THE WISDOM OF SEEKING CROWD WISDOM[1]

Reflections on the ethics of using crowdsourcing and crowdseeding to collect data in conflict zones

Anouk S. Rigterink[2] and Ghassan Baliki[3]

February 2019

Abstract

This paper reflects on the ethical implications of using crowd-based methods to collect information in situations of violent conflict. These methods involve obtaining information directly from individuals in

---

[1] The CCPE Syria project was funded by the ESRC UK, Grant number ES/L014955/1. We are indebted to our collaborators on the CCPE Syria project: Tilman Brück, Mary Kaldor and Rim Turkmani, and most of all to twelve participants and a field coordinator who must remain anonymous. In addition, we are grateful to Tilman Brück, Neil T.N. Ferguson, Denisa Kostovicova, Eleanor Knott, Jack McDonald, Dipali Mukhopadhyay, participants at the International Studies Association 2018 Annual Conference, three anonymous reviewers and our fellow authors of this special issue for helpful comments. All errors remain our own.

[2] University of Oxford. Department of Economics and Blavatnik School of Government. 5 Manor Rd, Oxford OX1 3UQ. Email: anouk.rigterink@economics.ox.ac.uk. Phone: +44(0)186581168.

[3] ISDC - International Security and Development Center gGmbH, Auguststr. 89, 10117 Berlin, Germany

conflict areas, who can choose when or how frequently to report, via some communication technology. Such methods have become mainstream in the humanitarian sector, and present a microcosm of ethical considerations surfacing when doing academic research in conflict areas. The paper proposes a typology of crowd-based projects, distinguishing them using two axes: selectivity in who makes up the 'crowd' and restrictions on access to information gathered. Building on authors' personal experience with a crowdseeding project, and published accounts of other crowd-based projects, this paper argues that the position of a project on these axes largely determines which ethical concerns are likely to arise. Furthermore, existing literature and processes to obtain ethics board approval for a project using crowd-based methods are strongly focussed on certain principles of procedural ethics, especially informed consent and avoiding harm to participants. Without taking away from the acute necessity to consider these aspects of ethics, this paper argues that researchers using crowd-based methods in conflict should also consider other applications of guiding ethical principles, including maximizing benefits and respect for individual autonomy through co-creation of research, and non-procedural dimensions of ethics, such as micro-ethics and "getting it right".

1. Introduction

The humanitarian sector and academic researchers are trying to 'harness the power of the crowd' to gather information on areas in violent conflict. With access to communication technology, 'the crowd' - individuals in conflict areas - could provide on-the-ground information in near real-time. This has given rise to new crowd-based methods for research and practice, including crowdsourcing and crowdseeding.

Crowd-based methods have become mainstream in the humanitarian sector (Butler, 2013) and have made inroads into academic research. The United Nations office coordinating the international response to humanitarian emergencies (UN-OCHA) regularly commissions crowdsourced maps of areas in crisis (Butler, 2013). Some legal scholars make an argument to include crowdsourcing in the legal curriculum, citing instances of crowdsourced information triggering investigation by the International Criminal Court (Poblet, 2013). Some even argue that crowdsourcing can be a tool for conflict prevention (Kahl, Mcconnell, & Tsuma, 2012). Van der Windt and Humphreys (2016) cautiously highlighted benefits of crowdseeding for academic research (Van der Windt & Humphreys, 2016).

Crowd-based methods present numerous ethical challenges (Pham and Vinck 2012; Martin-Shields 2013; Chamales and Baker 2011), which are worth investigating for two reasons. First, because of crowd-based methods' popularity among humanitarian responders. Second, because crowd-based methods present a microcosm of the ethical challenges academic researchers may encounter when using more mainstream research methods in situations of violent conflict. Crowd-based methods require sustained engagement with conflict-affected populations - normally associated with ethnography or longitudinal studies -, raise questions about representativeness of data collected – just like household surveys -, and use information and communication technologies (ICTs) – increasingly common in academic research in general.

This paper reflects on the ethics of using crowd-based methods in violent conflict. By crowd-based methods, we mean methods that rely on individuals residing in conflict areas to report first-hand information through some ICT, where these individuals choose the frequency and timing of reporting.

We argue that projects using crowd-based methods in conflict areas differ from one another on two axes. First, selectivity in who makes up the crowd: some projects put out an open call and allow anyone in conflict-affected areas to participate, whereas others purposefully select participants. Second, restrictions on access to collected information: some projects publish information collected on publicly accessible websites, whereas others tightly control who gets access to information collected.

The argument of this paper is two-fold. First, we argue that the position of a crowd-based project on these two axes largely determines which ethical issues will arise. Projects with no selectivity in participation and few controls on access to data will likely experience greater difficulties obtaining informed consent from participants, and rely on riskier strategies to minimize the risk of harm to participants and researchers. Projects with strong selectivity and strict controls on access to data might find obtaining informed consent less problematic, as well as adopt less risky strategies to minimize harm. Secondly, we note that existing literature on the ethics of crowd-based methods, and the process of obtaining Institutional Review Board (IRB) permission for a study using such methods, strongly emphasizes specific procedural aspects of ethics - particularly informed consent and ‘do no harm’. We argue that although these principles are crucial to ethically using crowd-based methods in conflict zones, doing so *also* requires considering other aspects of ethics. These include other applications of the three main ethical principles guiding research with human subjects (The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979), as well as non-procedural aspects of ethics. We highlight a number of these in this paper, which include: co-creation of research, maximizing benefits, micro-ethics, and “getting it right", the imperative for researchers to live up to the ethical standards of their profession [another paper, this issue].

Several authors went before us in pointing out ethical challenges involved in crowd-based research (Chamales & Baker, 2011; Martin-Shields, 2013; Pham & Vinck, 2012). The above works were published shortly after the advent of the method; a substantial number of new accounts were published since then, so this issue merits revisiting. Also, these papers commonly consider all crowd-based projects as a single category, and seldom highlight ethical considerations other than informed consent and minimizing harm to participants.

This paper therefore contributes to this literature in several ways. First, we consider how distinct types of crowd-based projects require different ethical considerations. As more of these projects are put in practice, it is becoming increasingly clear that there is no one-model-fits-all when it comes to ethics. Second, we revisit but then go beyond the discussion of informed consent and minimizing harm, in order to provide a more nuanced and comprehensive picture of the various ethical challenges encountered. Third, we rely on personal experience and new candid accounts of crowd-based projects to inform our analysis.

The remainder of this paper is organized as follows. Section 2 sets out the main case studies informing this paper. Section 3 sets out a typology of crowd-based projects. Sections 4 discusses the maxims of procedural ethics, whilst Section 5 goes beyond these to highlight other, less discussed applications and dimensions of ethics. Section 6 concludes.

2. Case study

This paper is primarily informed by the authors' personal experience during the project 'Crowdseeding Conflict and Peace Events in the Syrian conflict' (CCPE) [authors, 2017]. The aim of this project was to gather geo-referenced disaggregated data about the Syrian conflict. At the start of the project, the Syrian conflict was not covered by traditional event-based conflict datasets, which rely on the coding of media and secondary sources. The project was also motivated by the concern that media coding could result in biased data [another paper, this issue]. Finally, we wanted to capture and quantify localized peace agreements. Qualitative research had found local peace agreements and cease-fires to be frequent, but these were rarely reported in mainstream media, and not a category of events captured in conventional datasets (Turkmani, Kaldor, Elhamwi, Ayo, & Hariri, 2014).

