Journal of Conflict Resolution
http://jcr.sagepub.com/

Crowdseeding in Eastern Congo: Using Cell Phones to Collect Conflict Events Data in Real Time

Peter Van der Windt and Macartan Humphreys

Journal of Conflict Resolution published online 4 November 2014
DOI: 10.1177/0022002714553104

Original Article

Journal of Conflict Resolution
1-34

DOI: 10.1177/0022002714553104
jcr.sagepub.com

Crowdseeding in Eastern Congo: Using Cell Phones to Collect Conflict Events Data in Real Time

Peter Van der Windt[1] and Macartan Humphreys[2]

Abstract

Poor-quality data about conflict events can hinder humanitarian responses and bias academic research. There is increasing recognition of the role that new information technologies can play in producing more reliable data faster. We piloted a novel data-gathering system in the Democratic Republic of Congo in which villagers in a set of randomly selected communities report on events in real time via short message service. We first describe the data and assess its reliability. We then examine the usefulness of such "crowdseeded" data in two ways. First, we implement a downstream experiment on aid and conflict and find evidence that aid can lead to fewer conflict events. Second, we examine conflict diffusion in Eastern Congo and find evidence that key dynamics operate at very micro levels. Both applications highlight the benefit of collecting conflict data via cell phones in real time.

Keywords

conflict, foreign aid, humanitarian intervention, conflict management

[1]New York University Abu Dhabi, United Arab Emirates
[2]Columbia University, New York, NY, USA

Corresponding Author:
Peter Van der Windt, New York University Abu Dhabi, Saadiyat Island, Building A5, Office 147, Abu Dhabi, United Arab Emirates
Email: petervanderwindt@nyu.edu

Information technology is becoming increasingly visible as a means for gathering data on social trends. Web-based social networks are used to organize, mobilize, and coordinate activities—from riots in London to regime change in the Middle East. Satellite imagery is used by the United Nations High Commissioner for Refugees to track population movements during conflicts or droughts and by the International Criminal Court to learn about the presence of mass graves in Bosnia. And cell phones, combined with geo-mapping software, are used to create early warning systems. Information technology is now also employed prominently in research. Satellite imagery has been used to study economic growth (Henderson, Storeygard, and Weil 2011, 2012), Twitter to learn about the Gaza Conflict (Zeitzoff 2011), and geographic information system data to assess many elements of conflict (Lujala 2010; Raleigh and Hegre 2009; Lujala, Gleditsch, and Gilmore 2005; Weidmann, Rod, and Cederman 2010; Raleigh et al. 2010).

Despite broad interest in understanding patterns of political violence, the development of micro-level measures remains weak. There have been improvements in data collection on violence, but these data are typically at a high level of aggregation. As Verwimp, Justino, and Bruck (2009) argue, "At a fundamental level, conflict originates from individuals' behavior and their repeated interactions with their surroundings, in other words, from its micro-foundations" (p. 307). Recent work by Autesserre in the Democratic Republic of Congo (DRC) illustrates the point: although the war officially came to an end in 2003, thousands of civilians still die each week. A plausible reason is that international interventions have focused at the regional and national levels; ignoring the local level where conflicts over land and political power became self-sustaining and autonomous from the national and regional tracks (Autesserre 2009, 2010).

Micro-level conflict data are normally collected by conducting surveys or interviews. However, obtaining such data in real time using traditional approaches is difficult because conflict events cluster in areas that have three characteristics. First, they are insecure, which limits the ability of researchers to visit them. Second, they are often physically hard to access (indeed, difficult terrain is significantly related to a higher incidence of civil war; Fearon and Laitin 2003). Third, they may exhibit high levels of societal suspicion and distrust due to, for example, recent conflict experiences or the influx of new, unknown people due to displacement. These characteristics can create different types of biases that affect data quality.

First are selection biases. If areas are off-limits while conflict events take place, researchers often access and enlist the cooperation of only a (nonrandom) fraction of the research populations.[1] Such selection biases may arise in subtle ways—for example, even if all areas are covered by a survey, security considerations may result in scheduling adjustments, so that different regions are only visited when it is safe to do so. Also, testimonies are only taken from the living.

Second are risks of recall bias arising when data are gathered after the conclusion of conflict. Coughlin (1990) provides an overview of the literature. While time cues might lessen this bias in some cases, De Nicola and Giné (2014) find evidence that the use of time cues can sometimes exacerbate the problem. One example is the problem of "telescoping," where respondents run together several events and lose track of when they actually happened. The importance of prompt data collection then may depend on the duration of the conflict and the number of events taking place.

