tweets will ever be perfect. Automated analysis especially will struggle with non-literal language uses, such as sarcasm, pastiche, slang and spoofs.

Even if we can accurately measure tweets, what do they mean? We make the following observations:

Attitudinal indicators on Twitter may not represent underlying attitudes

There is no straightforward or easy relationship between even attitudinal expressions on Twitter, and the underlying inclinations of the tweeter. Twitter is a new medium: digital social platforms, including Twitter, are new social spaces, and are allowing the explosion and growth of any number of digital cultures and sub-cultures with distinct norms, ways of transacting and speaking. This exerts ‘medium effects’ on the message – social and cognitive influences on what is said. ‘Online disinhibition effect’ is one such influence – where statements made in online spaces, often because of the immediacy and anonymity of the platform, are more critical and rude, and less subject to offline social norms and etiquettes than statements made offline.

It is unclear how Twitter fits into people’s lives

To understand how attitudes on Twitter relate to people, it is important to understand how Twitter fits into people’s broader

lives, how they experience it, and when they use it. Social media, including Twitter, as a widespread habit as well as a technology, is constantly evolving. Our event-specific research was an attempt to fit attitudes on Twitter into how Twitter fits into people's lives. By providing context to situate attitudinal data from Twitter into a narrative of events, it also could then touch on causes, consequences and explanations of attitudes – the 'why' as well as the 'what'.

Current methods struggle to move from 'what?' to 'why?'

The generation of raw, descriptive enumeration of attitudes is not enough. Beyond this, researchers must engage with and contribute towards more general explanatory theories – abstract propositions and inferences about the social world in general, causes and explanations, even predictions – 'why?' and 'where next?', as well as 'what?'. Sociologists understanding meaning in this way often draw on different theories – from positivism to interpretivism and constructionism – each with their own ideas on how to expose the representational, symbolic or performative significance implied or contained in what is said.

Conclusion: a new type of attitudinal research

Digital observation cannot be considered in the same light as a representative poll. Our digital observation of the EU did not attempt to intervene within the EU – by convening a panel, mailing out interviews – to attempt to understand what the whole of the EU thinks. Rather, it lets a researcher observe a new, evolving digital forum of political expression, the conversations of the EU's energised, arguing digital-citizens as they otherwise and anyway talk about events.

This new technique to conduct attitudinal research has considerable strengths and weaknesses compared with conventional approaches to research. It is able to leverage more data about people than ever before, with hardly any delay and at very little cost. On the other hand, it uses new, unfamiliar technologies to measure new digital worlds, all of which are not well understood, producing event-specific, ungeneralisable

insights that are very different from what has until now been produced by attitudinal research in the social sciences.

We believe digital observation is a viable new way of beginning to realise the considerable research potential that Twitter has. It will continue to improve as the technology gets better, and our understanding of how to use and our sense of how digital observation fits in with other ways of researching attitudes become more sophisticated.

Overall

An interaction of qualitative and quantitative methods

Automated techniques are only able to classify social media data into one of a small number of preset categories at a certain (limited) level of accuracy for each message. Manual analysis is therefore almost always a useful and important component; in this report it is used to look more closely at a small number of randomly selected pieces of data drawn from a number of these categories. In scenarios when a deeper and subtler view of the social media data is required, the random selection of social media information can be drawn from a data pool, and sorted manually by an analyst into different categories of meaning.

Subject matter experts at every step

It is vital that attempts to collect and analyse ‘big data’ attitudes are guided by an understanding of what is to be studied: how people express themselves, the languages that are used, the social and political contexts that attitudes are expressed in, and the issues that they are expressed about. Analysts who understand the issues and controversies that surround the EU are therefore vital: to contextualise and explain the attitudes that are found on Twitter, and to help build the methods used to find and collect these attitudes.

For acquiring data

New roving, changeable sampling techniques

The collection of systemically biased data from Twitter is far from easy. The search terms that are used are vulnerable to the

fact that Twitter is chaotically subject to viral, short-term surging variations in the way that language is mobilised to describe any particular topic. During this study, a new data acquisition technique was piloted that attempted to reflect the changing and unstable way people discuss subjects on Twitter. The ‘information gain cascade’ was developed. It is a method intended to ‘discover’ words and phrases that coincide with, and therefore indicate, topics of interest. To do this, a sample of tweets on a topic is collected using high recall ‘originator terms’. A relevancy classifier is built for this stream in the usual way and applied to a large sample of tweets.

The terms (either words or phrases) that this classifier uses as the basis for classification are ranked based on their information gain: a measure of the extent to which the term aligns with the relevant or irrelevant classes. Terms that are randomly distributed between the relevant and irrelevant classes have low information gain, and terms that are much more likely to be in one class than another have high information gain. The terms that have high information gain in the relevant class are designated ‘candidate search terms’. Each candidate search term is then independently streamed, to create its own tweet sample, analysed on their own merits and then, on the decision of an analyst, either graduated to become full search terms, or discarded. This process iteratively ‘cascades’ to continuously construct a growing cloud of terms discovered to be coincident with the originator terms.

This approach allows the search queries used to arise from a statistical appreciation of the data themselves, rather than the preconceptions of the analyst. This method is designed to produce samples containing a large proportion of all conversations that might be of interest – high recall.

Automatic identification of twitcidents

An important but separate area of study is to detect the emergence of twitcidents automatically through statistically finding the ripples that they cast into the tweet stream.⁷⁹ This technology can be used to identify twitcidents as they occur, allowing for the research to be real time, and used reactively.

For analysis

Natural language processing classifiers should:

- be bespoke and event-driven rather than generic
- work with each other: classifiers, each making a relatively simple decision, can be collected into larger architectures of classifiers that can conduct more sophisticated analyses and make more complex overall decisions
- reflect the data: when categories to sort and organise data are applied *a priori*, there is a danger that they reflect the preconceptions of the analyst rather than the evidence. It is important that classifiers should be constructed to organise data along lines that reflect the data rather than the researcher's expectations; this is consistent with a well-known sociological method called grounded theory⁸⁰

For interpretation

- *Accepting uncertainty*: Many of the technologies that can now be used for Twitter produce probabilistic rather than definite outcomes. Uncertainty is therefore an inherent property of the new research methods in this area, and the insights they produce. Therefore there needs to be an increased comfortableness with confidence scores and systematically attached caveats in order to use them.
- *From metrics to meaning*: Of all aspects of attitudinal research on Twitter, the generation of meaningful insight that can be acted on requires the most development, and can add the most value. Attitudinal measurements must be contextualised within broader bodies of work in order to draw out causalities and more general insights.

For use: the creation of digital observatories

Organisations, especially representative institutions, now have the opportunity to listen cheaply to attitudes expressed on Twitter that matter to them. They should consider establishing digital observatories that are able to identify, collect and listen to digital voices, and establish ways for them to be appropriately reflected in how the organisation behaves, the decisions it makes

and the priorities it has. Digital observatories, constantly producing real-time information on how people are receiving and talking about events that are happening, could be transformative in how organisations relate to wider societies.

There must be clear understanding of how they can be used. In the face of the challenges that have just been outlined, the validation of attitudinal research on Twitter is especially important in two senses. Digital observation must:

- validate social media research by the source itself, such as through a common reporting framework that rates the ‘confidence’ in any piece of freestanding piece of research that points out potential vulnerabilities
- address biases in the acquisition and analysis of the information and caveats outcomes accordingly

Social media outputs must be cross-referenced and compared with more methodologically mature forms of offline research, such as ‘gold standard’ administered and curated data sets (such as Census data, and other sets held by the Office for National Statistics),⁸¹ and the increasing body of ‘open data’ that now exists on a number of different issues, from crime and health to public attitudes, finances and transport, or bespoke research conducted in parallel to research projects.⁸² The comparisons – whether as overlays, correlations, or simply reporting that can be read side by side – can be used to contextualise the safety of findings from social media research.

Digital observations must be weighed against other forms of insight. All attitudinal research methods have strengths and weaknesses – some are better able at reaching the groups that are needed, some produce more accurate or detailed results, some are quicker and some are cheaper. It is important to recognise the strengths and weaknesses of attitudinal research on Twitter, relative to the other methods of conducting this sort of research that exist, to be clear about where it fits into the methodological armoury of attitudinal researchers.

Annex: methodology

The methodology annex sets out a more detailed explanation and description of the methods used in this study, and how they performed.

Data collection

APIs

All data from Twitter were collected from its APIs. Twitter has three different APIs that are available to researchers. The 'search' API returns a collection of relevant tweets matching a specified query (word match) from an index that extends up to roughly a week in the past. Its 'filter' API continually produces tweets that contain one of a number of keywords to the researcher, in real time as they are made. Its 'sample' API returns a random sample of a fixed percentage of all public tweets in real time. Each of these APIs (consistent with the vast majority of all social media platform APIs) is constrained by the amount of data they will return. A public, free 'spritzer' account caps the search API at 180 calls every 15 minutes with up to 100 tweets returned per call; the filter API caps the number of matching tweets returned by the filter to no more than 1 per cent of the total stream in any given second, and the sample API returns a random 1 per cent of the tweet stream. Others use white-listed research accounts (known informally as 'the garden hose'), which have 10 per cent rather than 1 per cent caps on the filter and sample APIs, while still others use the commercially available 'firehose' of 100 per cent of daily tweets. With daily tweet volumes averaging roughly 400 million, many researchers do not find the spritzer account restrictions to be limiting to the number of tweets they collect (or need) on any particular topic.

Keywords

To gather data for this report, we accessed the search API that delivers already posted tweets that match a certain keyword, and a filter API that does the same in real time, as tweets are posted. Both of these APIs collect all public instances of a designated keyword being used in either the tweet or the user name of the tweeter. Both these APIs restrict the total number of tweets they will produce as a given total proportion of the total number of tweets that are sent. These ‘rate limits’ were never exceeded during the course of the project.