Research methods that allowed investigating these topics were limited: physically going to Syria posed grave risks to researchers and participants and was strongly opposed by the University's Health and Safety department. We saw crowd-based methods as a viable method to gather systematic data on the Syrian conflict.

Ultimately, the project aimed to increase our understanding of violent conflict in general, and the Syrian conflict in particular. At a minimum, we aimed to be able to assess the potential of crowd-based methods to gather reliable data on violent conflict, possibly in comparison to more traditional methods based on coding of secondary reports.

In practice, we partnered with two Syrian civil society organisations, which we unfortunately cannot name for security reasons. We chose these particular organisations, because they were already independently involved with reporting on the Syrian conflict, and thus well informed of the risks involved in this. These two organisations provided us with a 'crowd' of twelve participants, distributed geographically across cities and towns in Syria. These participants reported instances of violence, looting, kidnapping and detainments, and local peace agreements to us through a web platform. This platform was designed by a non-commercial third party, after extensive consultations with the university's specialist in data security, and did not make use of existing social media (e.g. Facebook, Google forms, Twitter).

Each participant received an individual training session on the use of the platform and best-practice in online security. This session also included a discussion of potential risks of participating.

Participants' first entry on the platform consisted of a consent statement. Participants were paid for their time; payment amount was not linked to amount of information entered into the platform. A representative of the civil society organisation was in regular contact with the participants throughout the project and acted as a go-between between participants and researchers. Access to the data collected was restricted to researchers, the civil society organisation representative and the designers of the platform; participants could not see it. No information that could identify individual participants was exchanged via the platform. During the nine months of data gathering, this project collected 3546 reports of events, including 2639 violent events and 476 'peace events'.

We sought and received approval for the CCPE Syria project from the ethics board of [author's university]. In practice, we also needed to obtain approval from representatives of the university's administration, who needed to sign the contract with the web platform builder, and the Memorandum of Understanding with the Syrian civil society organisation. While deciding whether to sign these, representatives of university administration made reference to ethical treatment of participants, as well as risk to the university itself.

In addition to the CCPE Syria project, this paper is informed by four impressively candid accounts of crowd-based projects: *Ushahidi Kenya*, which documented violence during the Kenyan elections of 2007-2008 (Okolloh, 2009); *Libya Crisis Map,* covering the conflict between the Ghaddafi regime and those seeking to oust it following the Arab Spring. This was commissioned by UN-OCHA but the project was executed by a group of volunteers called the Standby Task Force (Standby Task Force, 2011; Verity, 2011) ; *U-Shahid*, which monitored electoral fraud during Egypt's 2010 parliamentary elections (Meier, 2012); and *Voix des Kivus,* which gathered information about conflict-related instances in the Democratic Republic of Congo (DRC) (van der Windt, 2014; Van der Windt & Humphreys, 2016; Windt & Humphreys, 2012).

3. Crowd-based methods: a typology

3.1. **Commonalities**

There is no definitive definition of what constitutes a crowd-based project. The literature provides numerous definitions of crowdsourcing (see Table 1 for a selection). Projects using crowd-based methods in conflict-affected areas have three things in common however. First, they rely on direct reports from individuals residing in affected areas (the 'crowd'). Note that this is distinct from relying on a 'crowd' of online volunteers *not* resident in the conflict area to process or map reports. The latter, which we will call crowd-tasking for the purpose of this paper, is not necessarily an integral part of crowd-based projects, but was used to create the Libya Crisis Map. Second, the use of ICTs facilitates the transfer of information from the crowd to those collating it ('the researchers' for the purpose of this paper, although not all individuals leading a crowd-based project identify themselves as such). Using ICTs implies using some *technology* (e.g., SMS, websites, internet-based apps), which comes with a distinct set of ethical considerations, and it creates *physical distance* between the researcher and participant by obviating the need for the researcher to be physically present in the research area. Third, individuals making up the crowd consciously choose when or how frequently to report information and can in theory report information in real-time. This creates the opportunity for repeated interaction between researcher and participant, and gives the participant a degree of discretion, somewhat narrowing the *distance* between researcher and participant in terms of *power* (Fujii, 2012). 'Consciously' implies that participants are aware that they provide information to a specific project. Although some projects such as the Libya Crisis Map *also* collate information from publicly accessible social media such as Twitter or YouTube, we exclude this activity from our definition of crowd-based projects.

Table 1: Selected definitions of 'crowdsourcing'

Definition of ‘crowdsourcing’	Author
The process of getting the general public (the crowd) to be a source of specific resources	(Roberts & Marchais, 2017)
The practice of harnessing collective action and capability to achieve a particular task	(Perera, 2017)
Using a large amount of people to report a story	(Okolloh, 2009)
A combination of micro-tasking and citizen reporting	(Munro, 2013)
A data collection method that encourages the “crowd” (the public) to provide answers to questions or content for consumption through open transmission systems like mobile phones, social media and websites	(Martin-Shields, 2013)
Using technology to receive specific information from a crowd of people related to targeted topics or issues and then allowing for this information to be geographically mapped	(Kahl et al., 2012)

3.2. Axes of difference

Despite the above-mentioned similarities, crowd-based projects in violent conflict differ on two important aspects: restrictions to access to the data collected and the selection of the crowd. We illustrate these as two ‘axes of difference’, dividing the universe of crowd-based projects in four different quadrants (see Figure I).

Considering restrictions on access to data collected, at one extreme there are no such restrictions and data is open to all. The Ushahidi and U’Shahid projects made information available for all to access online, after removing any personally identifying details of individuals (Meier 2012; Okolloh 2009). These projects often create visually appealing maps which online viewers can tailor. At the other extreme, only researchers have access to the data. Access to data collected is seldom fully restricted to researchers: *Voix des Kivus,* shared information with trusted emergency response organisations and sent selected information back to participants, but only for their own village.

(van der Windt, 2014). For the CCPE project, we would consider requests to share de-identified data with *bona fide* researchers. Data collected for the Libya Crisis Map was initially only available to UN OCHA and volunteers crowd-tasked to process reports. Later, a "heavily redacted" version, which sported a 24 hour delay and starkly abbreviated reports, was made fully public (Standby Task Force, 2011).

The distinction between crowdsourcing and crowdseeding defines the second axis on which crowd-based projects vary. Crowdsourcing projects allow anyone who wishes to do so to report information, although in practice this is restricted to those individuals who have access to the ICT the project uses and who are aware of the project (Roberts & Marchais, 2017). The Libya Crisis Map for example, in practice appears to have heavily relied on a limited number of independent journalists in Libya (Standby Task Force, 2011). Allowing anyone to contribute implies that participants cannot receive any training and can be given only minimal instruction as to what to report and how to categorize it. Information collected in this manner requires 'post-processing': Ushahidi Kenya manually checked reports to remove inflammatory messages (Okolloh, 2009), and the Libya crisis map crowd-tasked volunteers globally to categorize messages into meaningful groups (Standby Task Force, 2011). Crowdseeding projects, by contrast, purposefully select participants. *Voix des Kivus* employed survey-like sampling methods (Van der Windt & Humphreys, 2016), which in an ideal world would ensure that the crowd participating in the project is a representative sample of the overall population. For the CCPE Syria project, we purposefully selected individuals based on their location within Syria, who we felt were especially motivated, non-partisan, and well aware of the risks involved. The U'Shahid project employed a mixture of both methods of participant selection (Meier, 2012). Purposefully selecting participants makes it possible to give training *a priori*, which reduces the need for post-processing. *Voix des Kivus* trained participants to use a range of numerical codes to categorize their messages (Van der Windt & Humphreys, 2016), and during the CCPE project, we trained participants according to a codebook [authors, 2017].