Third are risks of reporting bias that may arise when sensitive questions are asked in settings characterized by distrust and suspicion. Trust has to be created to obtain high-quality responses. However, this is especially difficult in conflict areas where data collection is often a one-off activity.[2] And, of course, reporting may also be biased because of deliberate attempts to manipulate information.[3]

We seek to investigate how information technology can overcome the problems associated with the collection of conflict data. To do so, we piloted a "crowdseeding" system that collected information via short message service (SMS) from preidentified informants in randomly sampled locations. The project, *Voix des Kivus*, was implemented between 2009 and 2011 in the war-torn province of South Kivu in Eastern DRC. Large parts of the province are hard to access and although there is reasonable cell phone coverage, access to cell phone technology is limited.

In using cell phone technology, we build off the "crowdsourcing" approach pioneered by groups such as Ushahidi. Under crowdsourcing, anyone can send an SMS to a central platform in which messages are gathered, stored, and visualized on a map. Though clearly a compelling system for "fire alarm" monitoring, there are reasons to worry that the data generated through such systems are not representative.[4] Reporters might send incorrect information—for example, they might overreport hardship hoping for humanitarian intervention. The data may also be unrepresentative for more innocent reasons: only people with access to a cell phone (and who have heard about the project) can send messages. Furthermore, the system will only receive messages from people who are willing to pay the cost of an SMS.[5]

The crowdseeding approach seeks to combine the strengths of crowdsourcing technologies to generate detailed real-time data, with the strengths of traditional approaches that rely on known sources and representative samples. Like crowdsourcing, crowdseeding can alleviate some of the concerns regarding selection biases highlighted previously. Rather than relying on discrete retrospective data gathering, these approaches rely on continuous information flows from individuals in the conflict region. Indeed, even if an area is off-limits to a research team, SMS can still be sent and received. Moreover, real-time data gathering removes concerns related to recall. Crowdseeding offers additional advantages relative to crowdsourcing. First, by working with a random set of villages, data from a crowdseeding system render data that are representative at the village level. Moreover, if the reporters are preidentified and given the means to provide information, the risk that bias will be introduced when reporters self-select into the system is reduced. Even in the event of forced displacements, the reporter can continue providing information. Finally, crowdseeding may alleviate concerns associated with distrust because, by working with preselected reporters over extended periods, the system allows both for the generation of trust and for a greater ability to verify received information. Repeated interactions can reduce the incentive for reporters to send incorrect information and allows for ready auditing and verification of data.

In this article, we introduce data from the *Voix des Kivus* system and explore its utility in practice. In the second section, we describe the implementation and feasibility of the system. We then present the data in the third section, describing the distribution of messages and assessing the overall data quality. Specifically, to assess data quality, we investigate whether time, reporter type, and event type affect whether an event is reported and compare the coverage to two existing conflict data sets collected by traditional survey methods and by Armed Conflict Location and Events Data Set (ACLED's) primarily media-based system. The fourth section provides two illustrations of the use of the data. In the "Application 1: The Conflict Impact of Development Programs" subsection, we assess the effects of a downstream experiment on aid and conflict, exploiting additional data on the presence of a randomized development intervention in the region. Despite the small sample, we find evidence for a negative relationship between aid and conflict in this set of villages. In the "Application 2: Spatial Clustering and Diffusion of Conflict" subsection, we use the data to assess the spatiotemporal distribution of violence events. Though identification is weaker for this analysis, the evidence is consistent with large local spillover effects that decay at distances of around fifty kilometers. Both applications suggest that the assessments are highly sensitive to the timing of measurement in a way that often cannot be assessed from traditional conflict data sources. We conclude in the fifth section by considering ethical and practical implications of scaling up this type of data system.

The *Voix des Kivus* System

The *Voix des Kivus* (literally "Voice of the Kivus") data collection system was piloted in the province of South Kivu in the war-torn East of the DRC over an eighteen-month period starting in August 2009. The pilot ran in eighteen villages spread over four territories: Kalehe (5), Mwenga (4), Uvira (3), and Walungu (6).[6] Villages 1 through 4 were selected purposely to ensure some conflict and some nonconflict areas; villages 5 through 18 were selected using stratified random sampling from these territories. Specifically, villages were randomly selected stratified by chiefdom and by the village's treatment status in a large development project (we discuss the project in more detail in the "Application 1: The Conflict Impact of Development Programs" subsection), taking village population sizes into account.[7] Figure 1 shows the approximate area of operation.[8] The headquarters of *Voix des Kivus* was Bukavu—the capital city of South Kivu and also the city where most of the province's nongovernmental organizations (NGOs) are located.