Acquiring data from Twitter on a particular topic through the use of keywords is a trade-off between precision and comprehensiveness. A very precise data collection strategy generally only returns tweets that are on-topic, but will likely miss some. A more comprehensive data collection strategy collects more of the tweets that are on-topic, but will likely include some which are off-topic. Individual words themselves, reflecting how and when they are used, can be inherently either precise or comprehensive. ‘Euro’ cuts across many different types of issues that are often discussed in high volumes, from the football competition to foreign exchange speculation. Others, like ‘Barroso’, are more often used specifically in the context of discussing José Manuel Barroso.

As noted above, precision and comprehensiveness are inherently conflicting properties of a sample, and a balance must be struck between them. To do this, the search strategy and exact search terms used for each stream were evolved over the early part of the project, before the final phase of data collection began. The search terms for each stream were incrementally crafted by analysts, who monitored how the addition of each term or specific, often topical, annotation of tweets (hashtags) influenced the tweets that were subsequently collected. Both strategies were tried out before final data collection started; in the first week, a high precision search strategy using only a single core term for each stream was used, in the second week a long list of related terms was used to achieve a high recall, and in the third, a balance was struck between both, where enough relevant tweets were collected

without flooding the stream with irrelevant ones. From the third week onwards, a final, balanced approach was taken in which only a short list of directly relevant scraper terms and hashtags was used per stream.⁸³

Each stream struck this balance differently. Some returned larger and generally less precise bodies of data, others smaller, more precise returns. The finalised search terms and the numbers that each produced are shown in tables 5 to 7. Between 5 March and 6 June 2013, we collected approximately 1.91 million tweets in English across the data streams, 1.04 million in French, and 328,800 in German.

Sampling on Twitter is an important example of the lack of clear methodological best practice in social media research. Current conventional sampling strategies on social media construct ‘hand-crafted’ or ‘incidental’ samples using inclusion criteria that are arbitrarily derived.⁸⁴ There are many reasons why a small body of keywords should not be expected to return a sociologically robust, systemically unbiased sample: they are likely to return data sets with ‘systemic bias’, wherein data have been systematically included or excluded in a systematic way; some words or hashtags may be most used by people who hold a particular political position, while other words or hashtags may be used by people who hold another; and unless both sets of words are equally identified and used to acquire a sample, the sample will be biased.

Table 5 shows the data volumes collected for search terms in English on the six themes studied.

Table 5 **The exact search terms used in English and total number of tweets per theme**

Theme	Exact search terms used	Total number
The euro	#Euro #ECB #eurozone #EMU #eurodebtcrisis eurozone ECB	537,741
European Commission	#Euco #EC #EUCommission #Euchat #EUAffairs EU Commission European Commission Euro Commission EUCCO	119,950
European Parliament	#EP #EUDatP #Europarl #Eparliament #EUParl European Parliament Europarl EU Parl EUParliament	124,849
ECHR	#fundamentalrights #ECHR #humanrights #ECTHR #ruleoflaw ECHR ECTHR European Court of Human Rights Euro judges Strasbourg court	278,917
José Manuel Barroso	#Barroso #Barroso #Barosso #euco #EUPresident Barroso Euro chief European Commission chief President of EU Commission EU Commission President	38,160

Table 5 (continued)

Theme	Exact search terms used	Total number
EU	#Eurogroup #EU #Europe #eurostat #Europeanproblems EU eurozone euro-wars E.U.	810,629

Table 6 shows the data volumes collected for search terms in French on the six stream topics studied.

Table 6 The exact search terms used in French and total number of tweets per theme

Theme	Exact search terms used	Total number
The euro	#euro #zoneeuro #eurocrise #bce #eurogroupe banque centrale européenne BCE zone euro eurocrates eurozone	311,529
European Commission	#CommissionEuropeenne #EuropeanCommission #EC #Barroso Commission européenne UE Commission Barroso président eurochief	12,323

Annex: methodology

Table 6 (continued)

Theme	Exact search terms used	Total number
European Parliament	#PE #ParlementEuropéen #ParlementEuropeen #EP #Europarl #pleniere Parlement Europeen Europarl MPE Eurodepute	11,490
ECHR	#CEDH #droitsdelhomme #Strasbourg #ECHR #ECtHR Cour europeenne des droits de l'homme CEDH Cour Europeenne Cour de Strasbourg ECtHR	53,430
José Manuel Barroso	#Barroso #europresident #Baroso #Barosso #CommissionEuropeenne Barroso Barosso commissaire europeen president de la commission eurocrate	7,412
EU	#UE #UnionEuropeenne #europe #EU #Bruxelles Union europeenne l'UE europe	647,830

Table 7 shows the data volumes collected for search terms in German on the six stream topics studied.

Table 7 **The exact search terms used in German and total number of tweets per theme**

Theme	Exact search terms used	Total number
The euro	#Eurozone #Eurobonds #euro #eurokrise #EuroGruppe Eurokrise Eurobonds Euro-Zone Europäische Zentralbank	37,098
European Commission	#Eukommission #Eukom #Kommission #Barroso #EU EU-Regierung Kommission Europäische Kommission Eurokraten EU-Kommissare	59,898
European Parliament	#Europaparlament #EP #MdEP #europarlament #europaischesparlament Europäisches Parlament EU-Parlament Europaparlament EU-Wahlen Europawahl	11,680
ECHR	#menschenrechte #EGMR #EuGHMR #EMRK #Strassburg Europäischer Gerichtshof für Menschenrechte Europäischer Gerichtshof für Menschenrechte EGMR EuGHMR EMRK Europäischen Menschenrechtskonvention	3,579

Table 7 (continued)

Theme	Exact search terms used	Total number
José Manuel Barroso	#Barroso #EUKrise #EUR #EU #StoppESM #EU-Sozialberichts #BarrosoBlunder #Euro-Krise #Eurokrise @Wahlalternativ1 @pittromi @Tiefseher @BarrosoEU @alpenwilli barroso Euro-Krise Eurokrise EU-Kommissionspräsident	
EU	#Europa #EuropäischeUnion #Europapolitik #Euro #EU Europa Europapolitik EU-Politik Europäische Union	216,577

Data analysis

For our study we used a web-hosted software platform, developed by the project team, called Method51, which uses NLP technology to allow the researcher to construct bespoke classifiers rapidly to sort defined bodies of tweets into categories (defined by the analyst).⁸⁵ To create each classifier we went through the following phases using this technology:

Phase 1 – Definition of categories

The formal criteria explaining how tweets should be annotated were developed. This, importantly, continued throughout the early interaction of the data: categories and definitions of

meaning were not arrived at *a priori*, but through relating the direct observation of the contours of the data with the overall research aims. These guidelines were provided to all the annotators working on the task.

Phase 2 – Creation of a gold-standard baseline

On the basis of this formal criteria, analysts manually annotated a set of around 100–200 ‘gold-standard’ tweets using Method51. This phase has two important functions. First, it measures the inter-annotator agreement: the extent to which two human beings agreed with each other on the correct categories for each of the tweets. A low (typically, below 80 per cent) inter-annotator agreement implies that the categories are incorrect: they either are not distinct enough to allow human beings to tell the difference between them dependably, or they do not fit the data, forcing the analyst to make imperfect, awkward and inconsistent categorisations. Second, gold-standard tweets provide a baseline of truth against which the classifier performance was tested.

Phase 3 – Training

The analyst manually annotated a set of tweets to train the machine learning classifier, through web access to the digital observation software interface. The number of tweets that were annotated depended on the performance of the classifier, which itself depended on the scenario. For some streams and for some classifiers, the decision the classifier was required to make, and the data it was required to make the decision on, was relatively straightforward. In others, the analytical challenge was more difficult, and required the creation of larger bodies of training data. Between 200 and 2,000 tweets were analysed for each stream.

Phase 4 – Performance review and modification

The performance of the classifier was reviewed, and examples of its outputs were read. Where feasible and necessary, the algorithm was modified to improve its performance.

Architecture of classifiers

The process above was followed, throughout the lifetime of the project, by 15 human annotators to create a specific ‘architecture’, or system of cooperating classifiers, for each stream. Each stream’s architecture was in the form of a cascade: a number of classifiers that were connected first to the tweets that were being automatically connected, and then with each other to create a coherent cascade of data.

Each architecture comprised at least six levels:

Level 1 – Collection of raw data:

All the tweets were collected through Twitter’s filter APIs, which matched the body of search terms for each tweet.

Level 2 – Language filter:

Raw data were first passed through a language filter to ensure that each tweet was in the correct language for the stream.

Level 3 – Relevancy filter:

All data in the correct language were passed through a ‘relevancy classifier’, an NLP algorithm trained to decide whether a tweet was relevant to the particular theme under which it was collected. The relevancy classifiers were meant to filter out any tweet that did not refer to the topic. For instance, if it was collected under the ‘Barroso’ theme, was the tweet about José Manuel Barroso, the President of the European Commission? The classifier was trained to categorise all tweets as either relevant or irrelevant. Tweets judged to be irrelevant by the classifier were discarded.⁸⁶

Level 4 – Attitudinal filter:

All tweets judged to be relevant were passed through an ‘attitudinal classifier’, an algorithm trained to categorise whether data were attitudinally relevant expressions by an EU citizen, or not. ‘Attitudinally relevant’ tweets were those that expressed, implied or included a non-neutral comment on the topic of the stream as defined for the relevancy classifier.⁸⁷

We only considered tweets that expressed the attitude of the poster as attitudinal; many of the tweets we found contained attitudinal statements from people other than the tweeter, which were quoted or paraphrased as such, but where it could not be assumed that this implied endorsement. All tweets judged to be the former were collected and stored. All tweets judged to be the latter were discarded.