Figure I: A classification of projects using crowd-based methods in conflict

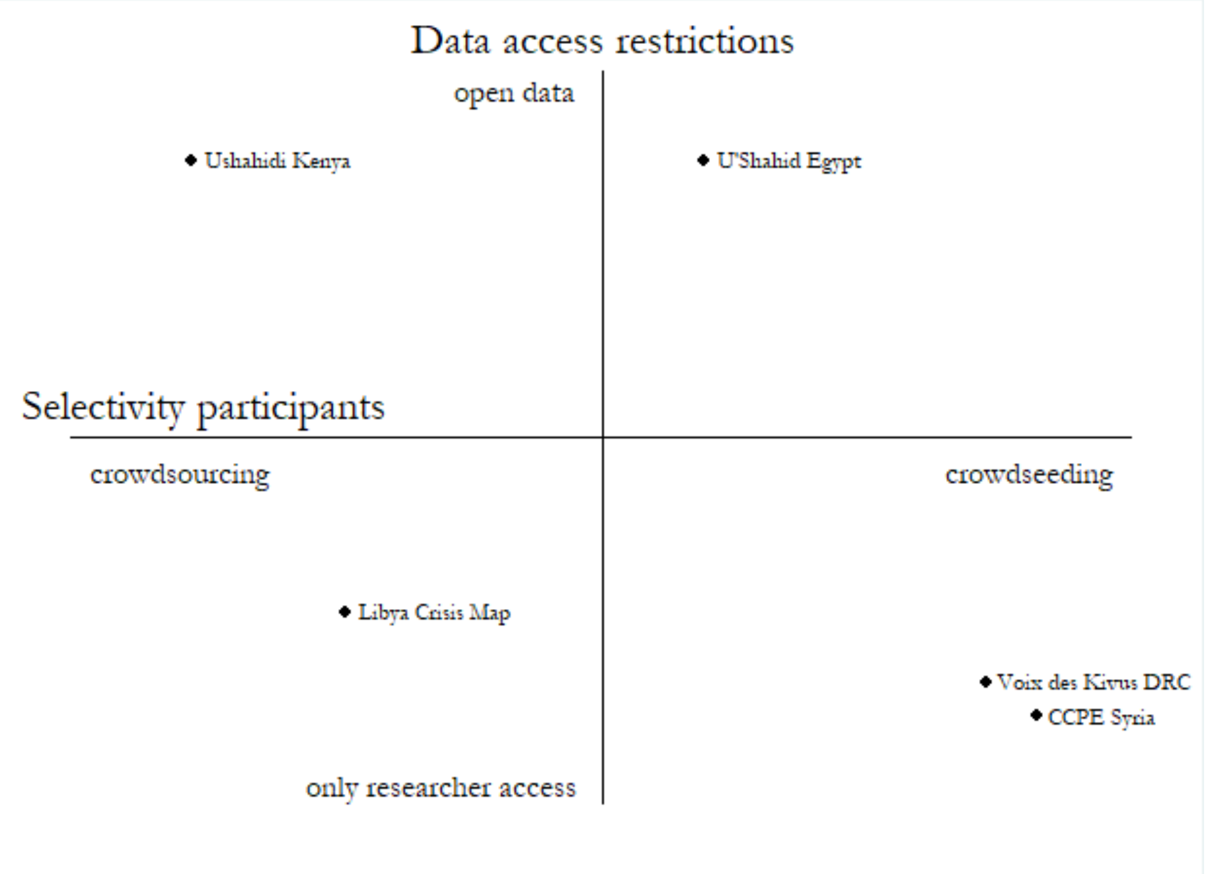

Figure I illustrates the two axes on which projects employing crowd-based methods in conflict situations differ and maps the main crowd-based projects informing this paper. Throughout this paper, we will argue that the position of a project along these two axes affects which questions of ethics will arise. These axes are meant to illustrate a continuum, not to imply that the difference between, for example, crowdseeding and crowdsourcing can be measured on some linear scale.

The typology in Figure I is specific to crowd-based projects in conflict situations: projects in other contexts may vary along other dimensions. For example, they may differ in whether participants remain anonymous to the wider public, or the sensitivity of information gathered in terms of whether it is publicly known or not. In practice, the crowd-based projects in conflict zones we reviewed do not vary along these dimensions. These projects generally consider all information, publicly known or not, as potentially sensitive and consequently keep all participants anonymous.

For example, although a bombardment of an area might be visible to all in the vicinity, revealing which individuals or communities are systematically willing to report bombardments might make those reporting vulnerable to retaliation by armed actors.

4. Dominant considerations in procedural ethics

To structure our analysis of the ethical considerations involved in using crowd-based methods in conflict-affected areas, we first review, and then go beyond, considerations that dominate the procedural dimension of ethics. Procedural ethics refers to seeking approval from an ethics committee for research involving humans (Guillemin & Gillam, 2004). From our experience of going through this process for the CCPE project, this almost exclusively involved a discussion of: (1) informed consent; and (2) minimizing risks to participants and avoiding any harm. There was also an emphasis on ethical considerations that followed from the use of technology. Existing work discussing the ethics of crowd-based methods similarly focusses on these aspects (Chamales & Baker, 2011; Martin-Shields, 2013; Pham & Vinck, 2012). The discussion in this section highlights how acutely relevant these maxims of procedural ethics are throughout all stages of executing a crowd-based project. However, we argue that researchers executing such projects should also consider further aspects of ethics, which will be discussed in Section 0.

4.1. **Informed consent**

Within our typology of crowd-based projects, the selectivity in who can participate largely influences the difficulties in obtaining informed consent from participants, whereas controls on access to the data influences *what* participants are consenting to.

Crowdsourcing, which calls on an anonymous crowd to report information, poses much greater difficulties for obtaining informed consent compared to crowdseeding. Some of these difficulties follow from the use of *technology*. If the project's ICT of choice is SMS, the participant in crowdsourcing projects is the first mover in terms of revealing information and a single SMS message is likely too short to convey a full consent statement. If the ICT used is web-based, it is possible to present potential participants with an online consent statement and have them acknowledge agreement with this statement by ticking one or more boxes. If implemented in crowdsourcing projects, these measures might still fall short of truly informed consent. Rightly or wrongly, there is a tendency to swiftly 'click through' online agreements: how often do we actually read the 'terms and conditions' we agree to before installing some software?

There exist ways to mitigate difficulties in obtaining informed consent that follow from technology. For example, consider consent obtained by having potential participants tick a single box indicating that they have read the terms and conditions that are provided in a document that needs separate downloading. This is less credibly 'fully informed' than consent obtained by having participants tick a box with a description for each potential risk that they are taking. Projects rarely provide information about the consent process in this level of detail (Meier, 2012; Okolloh, 2009), although it may have strong implications for how 'informed' the consent obtained really is.

The main reason why it is more difficult to obtain truly informed consent for crowdsourcing projects, however, follows from the *physical distance* between researchers and participants. When using crowdsourcing, the researchers and participants are *only* in contact through the ICT used. Even though the researcher is "responsible for ascertaining that the subject has comprehended the information [necessary to give informed consent]" (The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979), he or she has no alternative way to verify this. The researcher cannot read the consent statement out loud or observe first-hand clues as to whether the respondent has understood it. It is also difficult for the researcher to establish whether the potential participant is a person of diminished authority who requires additional protection: how for example, could we ensure that the participant is not a minor?

When participants to the project are purposefully selected, i.e., when a project engages in crowdseeding rather than crowdsourcing and falls on the East end of the selectivity of participation axis, obtaining informed consent poses fewer difficulties. Persons with diminished authorities can more easily be excluded from the project. Possible risks can be discussed with potential participants prior to participation and participants can more easily ask questions to the researcher. None of these measures necessary guarantee full understanding of risks on the part of the participants, but by purposefully selecting participants, researchers can limit participation to subgroups who are *likely* to be well-informed of the associated risk. This was one reason why we opted to include only individuals already reporting on the conflict in Syria in the participant pool for the CCPE project.

Rather than writing off all crowdsourcing as unethical by definition, because obtaining informed consent can be problematic, one might usefully consider whether some projects involving crowdsourcing could ethically be exempt from the requirement to obtain informed consent. Guidance on research on human subjects is helpful here. Informed consent may be waived in emergencies under certain conditions. Standout among those are that the potential participants facing life-threatening situations and that participation holds out the prospect of direct benefit (Office for Human Research Protections, 1996). How realistic the prospects for direct benefit are can be difficult to foresee *a priori* though. For example, the Libya Crisis Map was commissioned by UN-OCHA, which in its capacity of coordinator of humanitarian assistance might have been expected to lend support to participants. In practice however, UN-OCHA could not provide the volunteers working on the project with concrete examples of how they had acted upon information received through the Map (Standby Task Force, 2011; Verity, 2011).

The second axis of difference, control of access to information collected, largely determines what participants are consenting to. If data collected is freely accessible to all, this fact is easy to communicate to potential participants (although it may pose risks to them which they cannot fully foresee, see section 4.2). If access to data is controlled, it cannot always be foreseen at the moment of data collection who might request access to the data. If information is collected anonymously, it may be very difficult for participants to withdraw their consent for this use of the information *ex post*, if data collected is used in a manner they disagree with, or if a change in context (e.g. regime change) alters the risk involved in having participated. In the course the Libya Crisis Map project, participants protested when the data was later used by the International Criminal Court (Martin-Shields 2013), but they could only cease their participation as a response, not withdraw information already provided.

This highlights a broader tension between minimizing risks on the one hand, and obtaining *written* informed consent and allowing participants to withdraw consent on the other. The identity of participants may be best protected if information is gathered and stored anonymously, and their names are not languishing in some database that could potentially be infiltrated or hacked. However, if a participant withdraws consent, their particular entries have to be identifiable in order to remove them and *written* informed consent would usually require participants to sign their name. Guidance for human subjects research provides a clear set of circumstances under which *documentation* of informed consent could be waived: if the consent statement is the only documented link between participants and the project and if the principal risks to participants would be the potential harm resulting from a breach in confidentiality (NIH Office of Extramural Research, 2018). In the CCPE project, we argued that these criteria were met. However, although the university ethics board was well aware of the possibility to waive *documentation* of informed consent, representatives of university administration were less so. Hence, we encountered initial resistance against not documenting written informed consent.

In sum, crowdsourcing poses more difficulties in obtaining informed consent compared to crowdseeding. There is also a tension between *written* informed consent and minimizing risks, which is the subject of the next section.

4.2. ‘Do no harm’ and minimizing risks

As with informed consent, strategies employed to minimize risk to participants vary considerably depending on the position of the crowd-based projects along the two axes in Figure I. Crowd-based projects that publicize their existence, either to attract reports from the 'crowd' or by making their data publicly available, often rely on strategies that require large numbers of participants, whereas projects in the South-East quadrant of Figure II rely on careful management of the number of participants. The former strategy can be risky: participants might be subject to a greater risk of harm if participation in the project is not as high as envisioned. However, not all projects that have put participants at risk have been considered unethical *per se*, possibly because they hold greater potential for benefits to participants or because the researchers are themselves resident in a conflict area.

Projects relying on public calls for information (crowdsourcing) and popular attention for data collected (few controls on access) commonly employ a dual strategy to minimize risk for researchers and participants: maximize access to information collected and participation in the project so that malevolent actors would draw public censure if they threaten the project. We call this *safety in numbers*. U'Shahid Egypt deliberately employed high public visibility as a strategy to protect participants. Meier (2012) discusses why Egyptian National Security, although harassing and surveiling participants in this project, did not go beyond this to arrest them. Participants were mostly well-known bloggers with large numbers of followers on social media. Hence, arresting them would likely draw more attention to rather than detract from government election malpractice. Ushahidi Kenya likely relied on numbers to enable individual participants to remain anonymous. Large numbers of participants make it more difficult to identify an individual participant, even if there is open access to information collected. If the only reported violent event is located near of my front door, it might be easier to guess that I'm a participant to a crowd-based project than when there are many reported events near me.

Yet in the absence of large number of participants, risks are at their highest, especially if data is openly accessible. If a crowd-based project attracts the attention of malevolent actors, but is not well-known enough to cause public outcry, participants are at great risk of harm if they are easily identifiable. If there are not enough reports to obscure the identity of the few individual participants, this can cause real damage, to participants and researchers. In some crowd-based projects, the use of the 'safety in numbers' strategy has resulted considerable harm to both researchers and/or participants. The brain behind Ushahidi Kenya, Ory Okolloh was forced to flee from Kenya to South Africa after she received death threats. Participants in U'Shahid in Egypt had their cell phones tapped and social media accounts under surveillance and the head of the organisation leading the project was contacted and implicitly threatened by the Egyptian security services (Meier, 2012).

These examples illustrate two points. First, the possibility of harm to researchers and participants in crowd-based projects is very real, and the 'safety in numbers' strategy can fail to keep researchers and participants from harm. Second, these projects were not considered by all to be unethical *per se*. These harms were revealed in peer-reviewed publications (Meier, 2012), the Ushahidi Kenya project often is mentioned as a positive example (Kahl et al., 2012), and the Ushahidi platform is now used widely. These projects are considered in a positive light despite having apparently violated the principle of 'do no harm'. This may be explained because these projects have the potential to provide benefits to participants (see section 5.2) or because the distance between researchers and participants in terms of power is narrower (see section 5.1).

Projects in the South-East quadrant of Figure I, commonly use an alternative strategy to minimize risk of harm that we call *safety in obscurity*. These projects seed specific participants and control access to data collected so that the project goes unnoticed by malevolent actors. This risk minimisation strategy is likely to succeed if the number of participants remains well-managed by the project. When using this strategy, the risk of harm will likely increase with the number of participants. High network traffic to the obscured site and the incidence of password loss become more likely. These could draw attention and provide unwanted access to malevolent actors, respectively. A larger project likely requires more researchers or crowd-tasking volunteers, some of whom may have connections to malevolent actors despite screening - this was of particular concern to the Libya Crisis Map project (Standby Task Force, 2011). All these factors could expose the participants and researchers to harm, particularly with the increased suspicions associated with a secretive non-public project collecting data from a conflict zone.