Level 5 – Polarity:

All attitudinal data were passed through an algorithm to categorise tweets as ‘positive’, ‘negative’ or neutral in the nature of the sentiment expressed towards the theme of the stream. Double negative tweets that rejected criticism of the person or institution of interest were considered positive, while obvious sarcastically positive tweets as well as back-handed compliments were considered negative (eg ‘After ruining the European economy, Barroso finally realises austerity has reached its limit. Better late than never I guess’).⁸⁸ For lack of an appropriate category, tweets that simultaneously expressed a positive opinion about one aspect of the stream topic, but a negative one about another (for example, tweets attacking one but defending another MEP for the parliament stream) were marked as neutral.

Level 6+ – Event-specific analysis:

In some cases, additional classifiers were built to make highly bespoke categorisations of the data collected by specific streams in specific time-windows (see below). In these circumstances, a classifier was trained to classify relevant tweets into very context-specific categories of meaning.

Classifier performance

We tested the performance of all the classifiers used in the project by comparing the decisions they made against a human analyst making the same decisions about the same tweets. As stated above, phase 2 of classifier training involved the creation of a ‘gold-standard’ data set containing around 100–200 tweets for each classifier, annotated by a human annotator into the same categories of meaning as the algorithm was designed to do.

The performance of each classifier could then be assessed by comparing the decisions that it made on those tweets against the decisions made by the human analyst. There are three outcomes of this test, and each measures the ability of the classifier to make the same decisions as a human – and thus its overall performance – in a different way:

- *Recall*: The number of correct selections that the classifier makes as a proportion of the total correct selections it could have made. If there were ten relevant tweets in a data set, and a relevancy classifier successfully picks eight of them, it has a recall score of 80 per cent.
- *Precision*: The number of correct selections the classifier makes as a proportion of all the selections it has made. If a relevancy classifier selects ten tweets as relevant, and eight of them actually are indeed relevant, it has a precision score of 80 per cent.
- *Overall, or 'F1'*: All classifiers are a trade-off between recall and precision. Classifiers with a high recall score tend to be less precise, and vice versa. 'F1' equally reconciles performance and recall to create one, overall measurement of performance for the classifier. The F1 score is the harmonic mean of precision and recall.⁸⁹

Note precision and recall must be understood with reference to a particular target class, for example this would typically be the 'relevant' class for the relevancy classifier, and the 'attitudinal' class for the attitudinal classifier. This is particularly important when there are more than two classes, as in such cases there are distinct 'F1' scores for each of the possible target class. In tables 8–10 we show F1 scores for each language, with two scores shown for the sentiment classifiers, the first in cases where the target class is the 'positive' class, and the second where it is the 'negative' class. The performance of each of the decisions that a classifier makes can be drastically different: it can much more reliably select 'relevant' rather than 'irrelevant' tweets, or 'negative' rather than 'positive' ones.

Table 8 **Classifier scores for tweets in English**

Stream	Classifier 1: relevancy	Classifier 2: attitudinal	Classifier 3: sentiment: (positive-negative)
Euro	0.62	0.24	—
ECHR	0.81	—	—
Barroso	0.72	0.63	0.74-0.86
European Parliament	—	—	—
European Commission	0.86	0.19	—
EU	0.83	0.62	1.0

Table 9 **Classifier scores for tweets in French**

Stream	Classifier 1: relevancy	Classifier 2: attitudinal	Classifier 3: sentiment
Euro	0.69	—	—
ECHR	0.94	0.57	0.82-0.67
Barroso	0.68	0.54	0.17-0.84
European Parliament	0.97	0.49	0.13-0.71
European Commission	0.93	0.27	—
EU	0.64	—	—

Table 10 **Classifier scores for tweets in German**

Stream	Classifier 1: relevancy	Classifier 2: attitudinal	Classifier 3: sentiment
Euro	0.65	—	—
ECHR	0.45	—	—
Barroso	—	—	—
European Parliament	0.89	0.29	0.29-0.27
European Commission	0.69	—	—
EU	0.66	—	—

Classifier performance: event-specific data sets

We also produced a small number of event-specific classifiers for chapter 4 (case studies of real world events). These classifiers were trained on smaller data sets, but were specific to one event that caused a large surge in traffic surrounding an offline event:

- *Classifier 1*: European Commission opening on 22 May – whether the tweeter was ‘broadly optimistic’ or ‘broadly pessimistic’ about the ability of the European Commission to enact positive influence on the tweeter’s life; this had an F1 score of 0.63
- *Classifier 2*: whether, in the context of the proposal to suspend Britain’s membership of the European Convention on Human Rights temporarily in order to deport Abu Qatada, the tweeter was ‘broadly positive about the European Court of Human Rights’, or ‘broadly negative’; this had an F1 score of 0.68
- *Classifier 3*: whether tweets that were supportive of José Manuel Barroso’s criticism of France’s failure to enact meaningful budgetary reform on 15 May 2013; this had an F1 score for ‘supportive’ of 1.0 and 0.9 for ‘unsupportive’

Ethics

We consider that the two most important principles to consider for this work are whether informed consent is necessary to collect, store, analyse and interpret public tweets, and whether there are any possible harms to participants in including and possibly republishing their tweets, as part of a research project, which must be measured, managed and minimised.

Informed consent is widely understood to be required in any occasion of ‘personal data’ use when research subjects have an expectation of privacy. Determining the reasonable expectation of privacy someone might have is important in both offline and online research contexts. How to do this is not simple. The individual must expect the action to be private and this expectation must be accepted in society as objectively reasonable.

Within this frame, an important determination of an individual's expectation of privacy on social media is by reference to whether the individual has made any explicit effort or decision in order to ensure that third parties cannot access this information. In the UK, there are a number of polls and surveys that have gauged public attitudes on this subject, including a small number of representative, national level surveys. Some research suggests that some users have become increasingly aware of the privacy risks and have reacted by placing more of their social media content onto higher privacy settings with more restricted possible readerships.⁹⁰ Users are taking more care to manage their online accounts actively; figures for deleting comments, friends and tags from photos are all increasing, reported a Pew internet survey.⁹¹ Taken together, the surveys find that citizens are increasingly worried about losing control over what happens to their personal information, and the potential for misuse by governments and commercial companies.⁹² However, these surveys also show that it is less clear what people actually understand online privacy to entail. They found that there is no clear agreement about what constitutes personal or public data on the internet.⁹³

Applying these two principles to Twitter for our work we believe that those who tweet publicly available messages in general expect a low level of privacy. (This is not true of all social networks.) Twitter's terms of service and privacy policy both state: 'What you say on Twitter may be viewed all around the world instantly',⁹⁴ and the terms of service also states: 'We encourage and permit broad re-use of Content. The Twitter API exists to enable this.'⁹⁵ We believe that people have a relatively low expectation of privacy on Twitter, given recent court cases that have determined tweets are closely analogous to acts of publishing, and can thus also be prosecuted under laws governing public communications, including libel.

That does not remove the burden on researchers to make sure they are not causing any likely harm to users, given users have not given a clear, informed, express consent. Harm is difficult to measure in social media research. We drew a distinction in our research between key word searches and

named account searches. We built no detailed profiles about any online user, or offline person. This was partly a technological challenge: extraction tools need to be designed to avoid accidental extraction from non-public accounts, and new forms of collection – such as extracting profile information – might in some instances require explicit consent.

Notes

- 1 There are a number of new and emerging academic disciplines developing in this area, most notably computational sociology and digital anthropology.
- 2 It may also partly be a reflection of the network effect of social networks. For example, given the high proportion of English on Twitter, non-English users may also feel compelled to use English as well, to take part in conversations on the network.
- 3 Attitudinal research itself can often change the context of what is said, and in doing so introduce 'observation' or 'measurement' effects'. This is 'reactivity' – the phenomenon that occurs when individuals alter their behaviour when they are aware that they are being observed. People involved in a poll are often seen to change their behaviour in consistent ways: to be more acceptable in general, more acceptable to the researcher specifically, or in ways that they believe meet the expectations of the observers. See PP Heppner, BE Wampold and DM Kivlighan, *Research Design in Counseling*, Thompson, 2008, p 331.
- 4 See BG Glaser and AL Strauss, *The Discovery of Grounded Theory*, New Brunswick: AldineTransaction, 1967.
- 5 These are the six principles: research should be designed, reviewed and undertaken to ensure integrity, quality and transparency; research staff and participants must normally be informed fully about the purpose, methods and intended possible uses of the research, what their participation in the research entails and what risks, if any, are involved; the confidentiality of information supplied by research participants and the anonymity of respondents must be

- respected; research participants must take part voluntarily, free from any coercion; harm to research participants and researchers must be avoided in all instances; and the independence of research must be clear, and any conflicts of interest or partiality must be explicit. See ESRC, 'Framework for Research Ethics', latest version, Economic and Social Research Council Sep 2012, www.esrc.ac.uk/about-esrc/information/research-ethics.aspx (accessed 13 Apr 2014).
- 6 However, a growing group of internet researchers has issued various types of guidance themselves. See AoIR, *Ethical Decision-Making and Internet Research: Recommendations from the AoIR Ethics Working Committee (Version 2.0)*, Association of Internet Researchers, 2012, p 2.
 - 7 European Commission, Eurobarometer survey on trust in institutions, Nov 2013, http://ec.europa.eu/public_opinion/cf/showchart_column.cfm?keyID=2189&nationID=6,3,15,&startdate=2012.05&enddate=2013.11 (accessed 24 Apr 2014); I van Biezen, P Mair and T Poguntke (2012) 'Going, going... gone? The decline of party membership in contemporary Europe', *European Journal of Political Research* 51, no 1, 2012, pp 24–56.
 - 8 J Birdwell, F Farook and S Jones, *Trust in Practice*, London: Demos, 2009.
 - 9 European Commission, 'Public opinion in the European Union: first results', *Standard Eurobarometer* 78, Dec 2012, http://ec.europa.eu/public_opinion/archives/eb/eb78/eb78_first_en.pdf (accessed 10 Apr 2014).
 - 10 Pew Research Center, 'The sick man of Europe: the European Union', 13 May 2013, www.pewglobal.org/2013/05/13/the-new-sick-man-of-europe-the-european-union/ (accessed 10 Apr 2014).