During the CCPE project, we aimed to minimize risks to participants through 'safety in obscurity'. We aimed to keep the project as 'under the radar' as possible and to avoid as much as possible detectable links between individual participants and the project. This meant limiting public announcements about the existence of this project and strong controls on access to the data collected: there was no public facing website with a summary of this data. The data collection platform was designed to be unobtrusive and keep participants anonymous. Access to the platform was password protected. Participants were given online safety training, which included strategies to decrease vulnerability to hacking and instructions not to communicate identifiable information through the platform. The platform itself had an innocuous URL, used high-trafficked servers with encryption to transfer information, and was not identifiable for any casual visitors as a research project or indeed a project relating to Syria at all. We decided against using in the Syrian context media and software that had been used elsewhere: SMS messaging (too vulnerable to tracking), TOR software to conceal the user's location (use of this software was judged to attract more attention than it discouraged) and use of the Ushahidi open-source software (due to reported vulnerabilities to hackers and its recognizability as a data collection platform). None of the participants reported to have come to any harm because of their participation in the CCPE project. However, increased security measures, were only feasible given the small number of project participants.

Other research-oriented crowd-based projects used a similar strategy of safety in obscurity. In the *Voix des Kivus* project, researchers kept participating villages anonymous, even to other participating villages, although it did not prove possible to keep participants anonymous vis-à-vis their fellow villagers (Windt & Humphreys, 2012). The Libya Crisis Map project initially seems to have followed a 'safety in obscurity' strategy: the project did not make collected data publicly available initially, used a password protected platform, encouraged participants from Libya to only report via encrypted email (Standby Task Force & UN OCHA, 2011), and seems to have relied on a mix of reports from known local journalists and second-hand coding of reports on Twitter and YouTube, rather than a public call for information. The later decision to make data public at the request of UN-OCHA appears to have started a fierce debate among those executing the project on the extent to which this put participants in Libya at risk and discouraged participation by inhabitants of Libya (Standby Task Force, 2011)

It is tempting to associate crowdsourcing projects with high numbers of participants and crowdseeding with a small number of participants. Even though to date most of the crowdseeded projects have involved a smaller number of participants compared to crowdsourced ones, the number of participants is not necessarily determined by this typology. Crowdseeded projects work with a selected and trained group of participants, who in some cases are monetarily rewarded for their ongoing commitment and contributions. This implies that projects, depending on budgets and capacity, could in theory be able to include a larger number of participants. Similarly, crowdsourced projects do not by default attract large numbers of participants, although relaxing or removing the eligibility and selection criteria on who can participate is likely to drive these numbers up.

Figure II: Risk to participants, risk mitigation strategies, and number of participants shows how potential risks to participants vary with the number of participants and risk mitigation strategy used. Projects who implement the 'safety in obscurity' strategy to minimize risk and harm are more likely to succeed if the number of participants remains small. The level of risk increases systematically with higher number of participants. Projects who implement the 'safety in numbers' strategy are more likely to succeed in reducing harm if they attract a large number of participants.

It might be possible for projects to adopt different risk mitigation strategies at different points in time. Van der Windt (2014) suggests that *Voix des Kivus*, while initially adopting a ‘safety in obscurity’ strategy, could in time have been scaled up to be “known and supported by many”. In terms of Figure II, this appears to suggest a shift from a ‘safety in obscurity’ strategy to a ‘safety in numbers strategy’ at about the number of participants where the two lines intersect, resulting the “hump-shaped” relationship between scale and risk of harm that Van der Windt hypothesizes. Although possible in theory, how to calibrate such a swift with changes in the number of participants in practice remains an open question.

As depicted in Figure II, risks are likely not close to zero regardless of strategy used, or the number of participants involved. Therefore, it is helpful to go beyond procedural ethics and depict the risk minimization strategies vis-a-vis possible benefits of the crowd-based projects.

Figure II: Risk to participants, risk mitigation strategies, and number of participants

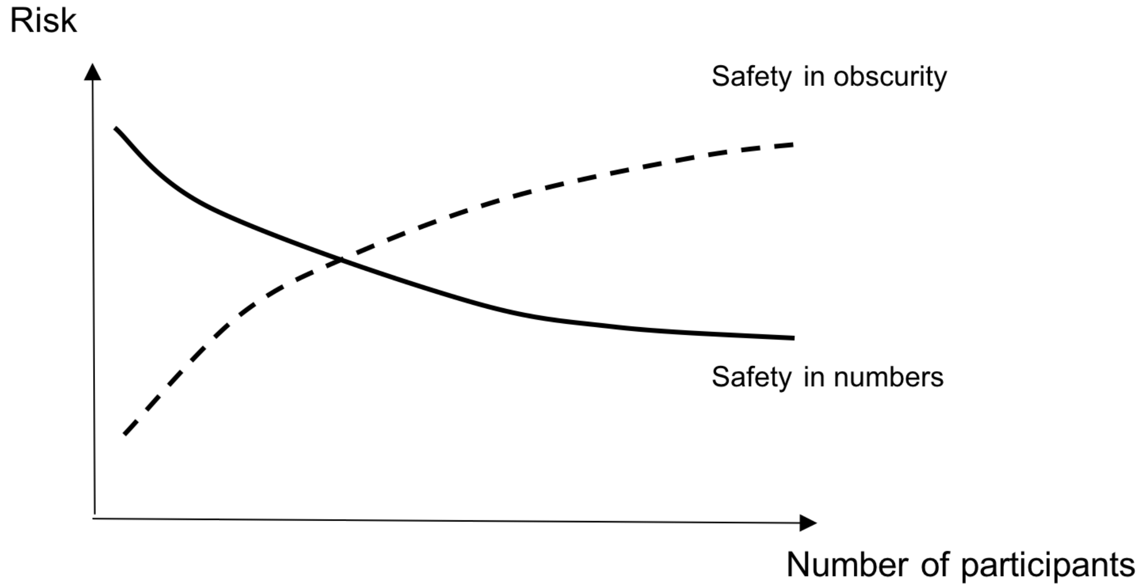

5. Beyond informed consent and harm

In existing literature on the ethics of crowd-based projects in conflict zones, and in our experience of obtaining IRB approval for the CCPE project, considerations of procedural ethics were strongly dominant. However, the principles of informed consent, 'do no harm' and minimizing risks to participants are applications of broader ethical principles. The Belmont report (The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1979) identifies three broad principles of ethics. First, respect for persons – treating individuals as autonomous agents and giving additional protection to those with diminished authority. Second, beneficence – including the principle of 'do no harm' and maximizing benefits to participants while minimizing risks. Third, justice – the fair distribution of risks and benefits from research.

Although informed consent and 'do no harm' are extremely important to the ethical execution of crowd-based projects in conflict zones, we argue that researchers should *in addition* consider other applications of the Belmont principles, specifically beneficence through maximizing possible benefits to participants, and respecting individuals' autonomy through co-creation of research [another paper, this issue]. In these applications of ethical principles, the *distance* between the researchers and participants in terms of *power* is paramount. Focus on procedural ethics may also obscure other dimensions of ethics. This section specifically considers micro-ethics (Guillemin and Gillam 2004) and ethics as articulated by professional codes of conduct.

5.1. Respecting autonomy through co-creation

In many crowd-based projects, the distinction between researchers and participants has been substantially blurred, regardless of their position on the two axes presented in Figure I: A classification of projects using crowd-based methods in conflict. This could happen if the 'researchers', the individuals setting up the project and collating the data, are themselves resident in the conflict-affected area, and thus are at least in theory part of the crowd. Ushahidi Kenya, which was set up by a Kenyan blogger, is an example of such a project (Okolloh, 2009). In the case of the CCPE Syria project, many of the 'participants' dedicated their time outside the project to investigating the same issue we 'researchers' did: the Syrian conflict. One might (and indeed participants themselves might) argue that they were researchers in their own right.

Crowd-based methods may therefore seem a very natural vehicle for the co-creation of research. Co-creation of research covers a range of practices, including giving participants a say in the research question or research process and/or ask participants for their interpretation of the research findings. It has been suggested that such co-creation is a strong and pro-active way of respecting participants' autonomy (Guillemin and Gillam 2004). Indeed, crowdsourcing has been called a "liberation technology" (Meier 2012), empowering those who were formerly passive recipients of humanitarian aid by giving them means to indicate which help they need (Hunt & Specht, 2019).

However, we want to caution against methodological optimism: co-creation is not achieved merely by virtue of using crowd-based methods, and these methods are not a "liberation technology" by definition. Meaningful co-creation arguably involves gauging what co-creation means to participants, which might involve some input on the participants' behalf in *what* data is collected and/or how it should be interpreted and used. None if this is *guaranteed* by using crowd-based methods, which only guarantee a narrowing of the power distance between researchers and participants in a very limited sense. The only thing that crowd-based methods ensure is that whoever is part of the 'crowd' has a say in when and how frequently to report.

In fact, incorporating full-blown co-creation in a crowd-based project will almost inevitably meet stumbling blocks, the nature of which will depend on the position of the project in the classification proposed. In crowdsourcing projects, it is unknown who will be a participant until the project and data collection platform are set up. Asking the participants what co-creation means to them, or giving them a say in which data gets collected is therefore very difficult. There is however some possibility for participants to organically shape what is being reported, as most crowdsourcing projects provide few instructions and restrictions to this. In addition, if crowd-based projects chose to make the data collected open access, this at least in theory gives participants the power to interpret and use the data as they see fit, although accounts of these projects do not give concrete examples of this. Crowdseeding projects, by selecting participants prior to the start of the project, can more easily elicit participants' input prior to data collection, but the risk involved in identification of a select group of participants often warrants that access to data is strictly controlled. Participants in the CCPE project gave input into the design of the data collection platform, and thus had some input in what data was being collected. But we as researchers decided against giving participants access to the data collected. The *Voix des Kivus* project made similar decisions (van der Windt, 2014). In most crowdseeding projects, the power to interpret and use the data remains in the hands of the researchers.

Although crowd-based project naturally blur the boundary between participants and researchers the method's potential to *guarantee* co-creation of research and close the power distance between researchers and participants has been overstated. Especially crowdsourcing projects may present an opportunity for meaningful co-creation, but whether this in fact emerges depends on researchers choosing a project design that surmounts barriers to co-creation.

5.2. **Beneficence through maximizing benefits**

Although 'do no harm' is the most prominent application of the principle of beneficence, a sometimes overlooked application of beneficence is to *maximize benefits*, whilst minimizing risks. Crowd-based methods have some potential to deliver benefits to participants, although we should be careful not to overstate this potential, or to raise unwarranted expectations. Researchers may be in a position of power vis-à-vis participants when deciding which benefits a crowd-based project should pursue, and any project in which participants are given some power to impose risks on a research institution is likely to be strongly opposed by this institution.

Four potential benefits of crowd-based projects have been put forward by the literature. First, *documentation*. This was the goal of, for example, the Ushahidi Kenya project: "[to create] a memorial, an archive of sorts, as a reminder […] so that Kenyans would hopefully avoid repeating history" (Okolloh, 2009). The second is *direct aid*. The Libya Crisis Map was set up by coordinator of humanitarian aid UN-OCHA, to improve its "situational awareness", which in turn could presumably improve aid provision (Verity, 2011). A third potential benefit is greater *awareness*. This may be awareness on the part of the population of the conflict-affected population itself: U'Shahid Egypt aimed to raise awareness of electoral fraud (Meier, 2012). Or it may be to draw international attention to the plight of a conflict-affected population: the Syria Tracker, an Ushahidi-like initiative in Syria, has such aims.[4] Fourth, analysing data from an individual conflict may facilitate research and thereby help create greater *knowledge* of conflict in general (the goal behind the CCPE project and *Voix des Kivus*) (Van der Windt & Humphreys, 2016; Windt & Humphreys, 2012).

It is not the purpose of this paper to evaluate the effectiveness of crowd-based methods in realizing these benefits. We might be rightfully sceptical of uncritical enthusiasm for these methods (Read, Taithe, & Mac Ginty, 2016). We will restrict ourselves to four remarks. First, realization of direct aid seems extremely unlikely in situations of violent conflict. UN-OCHA, despite having commissioned the Libya Crisis Map, considered the format in which data was presented ultimately not actionable, and numerous daily digests were blocked by its spam filters (Standby Task Force, 2011). *Voix des Kivus* aimed to partner with local organizations with a capacity to provide aid, but found these not responsive to information received (van der Windt, 2014). Even strong links to humanitarian or field partners seem no guarantee that crowd-based research can be translated into aid. Second, documentation and awareness may be more valuable to participants than researchers are aware of. Participants in a research project in El Salvador repeatedly stressed how important it was to them that their story was told (Wood, 2010). Third, it is difficult for a project to set the

---

[4] https://www.humanitariantracker.org/syria-tracker, accessed 19 March 2018.

relatively modest goals of documentation and awareness, without raising participants' expectations of direct aid. It may be difficult to imagine that a world informed about terrible events would choose to stand by, but in practice, this is often the most likely outcome in conflict situations. This imposes an obligation on the researchers beyond stressing that the project will not have direct benefits to participants, as expectations may be raised regardless of this. It is the responsibility of the researcher to ensure that participants are not motivated to take risks because of unrealistic expectations of benefits; a responsibility that may be hard to fulfil, especially if there is no direct contact between researchers and participants. Fourth, we consider the first three benefits to be activist-oriented, in the sense that they might accrue to project participants. More research primarily benefits the researchers, whose careers may be advanced by publication of such research, and the links between increased knowledge of conflict and benefits to participants are tenuous.

Researchers are likely in a position of power vis-a-vis participants when deciding which mix of the four potential benefits to pursue, and interest of participants and researchers may not naturally overlap. With some caveats, activist benefits are maximized if the largest possible crowd participates in them and if there are relatively few constraints to accessing the information gathered. This corresponds to North-West quadrant of Figure I. Documenting and raising awareness of atrocities is most effectively done if this is a public 'monument' accessible to all, and this account might be more complete if more people participate. This also gives respondents the largest audience for their story. It is therefore unsurprising that many activist-oriented projects are located in or near the North-West quadrant (i.e. Ushahidi Kenya and U'shahid Egypt). However, benefits to the researcher are often greater for projects in the South-East quadrant of Figure I, because it is easier to safeguard the quality of data for such projects (see Section 5.4). There may be only small areas of overlap where interests of participants pursuing activism and researchers coincide.