- 11 European Commission, 'Two years to go to the 2014 European elections', *Eurobarometer* 77, no 4, 2012, www.europarl.europa.eu/pdf/eurobarometre/2012/election_2012/eb77_4_ee2014_synthese_analytique_en.pdf (accessed 11 Apr 2014).
- 12 P Huyst, 'The Europeans of tomorrow: researching European identity among young Europeans', Centre for EU-studies, Ghent University, nd, http://aei.pitt.edu/33069/1/huyst._petra.pdf (accessed 11 Apr 2014).
- 13 M Henn and N Foard, 'Young people, political participation and trust in Britain', *Parliamentary Affairs* 65, no 1, 2012.
- 14 Eg J Sloam, 'Rebooting democracy: youth participation in politics in the UK', *Parliamentary Affairs*, 60, 2007.
- 15 D Zeng et al, 'Social media analytics and intelligence: guest editors' introduction', in *Proceedings of the IEEE Computer Society*, Nov–Dec 2010, p 13.
- 16 Emarketer, 'Where in the world are the hottest social networking countries?', 29 Feb 2012, www.emarketer.com/Article/Where-World-Hottest-Social-Networking-Countries/1008870 (accessed 11 Apr 2014).
- 17 Social-media-prism, 'The conversation', nd, www.google.co.uk/imgres?imgurl=http://spirdesign.no/wp-content/uploads/2010/11/social-media-prism.jpg&imgrefurl=http://spirdesign.no/blog/webdesignidentitet-og-trender/attachment/social-media-prism/&h=958&w=1024&sz=301&tbnid=EFQcS2D-zhOj8M:&tbnh=90&tbnw=96&zoom=1&usq=__VXussUcXEMznT42YLhgk6kOsPIk=&docid=hog_RAXkIYvcpM&sa=X&ei=gQBXUdeYO-iJoAXdyIHYAg&ved=oCEoQ9QEwAg&dur=47 (accessed 11 Apr 2014).

- 18 F Ginn, 'Global social network stats confirm Facebook as largest in US & Europe (with 3 times the usage of 2nd place)', Search Engine Land, 17 Oct 2011, <http://searchengineland.com/global-social-network-stats-confirm-facebook-as-largest-in-u-s-europe-with-3-times-the-usage-of-2nd-place-97337> (accessed 11 Apr 2014).
- 19 Emarketer, 'Twitter is widely known in France, but garners few regular users', 30 Apr 2013, www.emarketer.com/Article/Twitter-Widely-Known-France-Garners-Few-Regular-Users/1009851 (accessed 11 Apr 2014).
- 20 For a map of current Twitter languages and demographic data, see E Fischer, 'Language communities of Twitter', 24 Oct 2011, www.flickr.com/photos/walkingsf/6277163176/in/photostream/lightbox/ (accessed 10 Apr 2014); DMR, '(March 2014) by the numbers: 138 amazing Twitter statistics', Digital Market Ramblings, 23 Mar 2014, <http://expandedramblings.com/index.php/march-2013-by-the-numbers-a-few-amazing-twitter-stats/> (accessed 10 Apr 2014).
- 21 Slideshare, 'Media measurement: social media trends by age and country', 2011, www.slideshare.net/MML_Annabel/media-measurement-social-media-trends-by-country-and-age (accessed 11 Apr 2014).
- 22 Emarketer, 'Twitter grows stronger in Mexico', 24 Sep 2012, www.emarketer.com/Article/Twitter-Grows-Stronger-Mexico/1009370 (accessed 10 Apr 2014); Inform's Blog, 'Social media: how many people use Twitter and what do we think about it?', International Forum for Responsible Media Blog, 16 Jun 2013, <http://inform.wordpress.com/2013/06/16/social-media-how-many-people-use-twitter-and-what-do-we-think-about-it/> (accessed 11 Apr 2014).

- 23 Eg M Bamburic, 'Twitter: 500 million accounts, billions of tweets, and less than one per cent use their location', 2012, <http://betanews.com/2012/07/31/twitter-500-million-accounts-billions-of-tweets-and-less-than-one-per-cent-use-their-location/> (accessed 11 Apr 2014).
- 24 Beevolve, 'Global heatmap of Twitter users', 2012, www.beevolve.com/twitter-statistics/#a3 (accessed 11 Apr 2014).
- 25 European Commission, 'Political participation and EU citizenship: perceptions and behaviours of young people', nd, <http://eacea.ec.europa.eu/youth/tools/documents/perception-behaviours.pdf> (accessed 11 Apr 2014).
- 26 S Creasey, 'Perceptual engagement: the potential and pitfalls of using social media for political campaigning', London School of Economics, 2011, <http://blogs.lse.ac.uk/polis/files/2011/06/PERPETUAL-ENGAGEMENT-THE-POTENTIAL-AND-PITFALLS-OF-USING-SOCIAL-MEDIA-FOR-POLITICAL-CAMPAIGNING.pdf> (accessed 29 Apr 2014).
- 27 WH Dutton and G Blank, *Next Generation Users: The internet in Britain*, Oxford Internet Survey 2011 report, 2011, www.oii.ox.ac.uk/publications/oxis2011_report.pdf (accessed 3 Apr 2013).
- 28 Ibid.
- 29 J Bartlett et al, *Virtually Members: The Facebook and Twitter followers of UK political parties*, London: Demos 2013.
- 30 J Bartlett et al, *New Political Actors in Europe: Beppe Grillo and the M5S*, London: Demos, 2012; J Birdwell and J Bartlett, *Populism in Europe: CasaPound*, London: Demos, 2012; J Bartlett, J Birdwell and M Littler, *The New Face of Digital Populism*, London: Demos, 2011.

- 31 C McPhedran, 'Pirate Party makes noise in German politics', *Washington Times*, 10 May 2012, www.washingtontimes.com/news/2012/may/10/upstart-party-making-noise-in-german-politics/?page=all (accessed 11 Apr 2014).
- 32 T Postmes and S Brunsting, 'Collective action in the age of the internet: mass communication and online mobilization', *Social Science Computer Review* 20, issue 3, 2002; M Castells, 'The mobile civil society: social movements, political power and communication networks' in M Castells et al, *Mobile Communication and Society: A global perspective*, Cambridge MA: MIT Press, 2007.
- 33 G Blakeley, 'Los Indignados: a movement that is here to stay', *Open Democracy*, 5 Oct 2012, www.opendemocracy.net/georgina-blakeley/los-indignados-movement-that-is-here-to-stay (accessed 11 Apr 2014).
- 34 N Vallina-Rodriguez et al, 'Los Twindignados: the rise of the Indignados Movement on Twitter', in *Privacy, Security, Risk and Trust (PASSAT)*, 2012 International Conference on Social Computing (SocialCom), www.cl.cam.ac.uk/~nv240/papers/twindignados.pdf (accessed 11 Apr 2014).
- 35 GT Madonna and M Young, 'The first political poll', *Politically Uncorrected*, 18 Jun 2002, www.fandm.edu/politics/politically-uncorrected-column/2002-politically-uncorrected/the-first-political-poll (accessed 11 Apr 2014).
- 36 For example, are federal expenditures for relief and recovery too great, too little, or about right? Responses were as follows: 60 per cent too great; 9 per cent too little; 31 per cent about right. See '75 years ago, the first Gallup Poll', *Polling Matters*, 20 Oct 2010, <http://pollingmatters.gallup.com/2010/10/75-years-ago-first-gallup-poll.html> (accessed 11 Apr 2014).

- 37 Thereby avoiding a number of measurement biases often present during direct solicitation of social information, including memory bias, questioner bias and social acceptability bias. Social media, by contrast, is often a completely unmediated spectacle.
- 38 VM Schonberger and K Cukier, *Big Data*, London: John Murray, 2013.
- 39 Early and emerging examples of Twitterology were presented at the International Conference on Web Search and Data Mining 2008. It is important to note that there is a large difference between what are *current* capabilities, and what are *published* capabilities. We do not have access to a great deal of use-cases – including novel techniques, novel applications of techniques or substantive findings – that are either under development or extant but unpublished. Academic peer-reviewed publishing can take anywhere from six months to two years, while many commercial capabilities are proprietary. Furthermore, much social media research is conducted either by or on behalf of the social media platforms themselves, and never made public. The growing distance between development and publishing, and the increasing role of proprietary methodologies and private sector ownership and exploitation of focal data sets, are important characteristics of the social media research environment. Good examples include P Carvalho et al, ‘Liars and saviors in a sentiment annotated corpus of comments to political debates’ in *Proceedings of the Association for Computational Linguistics*, 2011, pp 564–68; N Diakopoulos and D Shamma, ‘Characterising debate performance via aggregated Twitter sentiment’ in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2010, pp 1195–8; S Gonzalez-Bailon, R Banchs and A Kaltenbrunner, ‘Emotional reactions and the pulse of public opinion: measuring the impact of political events on the sentiment of online discussions’, ArXiv e-prints, 2010, arXiv 1009.4019; G Huwang et al, ‘Conversational tagging in

Twitter’ in *Proceedings of the 21st ACM conference on Hypertext and Hypermedia*, 2010, pp 173–8; M Marchetti-Bowick and N Chambers, ‘Learning for microblogs with distant supervision: political forecasting with Twitter’ in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, 2012, pp 603–12; B O’Connor et al, ‘From tweets to polls: linking text sentiment to public opinion time series’ in *Proceedings of the AAAI Conference on Weblogs and Social Media*, 2010, pp 122–9; A Pak and P Paroubak, ‘Twitter as a corpus for sentiment analysis and opinion mining’ in *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, 2010; C Tan et al, ‘User-level sentiment analysis incorporating social networks’ in *Proceedings of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2011; A Tumasjan et al, ‘Election forecasts with Twitter: how 140 characters reflect the political landscape’, *Social Science Computer Review*, 2010. See also RE Wilson, SD Gosling and LT Graham, ‘A review of Facebook research in the social sciences’, *Perspectives on Psychological Science*, 7, no 3, 2012 pp 203–20.

- 40 Early and emerging examples of Twitterology were presented at the International Conference on Web Search and Data Mining, 2008.
- 41 European Commission, ‘Europeans and their languages’, *Special Eurobarometer* 243, Feb 2006, http://ec.europa.eu/public_opinion/archives/ebs/ebs_243_sum_en.pdf (accessed 11 Apr 2014).
- 42 It is also possible to acquire a large amount of social media data via licensed data providers. These are often third party resellers.
- 43 Some APIs can deliver historical data, stretching back months or years, while others only deliver very recent content. Some deliver a random selection of social media

data taken from the platform; others deliver all data that match the queries – usually keywords selected by the analyst to be present in the post or tweet – provided by the researcher. In general, all APIs produce data in a consistent, ‘structured’ format, in large quantities.

- 44 Twitter has three different APIs available to researchers. The search API returns a collection of relevant tweets matching a specified query (word match) from an index that extends up to roughly a week in the past. Its filter API continually produces tweets that contain one of a number of keywords to the researcher, in real time as they are made. Its sample API returns a random sample of a fixed percentage of all public tweets in real time. Each of these APIs (consistent with the vast majority of all social media platform APIs) is constrained by the amount of data they will return. A public, free ‘spritzer’ account caps the search API at 180 calls every 15 minutes with up to 100 tweets returned per call; the filter API caps the number of matching tweets returned by the filter to no more than 1 per cent of the total stream in any given second; and the sample API returns a random 1 per cent of the tweet stream. Others use white-listed research accounts (known informally as ‘the garden hose’), which have 10 per cent rather than 1 per cent caps on the filter and sample APIs, while still others use the commercially available ‘firehose’ of 100 per cent of daily tweets. With daily tweet volumes averaging roughly 400 million, many researchers do not find the spritzer account restrictions to be limiting to the number of tweets they collect (or need) on any particular topic.
- 45 S Fodden, ‘Anatomy of a tweet: metadata on Twitter’, Slaw, 17 Nov 2011, www.slaw.ca/2011/11/17/the-anatomy-of-a-tweet-metadata-on-twitter/ (accessed 11 Apr 2014); R Krikorian, ‘Map of a Twitter status object’, 18 Apr 2010, www.slaw.ca/wp-content/uploads/2011/11/map-of-a-tweet-copy.pdf (accessed 11 Apr 2014).

- 46 Acquiring data from Twitter on a particular topic is a trade-off between precision and comprehensiveness. A precise data collection strategy only returns tweets that are on-topic, but is likely to miss some. A comprehensive data collection strategy collects all the tweets that are on-topic, but is likely to include some which are off-topic. Individual words themselves can be inherently either precise or comprehensive, depending on how and when they are used.
- 47 Ibid.
- 48 The choice of these keywords and hashtags for each topic in each language was made in a quick manual review of the data collected in the early stages of the project. The inclusion of these terms was meant to bring in conversations that were relevant to the stream but did not explicitly reference the topic by its full name, without overwhelming the streams with irrelevant data. For a full list of scraper terms used per stream, see the annex.
- 49 AoIR, *Ethical Decision-Making and Internet Research*; J Bartlett and C Miller, 'How to measure and manage harms to privacy when accessing and using communications data', submission by the Centre for the Analysis of Social Media, as requested by the Joint Parliamentary Select Committee on the Draft Communications Data Bill, Oct 2012, www.demos.co.uk/files/Demos%20CASM%20submission%20on%20Draft%20Communications%20Data%20bill.pdf (accessed 11 Apr 2014).
- 50 It may also partly be a reflection of the network affect of social networks. For example, given the high proportion of tweets in English on Twitter, non-English users may also feel compelled to use English as well to take part in conversations on the network.

- 51 Emarketer, 'Twitter grows stronger in Mexico'; Inform's Blog, 'Social media'.
- 52 Given the historical nature of our data set, each twitcident was identified from a single data stream, rather than across Twitter as a whole (which would be a far better way of collecting data relating to an event). See discussion in chapter 4.
- 53 'Dix milliards d'euros pour sauver Chypre', *Libération*, 16 Mar 2013, www.liberation.fr/economie/2013/03/16/chypre-cinquieme-pays-de-la-zone-euro-a-beneficier-de-l-aide-internationale_889016 (accessed 11 Apr 2014); I de Foucaud, 'Chypre: un sauvetage inédit à 10 milliards d'euros', *Le Figaro*, 16 Mar 2013, www.lefigaro.fr/conjoncture/2013/03/16/20002-20130316ARTFIG00293-chypre-un-sauvetage-inedit-a-10-milliards-d-euros.php (accessed 11 Apr 2014); 'A Chypre, la population sous le choc, le président justifie les sacrifices', *Le Monde*, 17 Mar 2014, www.lemonde.fr/europe/article/2013/03/16/a-chypre-la-population-dans-l-incertitude-apres-l-annonce-du-plan-de-sauvetage_1849491_3214.html (accessed 11 Apr 2014).
- 54 'Hitting the savers: Eurozone reaches deal on Cyprus bailout', *Spiegel International*, 16 Mar 2013, www.spiegel.de/international/europe/savers-will-be-hit-as-part-of-deal-to-bail-out-cyprus-a-889252.html (accessed 11 Apr 2014)
- 55 This echoed much of the early press coverage, especially in Germany, with the *Frankfurter Allgemeine* stating 'Zyperns Rettung Diesmal bluten die Sparer' ('Cyprus rescue bleeding time savers')
- 56 'Does the bailout deal mean the worst is over for Cyprus? – poll', *Guardian*, 25 Mar 2013, www.theguardian.com/business/poll/2013/mar/25/bailout-deal-worst-over-cyprus-poll (accessed 11 Apr 2014).

- 57 Pew Research Center, *The New Sick Man of Europe: The European Union*, 2013, www.pewglobal.org/files/2013/05/Pew-Research-Center-Global-Attitudes-Project-European-Union-Report-FINAL-FOR-PRINT-May-13-2013.pdf (accessed 11 Apr 2014).
- 58 YouGov survey results, fieldwork 21–27 Mar 2013, http://d25d2506sfb94s.cloudfront.net/cumulus_uploads/document/eh65gpseiv/YG-Archive_Eurotrack-March-Cyprus-EU-representatives-Easter.pdf (accessed 11 Apr 2014).
- 59 From a background level of 117 tweets on 9 March 2013, 141 on 10 March and 288 on 11 March, there is an increase to 391 on 12 March, 844 on 13 March and a peak of 1,786 on 14 March.
- 60 See the section ‘classifier performance’ in the annex for a discussion of its accuracy.
- 61 ‘Affichette “casse-toi pov” con”: la France condamnée par la CEDH’, *Le Monde*, 14 Mar 2013, www.lemonde.fr/societe/article/2013/03/14/affichette-casse-toi-pov-con-la-france-condamnee-par-la-cedh_1847686_3224.html (accessed 11 Apr 2014).
- 62 European Court of Human Rights, ‘*Affaire Eon c. France*’, requête 26118/10, 14 Mar 2013, [http://hudoc.echr.coe.int/sites/fra/pages/search.aspx?i=001-117137#{'itemid': \['001-117137'\]}](http://hudoc.echr.coe.int/sites/fra/pages/search.aspx?i=001-117137#{'itemid': ['001-117137']}) (accessed 11 Apr 2014).
- 63 W Jordan, ‘Public: ignore courts and deport Qatada’, YouGov, 26 Apr 2013, <http://yougov.co.uk/news/2013/04/26/brits-ignore-courts-and-deport-qatada/> (accessed 24 Apr 2014).