It is tempting to conclude that crowd-based projects should “cut out the middlemen” – the researchers – and should be led by activists residing in conflict areas who are best placed to decide which benefits to pursue at which risks. Ushahidi Kenya is a good example of such a project, and has likely received the praise is has because of this. The more crowd-based projects have become mainstream however, the fewer of these ideal-type projects appear to take place. Hunt and Specht (2019) review 51 crowd-based projects, most of which sport only a handful of organizations as “main players” (Hunt & Specht, 2019).

If interests of activists and researchers were to overlap, or if participants were given some power to co-determine which benefits should be pursued, our experience during the CCPE Syria project suggests that institutions employing researchers are extremely unlikely to accept a situation where participants have some power to impose risks on this institution. In our case, university administration identified two risks to the university. Actors in the Syrian conflict engaged in cyber warfare, for example through the Syrian Electronic Army (SEA). SEA is a group of hackers that has proven effective in mounting online attacks against media outlets[5]. This raised concerns that the SEA might hack the university’s systems and disrupts its website, email accounts, etc. in response to the CCPE project. Another concern was whether the university could be held liable for any harm to participants, especially if there existed a *de facto* employment relationship (not a unique consideration (Perera, 2017)). Even if we had wanted to do so, it is extremely unlikely that the institution who employed us would have allowed us to pursue any kind of activist benefit.

This puts researchers and their home institutions in an uncomfortable position. It makes pursuing co-creation of research and pursuing direct activist benefits to participants difficult. Going one step further, applying the ethical principle of justice, we might argue that researchers and research

---

[5] https://www.theguardian.com/technology/2013/apr/29/hacking-guardian-syria-background, accessed 12 June 2016.

institutions should cooperate with activists if they are in a better position to bear the risks involved. Outside research institutions might lend credibility to open access crowdsourced data, may be less vulnerable to malevolent actors, while being better able to host a secure data collection platform. Some even go as far as to argue that social science researchers have a *duty* of activism, in the sense of critically and normatively reflecting on our social world (Sleat, 2017). If we accept this argument, the incentive for institutions employing researchers to avoid risks to itself impedes executing this duty.

5.3. **Micro-ethics**

Micro-ethics refer to difficult, unpredictable situations that may arise between a researcher and a participant that have a clear ethical dimension – or "ethically important moments" (Guillemin & Gillam, 2004). Guillemin and Gillam (2004) give the example of a participant in a study on heart disease disclosing to a researcher that she recently found out her husband has sexually abused her daughter. The authors argue that the 'big' principles of ethics guiding human subject research, although relevant, do not furnish the researcher with an immediate response to this disclosure.

Crowd-based methods amplify micro-ethics, as they create a direct line of communication between participants and researchers, which the participant may use at any time. Any information that the participant reveals may trigger an "ethically important moment". But there is not necessarily someone paying attention at the other end of the line. Researchers don't commonly monitor incoming reports 24 hours a day, although the Libya Crisis Map may be an exception. Even if they do, they may not be in a position to offer assistance to participants, or indeed respond to them at all. Van der Windt and Humphreys (2012) describe how the *Voix de Kivus* project was originally more ambitious in size, but never reached the planned scale and was then discontinued, partially because of a deep uneasiness on the part of the researchers. They were receiving messages about atrocities happening in their research area, but could not adequately respond themselves, or get humanitarian organisations in the area to provide a response.

Guillemin and Gillam (2004) suggest that there is no single rule that helps researcher "deal with" micro-ethics, but that awareness of the potential for and sensitivity to ethically important moments, and developing means to respond to these moments, are important. In the context of crowd-based projects, Van der Windt and Humphreys (2012), through their candid account of how they saw ethically important moments arise that they had not anticipated prior to their project, have given other researchers the opportunity to become aware of the potential for micro-ethics earlier. Strong partnerships between researchers and field partners, who may not be able to provide direct aid, but at least have the capacity to respond to participants, may go some way towards improving responsiveness when ethically important moments occur (Windt & Humphreys, 2012).

5.4. **Professional ethics: getting it right**

A final ethical dimension of using crowd-based methods is professional ethics of the researcher. Professional ethics would dictate that researchers ensure that data collected is truthful and representative, but also that the overall universe of research on conflict covers all conflicts and does not duplicate efforts, placing an undue burden on conflict-affect populations. [Another paper, this issue] call this the imperative to "get it right". Crowdseeding promises more veracity and representativeness of data, but this may come at the expense of covering only conflicts where crowdseeding is feasible and following a 'safety in obscurity' strategy may risk duplication of efforts.

It goes beyond the scope of this paper to provide a comprehensive assessment of quality of data gathered through crowd-based methods in comparison to other methods, such as media-based event data or surveys. The most reliable evidence on this is provided by Van Der Windt and Humphreys (2016), who suggest that a dataset based on coding secondary sources (ACLED) vastly underreports conflict events compared to crowdseeding (27 events in ACLED versus 1439 according to *Voix des Kivus*). The benefits vis-à-vis survey-based measures of conflict are less clear in the case of *Voix des Kivus*. This paper cannot provide such a comparison for CCPE data, as we do not have access to other data covering the same time period.

There is an argument to be made that crowdseeding, especially if it relies on survey-style sampling methods, provides higher quality data in terms of veracity and representativeness. *Voix des Kivus* used survey-like sampling methods to include a representative sample of villages in the project (Van der Windt and Humphreys 2016). This project also provided 'seeded' respondents with a mobile phone. In crowdsourcing projects, the sample of participants would be limited to those aware of the project, willing to send in reports, and with access to whichever ICT the project is using. These projects might therefore be subject to biases. An additional argument in favour of using crowdseeding rather than crowdsourcing for the CCPE project was veracity of the reports. We were very aware that reports from an anonymous 'crowd' might in fact constitute attempts by conflict actors to skew our data. The Syrian regime certainly attempted this when it came to mainstream media, which also raised concerns about the accuracy of data based on coding events from secondary sources (Baliki, 2017). Some crowdsourcing projects, notably Ushahidi Kenya and U'Shahid Egypt, use investigators on the ground to verify the veracity of reports by the 'crowd'. This may help address issues of veracity of reports received, but it does not address the bias from participants self-selecting into or out of participation.

However, "getting it right" does not only depend on the representativeness and veracity of information collected, but also on coverage of different conflicts and avoiding duplication. Crowd-based methods are primarily implemented to overcome the restriction to access hard-to-reach areas in conflict-affected settings. Given the difficulty of using alternative methods, the use of crowdsourcing can reduce the preparation time needed at inception and design phases, which allows for a speedy response. Yet a speedy response can compromise data quality, as participants are not a representative sample of the population, and opportunities for training are limited. Although making data open access comes with substantial risk, maintaining strict controls on access to data risks duplication: two research projects may both subject participants to risks, whereas this would not have been necessary if they had shared data. The ethical conundrum of "getting it right", hence, boils down to two questions: (1) if any data (regardless of quality and accuracy) is better than no data; and (2) how much risk should be taken to ensure the quality of the data and its use by other researchers.

6. Conclusion

This paper reflects on the ethics of using crowd-based methods in conflict-affected situations. It argues that not all these projects can be treated alike when it comes to ethics: which ethical issues arise is largely determined by the degree to which a project selects participants and controls access to information collected. It furthermore argues that ethical considerations dominant in procedural ethics are extremely important to the ethical execution of crowd-based projects, but that researchers should also consider alternative applications of guiding ethical principles, as well as non-procedural dimensions of ethics.

Implementing a crowd-based project should involve careful consideration of how to obtain truly informed consent from participants, and how to minimize the risk of any harm coming to them. But this is the starting point, rather than the end point of ethical considerations. Going beyond the maxims of procedural ethics likely involves difficult trade-offs: between benefits to participants and benefits to researchers and the institutions that employ them, between maximizing benefits by avoiding duplicate efforts and risks to participants, between different crowd-based methods and their potential for meaningful co-creation, between data quality and timeliness, between no data or a data-collection process in which the researcher can only respond with difficulty to "ethically significant moments".

This paper identifies a need to reflect further on the differences (if any exist) between the ethics of research and the ethics of activism, especially to the extent that they engage in the same activities.

Further research could also usefully explore the quality of crowdsourced or crowdseeded information vis-à-vis other methods of collecting data in conflict situations.

Many of the ethical considerations discussed in this paper are not unique to crowd-based methods. However, crowd-based methods embody a unique combination of ethical considerations usually found in a set of other research methods: micro-ethics - a term coined by researchers employing in-depth interviews or longitudinal studies -, representativeness of the data commonly required from survey-based methods, secure collection and storage of data associated with any project using some technology, and delivering benefits to participants, which we would associate more with activism than with research. In this sense, crowd-based methods are indeed a microcosm of the ethics of research in conflict situations.

7. References

Baliki, G. (2017). Measuring Media Bias in Conflict: Evidence from Syria. In *Empirical Advances in the Measurement and Analysis of Violent Conflict*. Berlin: Humboldt-Universität zu Berlin. https://doi.org/10.18452/18363

Butler, D. (2013). Crowdsourcing goes mainstream in typhoon response. *Nature News*. https://doi.org/10.1038/nature.2013.14186

Chamales, G., & Baker, R. (2011). Securing crisis maps in conflict zones. In *Proceedings - 2011 IEEE Global Humanitarian Technology Conference, GHTC 2011* (pp. 426–430). https://doi.org/10.1109/GHTC.2011.47

Fujii, L. A. (2012). Research ethics 101: Dilemmas and responsibilities. *PS - Political Science and Politics*, *45*(4), 717–723. https://doi.org/10.1017/S1049096512000819

Guillemin, M., & Gillam, L. (2004). Ethics, reflexivity, and “Ethically important moments” in research. *Qualitative Inquiry*, *10*(2), 261–280. https://doi.org/10.1177/1077800403262360

Hunt, A., & Specht, D. (2019). Crowdsourced mapping in crisis zones: collaboration, organisation and impact. *Journal of International Humanitarian Action*, *4*(1), 1. https://doi.org/10.1186/s41018-018-0048-1

Kahl, A., Mcconnell, C., & Tsuma, W. (2012). Crowdsourcing as a Tool in Conflict Prevention. *Conflict Trends*, 27–34. Retrieved from https://www.peaceportal.org/documents/130225323/0/CrowdsourcingasaToolinConflictPrevention_publ.pdf

Martin-Shields, C. (2013). The technologist's Dilemma: Ethical Challenges of Using Crowdsourcing Technology in Conflict and Disaster-Affected regions. *Georgetown Journal of International Affairs*, (Summer), 157–163. https://doi.org/10.3366/ajicl.2011.0005

Meier, P. (2012). Ushahidi as a Liberation Technology. In L. Diamond & M. F. Plattner (Eds.), *Liberation technology: Social media and the struggle for democracy* (pp. 95–109). Baltimore: JHU Press.

Munro, R. (2013). *Crowdsourcing and the crisis-affected community: Lessons learned and looking forward from Mission 4636. Information Retrieval* (Vol. 16). https://doi.org/10.1007/s10791-012-9203-2

NIH Office of Extramural Research. (2018). Protecting Human Research Participants - Web-based training course.

Office for Human Research Protections. (1996). Informed Consent Requirements in Emergency Research. *OPRR Letter*. Retrieved from http://www.hhs.gov/ohrp/regulations-and-policy/guidance/emergency-research-informed-consent-requirements/

Okolloh, O. (2009). Ushahidi, or "testimony": Web2.0 tools for crowdsourcing crisis information. In H. Ashley, J. Corbett, D. Jones, B. Garside, & G. Rambaldi (Eds.), *Change at hand: Web 2.0 for development* (Vol. 59, pp. 8–20). Retrieved from http://www.ingentaconnect.com/content/iiedpla/pla/2009/00000059/00000001/art00002

Perera, S. (2017). To Boldly Know: Knowledge, Peacekeeping and Remote Data Gathering in Conflict-Affected States. *International Peacekeeping*, *24*(5), 803–822. https://doi.org/10.1080/13533312.2017.1383566

Pham, P. N., & Vinck, P. (2012). Conflict Early Warning Systems, Public Health, and Human Rights. *Health and Human Rights*, *14*(2), 106–117.

Poblet, M. (2013). Visualizing the Law: Crisis mapping as an open tool for legal practice. *Journal of Open Access to Law*, *1*(1), 8–23. https://doi.org/10.3868/s050-004-015-0003-8

Read, R., Taithe, B., & Mac Ginty, R. (2016). Data hubris? Humanitarian information systems and the mirage of technology. *Third World Quarterly*, *37*(8), 1314–1331. https://doi.org/10.1080/01436597.2015.1136208

Roberts, T., & Marchais, G. (2017). *Assessing the Role of Social Media and Digital Technology in Violence Reporting* (IDS Working Paper Series No. 492).

Sleat, M. (2017). Responsible to whom. Obligations to participants and society in social science research. In R. Iphofen (Ed.), *Finding Common Ground: Consensus in Research Ethics Across the Social Sciences* (pp. 37–46). Bingley: Emerald Group Publishing.

Standby Task Force. (2011). Libya Crisis Map Deployment 2011 Report. Retrieved from http://blog.standbytaskforce.com/2011/09/01/libya-crisis-map-report/

Standby Task Force, & UN OCHA. (2011). *Security concerns for Libya Crisis Map*. Retrieved from https://wiki.openstreetmap.org/w/images/7/76/Security_Concerns_for_Libya_Crisis_Map.pdf

The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. (1979). *The Belmont report ethical principles and guidelines for the protection of human subjects*. Retrieved from https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html

Turkmani, R., Kaldor, M., Elhamwi, W., Ayo, J., & Hariri, N. (2014). *Hungry for peace. Positives and pitfalls of local truces and ceasefires in Syria*.

van der Windt, P. (2014). From crowdsourcing to crowdseeding. The cutting edge of empowerment? In *Bits and atoms. Information and communication technology in areas of limited statehood*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof

Van der Windt, P., & Humphreys, M. (2016). Crowdseeding in Eastern Congo: Using Cell Phones to Collect Conflict Events Data in Real Time. *Journal of Conflict Resolution*, *60*(4), 748–781. https://doi.org/10.1177/0022002714553104

Verity, Y. A. (2011). OCHA Lessons Learned - Collaboration with VTCs in Libya and Japan. Retrieved from http://digitalhumanitarians.com/sites/default/files/resource-field_media/OCHALessonsLearnedCollaborationwithVTCsinLibyaandJapanFinalNov2011.pdf

Windt, P. Van Der, & Humphreys, M. (2012). Voix des Kivus : Reflections on a Crowdseeding Approach to Conflict Event Data Gathering, (February), 1–8.

Wood, R. M. (2010). Rebel capability and strategic violence against civilians. *Journal of Peace Research*, *47*(5), 601–614. https://doi.org/10.1177/0022343310376473