- 64 Ipsos MORI, ‘Public blamed ECHR over the Home Secretary for Qatada delays’, 26 Apr 2013, www.ipsos-mori.com/researchpublications/researcharchive/2964/Public-blamed-ECHR-over-the-Home-Secretary-for-Abu-Qatada-delays.aspx (accessed 24 Apr 2014).
- 65 Average calculated across March, April and May.
- 66 ‘Récession: “La situation est grave”, juge Hollande’.
- 67 ‘Barroso: “la France doit présenter des reformes crédibles”’.
- 68 ‘José Manuel Barroso: “Être contre la mondialisation, c’est cracher contre le vent”’.
- 69 ‘Hollande ne va pas passer un “examen” à Bruxelles, souligne Barroso’.
- 70 ‘François Hollande au révélateur de la Commission européenne: le président de la République a rencontré les 27 commissaires européens à Bruxelles pour évoquer les réformes structurelles réclamées à la France’.
- 71 This was the highest performing classifier trained during the project – with a far higher accuracy than the generic attitudinal classifiers that attempted to make more generic decisions over a longer term.
- 72 See, for instance, O’Connor et al, ‘From tweets to polls’. The authors collected their sample using just a few keyword searches. Some more promisingly methodical approaches also exist: see J Leskovec, J Kleinberg and C Faloutsos, ‘Grapphevolution: densification and shrinking diameters’, *Data* 1, no 1, Mar 2007, www.cs.cmu.edu/~jure/pubs/powergrowth-tkdd.pdf (accessed 16 Apr 2012); J Leskovec and C Faloutsos, ‘Sampling from large graphs’ in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006,

- www.stat.cmu.edu/~fienberg/Stat36-835/Leskovec-sampling-kdd06.pdf (accessed 17 Apr 2012); P Rusmevichientong et al, 'Methods for sampling pages uniformly from the world wide web' in *Proceedings of the AAAI Fall Symposium on Using Uncertainty Within Computation*, 2001, pp 121–8.
- 73 D Singer, 'Forget the 80/20 principle, with Twitter it is 79/7', *Social Media Today*, 25 Feb 2010, <http://socialmediatoday.com/index.php?q=SMC/177538> (accessed 11 Apr 2014).
- 74 European Union, 'Twitter accounts', nd, http://europa.eu/contact/take-part/twitter/index_en.htm (accessed 11 Apr 2014).
- 75 S Bennett, 'Who uses Twitter? Young, affluent, educated non-white males, suggests data [study]', *All Twitter*, 6 Aug 2013, www.mediabistro.com/alltwitter/twitter-users-2013_b47437 (accessed 11 Apr 2014).
- 76 M Bulmer, 'Facts, concepts, theories and problems' in M Bulmer (ed.), *Sociological Research Methods: An introduction*, London: Macmillan, 1984.
- 77 Surveys often tap attitudes by using a sophisticated barrage of different indicators and different ways of measuring them. The Likert scale measures intensity of feelings (usually measured on a scale from 1 to 5) on a number of different specific questions to gauge an underlying attitude. A body of work around question design has produced settled dos and don'ts aimed at avoiding the unreliable measurement of attitudinal indicators. Questions are avoided if they are too long, ambiguous, leading, general, technical or unbalanced, and many surveys use specific wordings of questions drawn from 'question banks' designed to best practice standards for use by major surveys.

- 78 S Jeffares, ‘Coding policy tweets’, paper presented to the social text analysis workshop, University of Birmingham, 28 Mar 2012.
- 79 S Wibberley and C Miller, ‘Detecting events from Twitter: situational awareness in the age of social media’ in C Hobbs, M Matthew and D Salisbury, *Open Source Intelligence in the Twenty-first Century: New approaches and opportunities*, Palgrave MacMillan, forthcoming 2014.
- 80 Glaser and Strauss, *The Discovery of Grounded Theory*.
- 81 COSMOS platform.
- 82 Open Knowledge Foundation, ‘Open data – an introduction’, nd, <http://okfn.org/opendata/> (accessed 11 Apr 2014).
- 83 The choice of these keywords and hashtags for each topic in each language was made on the basis of a quick manual review of the data that were collected in the early stages of the project. The inclusion of these terms was meant to bring in conversations that were relevant to the stream but did not explicitly reference the topic by its full name, without overwhelming the streams with irrelevant data. For a full list of scraper terms used per stream see annex.
- 84 Marchetti-Bowick and Chambers, ‘Learning for microblogs with distant supervision’; O’Connor et al, ‘From tweets to polls’.
- 85 Method51 is a software suite developed by the project team over the last 18 months. It is based on an open source project called DUALIST. See B Settles, ‘Closing the loop: fast, interactive semi-supervised annotation with queries on features and instances’, *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2011, pp 1467–78. Method51 enables non-technical analysts to build machine-learning classifiers. The most important feature is the speed wherein accurate classifiers can be built.

Classically, an NLP algorithm would require many thousands of examples of ‘marked-up’ tweets to achieve reasonable accuracy. This is expensive and takes days to complete. However, DUALIST innovatively uses ‘active learning’ (an application of information theory that can identify pieces of text that the NLP algorithm would learn most from) and semi-supervised learning (an approach to learning that not only learns from manually labelled data, but also exploits patterns in large unlabelled data sets). This radically reduces the number of marked-up examples from many thousands to a few hundred. Overall, in allowing social scientists to build and evaluate classifiers quickly, and therefore to engage directly with big social media data sets, the AAF makes possible the methodology used in this project.

- 86 On the one hand, we have been fairly inclusive on the relevancy level, in that all discussions on something directly related to the topic were usually included as relevant. For example, for the European Parliament stream, all tweets about individual MEPs were considered relevant, as were tweets about individual commissioners for the European Commission stream. Similarly, for the European Court of Human Rights stream, tweets about the European Convention of Human Rights, on which the Court’s jurisdiction is based, were included. Anything about the management of the euro by the Eurozone countries and the European Central Bank, as well as euro-induced austerity, was considered relevant for the euro stream. On the other hand, some tweets that directly referred to the stream topic were considered irrelevant, because they did not match our criteria of interest in the six streams as they relate to the European project. For example, tweets that referred to the European Union purely as a geographical area, as a shorthand for a group of countries, without referring in any sense to this group of countries as belonging to a political union, were marked as irrelevant (eg ‘Car sales in the EU have gone down 20 per cent’). Similarly, tweets referring to the euro from

a purely financial perspective, quoting solely the price of things in euros or exchange rates, were irrelevant.

- 87 For example, for the European Parliament stream, tweets that expressed an opinion about its decisions, discussions taking place in the Parliament, individual MEPs and ‘lobbying’ directed at it (eg ‘@EP: please outlaw pesticides and save the bees!’) were considered attitudinal.
- 88 For the Parliament and Commission streams, positive or negative comments on individual MEPs and commissioners and specific decisions taken by each institution were marked as such.
- 89 The harmonic mean of p and r is equal to $p \times q, p + q$.
- 90 Wilson et al, ‘A review of Facebook research in the social sciences’.
- 91 M Madden, ‘Privacy management on social media sites’, Pew Research Center, 2012, www.pewinternet.org/~media//Files/Reports/2012/PIP_Privacy_management_on_social_media_sites_022412.pdf (accessed 11 Apr 2014).
- 92 J Bartlett, *The Data Dialogue*, London: Demos, 2010. Also see J Bartlett and C Miller, *Demos CASM Submission to the Joint Committee on the Draft Communications Data Bill*, Demos, 2012
- 93 Bartlett, *The Data Dialogue*. This is based on a representative population level poll of circa 5,000 people. See also Bartlett and Miller, ‘How to measure and manage harms to privacy when accessing and using communications data’.
- 94 Twitter, ‘Terms of Service’, 2012, www.twitter.com/tos (accessed 11 Apr 2014); Twitter, ‘Twitter Privacy Policy’, 2013, www.twitter.com/privacy (accessed 11 Apr 2014).
- 95 Twitter, ‘Terms of Service’.

References

‘75 years ago, the first Gallup Poll’, *Polling Matters*, 20 Oct 2010, <http://pollingmatters.gallup.com/2010/10/75-years-ago-first-gallup-poll.html> (accessed 11 Apr 2014).

‘A Chypre, la population sous le choc, le président justifie les sacrifices’, *Le Monde*, 17 Mar 2014, www.lemonde.fr/europe/article/2013/03/16/a-chypre-la-population-dans-l-incertitude-apres-l-annonce-du-plan-de-sauvetage_1849491_3214.html (accessed 11 Apr 2014).

‘Affichette “casse-toi pov’ con”: la France condamnée par la CEDH’, *Le Monde*, 14 Mar 2013, www.lemonde.fr/societe/article/2013/03/14/affichette-casse-toi-pov-con-la-france-condamnee-par-la-cedh_1847686_3224.html (accessed 11 Apr 2014).

‘Dix milliards d’euros pour sauver Chypre’, *Libération*, 16 Mar 2013, www.liberation.fr/economie/2013/03/16/chypre-cinquieme-pays-de-la-zone-euro-a-beneficier-de-l-aide-internationale_889016 (accessed 11 Apr 2014).

‘Does the bailout deal mean the worst is over for Cyprus? – poll’, *Guardian*, 25 Mar 2013, www.theguardian.com/business/poll/2013/mar/25/bailout-deal-worst-over-cyprus-poll (accessed 11 Apr 2014).

‘Hitting the savers: Eurozone reaches deal on Cyprus bailout’, *Spiegel International*, 16 Mar 2013, www.spiegel.de/international/europe/savers-will-be-hit-as-part-of-deal-to-bail-out-cyprus-a-889252.html (accessed 39 Apr 2014).

AoIR, *Ethical Decision-Making and Internet Research: Recommendations from the AoIR Ethics Working Committee (Version 2.0)*, Association of Internet Researchers, 2012, p 2.

Bamburic M, 'Twitter: 500 million accounts, billions of tweets, and less than one per cent use their location', 2012, <http://betanews.com/2012/07/31/twitter-500-million-accounts-billions-of-tweets-and-less-than-one-per-cent-use-their-location/> (accessed 11 Apr 2014).

Bartlett J, *The Data Dialogue*, London: Demos, 2010.

Bartlett J, Birdwell J and Littler M, *The New Face of Digital Populism*, London: Demos, 2011.

Bartlett J and Miller C, 'How to measure and manage harms to privacy when accessing and using communications data', submission by the Centre for the Analysis of Social Media, as requested by the Joint Parliamentary Select Committee on the Draft Communications Data Bill, Oct 2012, www.demos.co.uk/files/Demos%20CASM%20submission%20on%20Draft%20Communications%20Data%20bill.pdf (accessed 11 Apr 2014).

Bartlett J et al, *New Political Actors in Europe: Beppe Grillo and the M5S*, London: Demos, 2012.

Bartlett J et al, *Virtually Members: The Facebook and Twitter followers of UK political parties*, London: Demos 2013.

Beevolve, 'Global heatmap of Twitter users', 2012, www.beevolve.com/twitter-statistics/#a3 (accessed 11 Apr 2014).

Bennett S, 'Who uses Twitter? Young, affluent, educated non-white males, suggests data [study]', All Twitter, 6 Aug 2013, www.mediabistro.com/alltwitter/twitter-users-2013_b47437 (accessed 11 Apr 2014).

Birdwell J and Bartlett J, *Populism in Europe: CasaPound*, London: Demos, 2012.

Birdwell J, Farook F and Jones S, *Trust in Practice*, London: Demos, 2009.

Blakeley G, 'Los Indignados: a movement that is here to stay', Open Democracy, 5 Oct 2012, www.opendemocracy.net/georgina-blakeley/los-indignados-movement-that-is-here-to-stay (accessed 11 Apr 2014).

Bulmer M, 'Facts, concepts, theories and problems' in M Bulmer (ed.), *Sociological Research Methods: An introduction*, London: Macmillan, 1984.

Carvalho P et al, 'Liars and saviors in a sentiment annotated corpus of comments to political debates' in *Proceedings of the Association for Computational Linguistics*, 2011, pp 564–68.

Castells M, 'The mobile civil society: social movements, political power and communication networks' in M Castells et al, *Mobile Communication and Society: A global perspective*, Cambridge MA: MIT Press, 2007.

Creasey S, 'Perceptual engagement: the potential and pitfalls of using social media for political campaigning', London School of Economics, 2011, <http://blogs.lse.ac.uk/polis/files/2011/06/PERPETUAL-ENGAGEMENT-THE-POTENTIAL-AND-PITFALLS-OF-USING-SOCIAL-MEDIA-FOR-POLITICAL-CAMPAIGNING.pdf> (accessed 29 Apr 2014).

de Foucaud I, 'Chypre: un sauvetage inédit à 10 milliards d'euros', *Le Figaro*, 16 Mar 2013, www.lefigaro.fr/conjoncture/2013/03/16/20002-20130316ARTFIG00293-chypre-un-sauvetage-inedit-a-10-milliards-d-euros.php (accessed 11 Apr 2014).

Diakopoulos N and Shammar D, 'Characterising debate performance via aggregated Twitter sentiment' in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2010, pp 1195–8.

References

DMR, '(March 2014) by the numbers: 138 amazing Twitter statistics', Digital Market Ramblings, 23 Mar 2014, <http://expandedramblings.com/index.php/march-2013-by-the-numbers-a-few-amazing-twitter-stats/> (accessed 10 Apr 2014).

Dutton WH and Blank G, *Next Generation Users: The internet in Britain*, Oxford Internet Survey 2011 report, 2011, www.oii.ox.ac.uk/publications/oxis2011_report.pdf (accessed 3 Apr 2013).

Emarketer, 'Twitter grows stronger in Mexico', 24 Sep 2012, www.emarketer.com/Article/Twitter-Grows-Stronger-Mexico/1009370 (accessed 10 Apr 2014).

Emarketer, 'Twitter is widely known in France, but garners few regular users', 30 Apr 2013, www.emarketer.com/Article/Twitter-Widely-Known-France-Garners-Few-Regular-Users/1009851 (accessed 11 Apr 2014).

Emarketer, 'Where in the world are the hottest social networking countries?', 29 Feb 2012, www.emarketer.com/Article/Where-World-Hottest-Social-Networking-Countries/1008870 (accessed 11 Apr 2014).

ESRC, 'Framework for Research Ethics', latest version, Economic and Social Research Council, Sep 2012, www.esrc.ac.uk/about-esrc/information/research-ethics.aspx (accessed 13 Apr 2014).

European Commission, Eurobarometer survey on trust in institutions, Nov 2013, http://ec.europa.eu/public_opinion/cf/showchart_column.cfm?keyID=2189&nationID=6,3,15,&startdate=2012.05&enddate=2013.11 (accessed 24 Apr 2014).

European Commission, 'Europeans and their languages', *Special Eurobarometer* 243, Feb 2006, http://ec.europa.eu/public_opinion/archives/ebs/ebs_243_sum_en.pdf (accessed 11 Apr 2014).

European Commission, 'Political participation and EU citizenship: perceptions and behaviours of young people', nd, <http://eacea.ec.europa.eu/youth/tools/documents/perception-behaviours.pdf> (accessed 11 Apr 2014).

European Commission, 'Public opinion in the European Union: first results', *Standard Eurobarometer 78*, Dec 2012, http://ec.europa.eu/public_opinion/archives/eb/eb78/eb78_first_en.pdf (accessed 10 Apr 2014).

European Commission, 'Two years to go to the 2014 European elections', *Eurobarometer 77*, no 4, 2012, www.europarl.europa.eu/pdf/eurobarometre/2012/election_2012/eb77_4_ee2014_synthese_analytique_en.pdf (accessed 11 Apr 2014).

European Court of Human Rights, '*Affaire Eon c. France*', requête 26118/10, 14 Mar 2013, [http://hudoc.echr.coe.int/sites/fra/pages/search.aspx?i=001-117137#{'itemid':\['001-117137'\]}](http://hudoc.echr.coe.int/sites/fra/pages/search.aspx?i=001-117137#{'itemid':['001-117137']}) (accessed 11 Apr 2014).

European Union, 'Twitter accounts', nd, http://europa.eu/contact/take-part/twitter/index_en.htm (accessed 11 Apr 2014).

Fischer E, 'Language communities of Twitter', 24 Oct 2011, www.flickr.com/photos/walkingsf/6277163176/in/photostream/lightbox/ (accessed 10 Apr 2014).

Fodden S, 'Anatomy of a tweet: metadata on Twitter', Slaw, 17 Nov 2011, www.slaw.ca/2011/11/17/the-anatomy-of-a-tweet-metadata-on-twitter/ (accessed 11 Apr 2014).

Ginn F, 'Global social network stats confirm Facebook as largest in US & Europe (with 3 times the usage of 2nd place)', Search Engine Land, 17 Oct 2011, <http://searchengineland.com/global-social-network-stats-confirm-facebook-as-largest-in-u-s-europe-with-3-times-the-usage-of-2nd-place-97337> (accessed 11 Apr 2014).

Glaser BG and Strauss AL, *The Discovery of Grounded Theory*, New Brunswick: AldineTransaction, 1967.

Gonzalez-Bailon S, Banchs R and Kaltenbrunner A, 'Emotional reactions and the pulse of public opinion: measuring the impact of political events on the sentiment of online discussions', ArXiv e-prints, 2010, arXiv 1009.4019.

Henn M and Foard N, 'Young people, political participation and trust in Britain', *Parliamentary Affairs* 65, no 1, 2012.

Heppner PP, Wampold BE and Kivlighan DM, *Research Design in Counseling*, Thompson, 2008, p 331.

Huwang G et al, 'Conversational tagging in Twitter' in *Proceedings of the 21st ACM conference on Hypertext and Hypermedia*, 2010, pp 173–8.

Huyst P, 'The Europeans of tomorrow: researching European identity among young Europeans', Centre for EU Studies, Ghent University, nd, http://aei.pitt.edu/33069/1/huyst._petra.pdf (accessed 11 Apr 2014).

Inform's Blog, 'Social media: how many people use Twitter and what do we think about it?', International Forum for Responsible Media Blog, 16 Jun 2013, <http://inform.wordpress.com/2013/06/16/social-media-how-many-people-use-twitter-and-what-do-we-think-about-it/> (accessed 11 Apr 2014).

Ipsos MORI, 'Public blamed ECHR over the Home Secretary for Qatada delays', 26 Apr 2013, www.ipsos-mori.com/researchpublications/researcharchive/2964/Public-blamed-ECHR-over-the-Home-Secretary-for-Abu-Qatada-delays.aspx (accessed 24 Apr 2014).

Jeffares S, 'Coding policy tweets', paper presented to the social text analysis workshop, University of Birmingham, 28 Mar 2012.

Jordan W, 'Public: ignore courts and deport Qatada', YouGov, 26 Apr 2013, <http://yougov.co.uk/news/2013/04/26/brits-ignore-courts-and-deport-qatada/> (accessed 24 Apr 2014).

Krikorian R, 'Map of a Twitter status object', 18 Apr 2010, www.slaw.ca/wp-content/uploads/2011/11/map-of-a-tweet-copy.pdf (accessed 11 Apr 2014).

Leskovec J and Faloutsos C, 'Sampling from large graphs' in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006, www.stat.cmu.edu/~fienberg/Stat36-835/Leskovec-sampling-kddo6.pdf (accessed 17 Apr 2012).

Leskovec J, Kleinberg J and Faloutsos C, 'Grapphevolution: densification and shrinking diameters', *Data* 1, no 1, Mar 2007, www.cs.cmu.edu/~jure/pubs/powergrowth-tkdd.pdf (accessed 16 Apr 2012).

Madden M, 'Privacy management on social media sites', Pew Research Center, 2012, www.pewinternet.org/~media//Files/Reports/2012/PIP_Privacy_management_on_social_media_sites_022412.pdf (accessed 11 Apr 2014).

Madonna GT and Young M, 'The first political poll', *Politically Uncorrected*, 18 Jun 2002, www.fandm.edu/politics/politically-uncorrected-column/2002-politically-uncorrected/the-first-political-poll (accessed 11 Apr 2014).

Marchetti-Bowick M and Chambers N, 'Learning for microblogs with distant supervision: political forecasting with Twitter' in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, 2012, pp 603–12.

McPhedran C, 'Pirate Party makes noise in German politics', *Washington Times*, 10 May 2012, www.washingtontimes.com/news/2012/may/10/upstart-party-making-noise-in-german-politics/?page=all (accessed 11 Apr 2014).

- O'Connor B et al, 'From tweets to polls: linking text sentiment to public opinion time series' in *Proceedings of the AAAI Conference on Weblogs and Social Media*, 2010, pp 122–9.
- Open Knowledge Foundation, 'Open data – an introduction', nd, <http://okfn.org/opendata/> (accessed 11 Apr 2014).
- Pak A and Paroubak P, 'Twitter as a corpus for sentiment analysis and opinion mining' in *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, 2010.
- Pew Research Center, *The New Sick Man of Europe: The European Union*, 2013, www.pewglobal.org/files/2013/05/Pew-Research-Center-Global-Attitudes-Project-European-Union-Report-FINAL-FOR-PRINT-May-13-2013.pdf (accessed 11 Apr 2014).
- Pew Research Center, 'The sick man of Europe: the European Union', 13 May 2013, www.pewglobal.org/2013/05/13/the-new-sick-man-of-europe-the-european-union/ (accessed 10 Apr 2014).
- Postmes T and Brunsting S, 'Collective action in the age of the internet: mass communication and online mobilization', *Social Science Computer Review* 20, issue 3, 2002.
- Rusmevichientong P et al, 'Methods for sampling pages uniformly from the world wide web' in *Proceedings of the AAAI Fall Symposium on Using Uncertainty Within Computation*, 2001, pp 121–8.
- Schonberger VM and Cukier K, *Big Data*, London: John Murray, 2013.
- Settles B, 'Closing the loop: fast, interactive semi-supervised annotation with queries on features and instances', *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2011, pp 1467–78.

Singer D, 'Forget the 80/20 principle, with Twitter it is 79/7', *Social Media Today*, 25 Feb 2010, <http://socialmediatoday.com/index.php?q=SMC/177538> (accessed 11 Apr 2014).

Slideshare, 'Media measurement: social media trends by age and country', 2011, www.slideshare.net/MML_Annabel/media-measurement-social-media-trends-by-country-and-age (accessed 11 Apr 2014).

Sloam J, 'Rebooting democracy: youth participation in politics in the UK', *Parliamentary Affairs*, 60, 2007.

Social-media-prism, 'The conversation', nd, www.google.co.uk/imgres?imgurl=http://spirdesign.no/wp-content/uploads/2010/11/social-media-prism.jpg&imgrefurl=http://spirdesign.no/blog/webdesignidentitet-og-trender/attachment/social-media-prism/&h=958&w=1024&sz=301&tbnid=EFQcS2D-zhOj8M:&tbnh=90&tbnw=96&zooom=1&usg=__VXussUcXEMznT42YLhgk6kOsPIk=&docid=hog_RAXkIYvcpM&sa=X&ei=9QBXUdeYO-iJoAXdyIHYAg&ved=oCEoQ9QEwAg&dur=47 (accessed 11 Apr 2014).

Tan et al C, 'User-level sentiment analysis incorporating social networks' in *Proceedings of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2011.

Tumasjan A et al, 'Election forecasts with Twitter: how 140 characters reflect the political landscape', *Social Science Computer Review*, 2010.

Twitter, 'Terms of Service', 2012, www.twitter.com/tos (accessed 11 Apr 2014).

Twitter, 'Twitter Privacy Policy', 2013, www.twitter.com/privacy (accessed 11 Apr 2014).

References

Vallina-Rodriguez N et al, 'Los Twindignados: the rise of the Indignados Movement on Twitter', in *Privacy, Security, Risk and Trust (PASSAT)*, 2012 International Conference on Social Computing (SocialCom), www.cl.cam.ac.uk/~nv240/papers/twindignados.pdf (accessed 11 Apr 2014).

van Biezen I, Mair P and Poguntke T, 'Going, going... gone? The decline of party membership in contemporary Europe', *European Journal of Political Research* 51, no 1, 2012, pp 24–56.

Wibberley S and Miller C, 'Detecting events from Twitter: situational awareness in the age of social media' in Hobbs C, Matthew M and Salisbury D, *Open Source Intelligence in the Twenty-first Century: New approaches and opportunities*, Palgrave Macmillan, 2014.

Wilson RE, Gosling SD and Graham LT, 'A review of Facebook research in the social sciences', *Perspectives on Psychological Science* 7, no 3, pp 203–20.

Zeng D et al, 'Social media analytics and intelligence: guest editors' introduction', in *Proceedings of the IEEE Computer Society*, Nov–Dec 2010, p 13.

Demos – Licence to Publish

The work (as defined below) is provided under the terms of this licence ('licence'). The work is protected by copyright and/or other applicable law. Any use of the work other than as authorized under this licence is prohibited. By exercising any rights to the work provided here, you accept and agree to be bound by the terms of this licence. Demos grants you the rights contained here in consideration of your acceptance of such terms and conditions.

1 Definitions

- A **'Collective Work'** means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this Licence.
- B **'Derivative Work'** means a work based upon the Work or upon the Work and other pre-existing works, such as a musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work or a translation from English into another language will not be considered a Derivative Work for the purpose of this Licence.
- C **'Licensor'** means the individual or entity that offers the Work under the terms of this Licence.
- D **'Original Author'** means the individual or entity who created the Work.
- E **'Work'** means the copyrightable work of authorship offered under the terms of this Licence.
- F **'You'** means an individual or entity exercising rights under this Licence who has not previously violated the terms of this Licence with respect to the Work, or who has received express permission from Demos to exercise rights under this Licence despite a previous violation.

2 Fair Use Rights

Nothing in this licence is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3 Licence Grant

Subject to the terms and conditions of this Licence, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) licence to exercise the rights in the Work as stated below:

- A to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;
- B to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works; The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensor are hereby reserved.

4 Restrictions

The licence granted in Section 3 above is expressly made subject to and limited by the following restrictions:

- A You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this Licence, and You must include a copy of, or the Uniform Resource Identifier for, this Licence with every copy or phonorecord of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this Licence or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this Licence and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this Licence Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this Licence. If You create a Collective Work, upon notice from any Licensor You must, to the extent practicable, remove from the Collective Work any reference to such Licensor or the Original Author, as requested.
- B You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital

filesharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.

- c If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Collective Works, You must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied. Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5 Representations, Warranties and Disclaimer

- A By offering the Work for public release under this Licence, Licensor represents and warrants that, to the best of Licensor's knowledge after reasonable inquiry:
 - i Licensor has secured all rights in the Work necessary to grant the licence rights hereunder and to permit the lawful exercise of the rights granted hereunder without You having any obligation to pay any royalties, compulsory licence fees, residuals or any other payments;
 - ii The Work does not infringe the copyright, trademark, publicity rights, common law rights or any other right of any third party or constitute defamation, invasion of privacy or other tortious injury to any third party.
- B except as expressly stated in this licence or otherwise agreed in writing or required by applicable law, the work is licenced on an 'as is' basis, without warranties of any kind, either express or implied including, without limitation, any warranties regarding the contents or accuracy of the work.

6 Limitation on Liability

Except to the extent required by applicable law, and except for damages arising from liability to a third party resulting from breach of the warranties in section 5, in no event will licensor be liable to you on any legal theory for any special, incidental, consequential, punitive or exemplary damages arising out of this licence or the use of the work, even if licensor has been advised of the possibility of such damages.

7 Termination

- A This Licence and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this Licence. Individuals or entities who have received Collective Works from You under this Licence, however, will not have their licences terminated provided such individuals or entities remain in full compliance with those licences. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this Licence.
- B Subject to the above terms and conditions, the licence granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different licence terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this Licence (or any other licence that has been, or is required to be, granted under the terms of this Licence), and this Licence will continue in full force and effect unless terminated as stated above.

8 Miscellaneous

- A Each time You distribute or publicly digitally perform the Work or a Collective Work, Demos offers to the recipient a licence to the Work on the same terms and conditions as the licence granted to You under this Licence.
- B If any provision of this Licence is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Licence, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.
- C No term or provision of this Licence shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- D This Licence constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This Licence may not be modified without the mutual written agreement of Demos and You.

This project was supported by



and conducted in partnership with



Over the last decade European citizens have gained a digital voice. Close to 350 million people in Europe currently use social networking sites, with more of us signing into a social media platform at least once a day than voted in the last European elections. EU citizens have transferred many aspects of their lives onto these social media platforms, including politics and activism. Taken together, social media represent a new digital commons where people join their social and political lives to those around them.

This paper examines the potential of listening to these digital voices on Twitter, and the consequences for how EU leaders apprehend, respond to and thereby represent their citizens. It looks at how European citizens use Twitter to discuss issues related to the EU and how their digital attitudes and views evolve in response to political and economic crises. It also addresses the many formidable challenges that this new method faces: how far it can be trusted, when it can be used, the value such use could bring and how its use can be publicly acceptable and ethical.

We have never before had access to the millions of voices that together form society's constant political debate, nor the possibility of understanding them. This report demonstrates how capturing and understanding these citizen voices potentially offers a new way of listening to people, a transformative opportunity to understand what they think, and a crucial opportunity to close the democratic deficit.

Jamie Bartlett is Director of the Centre for the Analysis of Social Media (CASM) at Demos. Carl Miller is Research Director at CASM. David Weir is a Professor of Computer Science at the University of Sussex. Jeremy Reffin and Simon Wibberley are Research Fellows in the Department of Informatics at the University of Sussex.

ISBN 978-1-909037-63-2 £10

© Demos 2014