Within these villages, the *Voix des Kivus* system worked as follows. In each village, we identified three reporters: one representing the traditional leadership (the

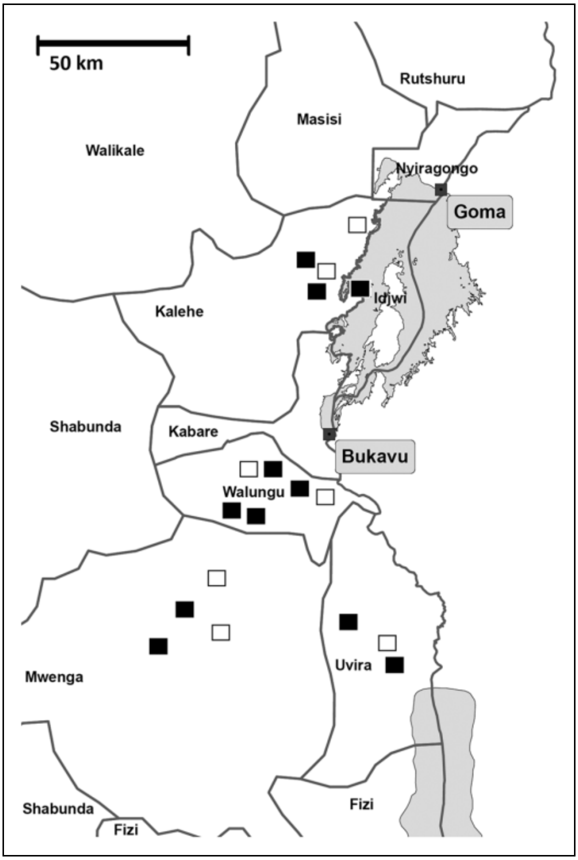

Figure 1. Map illustrating general area of *Voix des Kivus* operation.

Note: Squares denote approximate location of *Voix des Kivus* villages, accurate only up to the level of the chiefdom. Solid (hollow) squares denote villages that do (do not) participate in the development project discussed in the "Application 1: The Conflict Impact of Development Programs" section.

chief of the village or his appointee), one representing women's groups (the head of the women's association or her appointee), and one elected by the community.[9]

Reporters were given a phone and trained on how to send messages to the system. They were provided with a codesheet that lists possible events that can take place in the village, organized in ten categories: (1) presence of military forces, (2) attacks on the village, (3) deaths related to armed combat, (4) local violence and property loss, (5) displacement, (6) health events, (7) natural disasters, (8) development and NGO activities, (9) social events, and (10) special codes. An overview of these codes is given in Table 1.[10] Users could send a simple message containing one of these codes or they could send a standard "full text" message giving more particulars.

Table 1. *Voix des Kivus* Codebook and Summary Information.

Code	Event	Mean	SD	Minimum	Maximum	Total
10	Presence of military forces	0.04	0.21	0	1	9
11	Presence of MONUSCO	1.37	3.07	0	17	277
12	Presence of FARDC	1.79	4.49	0	23	362
13	Presence of other armed groups	0.18	0.55	0	3	37
20	Attacks (use of violence by an external group)	0.02	0.14	0	1	4
21	Attacks on village by MONUSCO	0.04	0.30	0	3	8
22	Attacks on village by FARDC	0.27	0.80	0	7	54
23	Attacks on village by rebel group	0.21	0.66	0	4	43
24	Attacks on the village by unknown group	0.09	0.39	0	3	18
30	Deaths related to armed combat	0.08	0.59	0	7	16
31	Civilian deaths (man)	0.45	1.26	0	11	90
32	Civilian deaths (woman)	0.46	1.71	0	18	92
33	Civilian deaths (child)	0.25	0.94	0	9	51
34	MONUSCO deaths	0.01	0.14	0	2	2
35	FARDC deaths	0.18	0.74	0	6	36
36	Other rebel group deaths	0.02	0.23	0	3	5
40	Local violence and property loss	0.03	0.20	0	2	6
41	Rioting	0.37	1.25	0	10	74
42	Looting/property damage	1.15	2.47	0	19	233
43	Violence between villagers	0.20	0.84	0	9	41
44	Violence between villagers due to a land conflict	0.62	2.54	0	25	126
45	Domestic violence	0.11	0.52	0	5	23
46	Ethnic violence	0.02	0.14	0	1	4
47	Forced labor by FARDC	0.30	0.95	0	6	60
48	Forced labor by other army groups	0.08	0.54	0	7	16
50	Displacement	0.00	0.07	0	1	1
51	Kidnapping by men in uniform	0.45	1.15	0	8	91
52	Kidnapping by rebel group	0.22	1.07	0	13	44
53	Arrival of Refugees or IDPs	0.16	0.59	0	4	33
54	Departure of villagers as IDPs	0.18	0.71	0	7	37
55	Disappearances	0.03	0.17	0	1	6
56	Villagers were forced to move	0.05	0.33	0	3	10
57	Villagers decided themselves to move	0.02	0.17	0	2	4
60	Health	0.09	0.36	0	3	18
61	New outbreak of disease	1.37	2.44	0	15	277
62	Civilian death due to disease	2.33	3.33	0	16	470
63	Civilian death due to natural causes	0.37	1.68	0	22	74
64	Sexual violence against women	0.23	0.69	0	5	47
65	Sexual violence against men	0.26	1.68	0	17	52
70	Natural disasters	0.16	0.68	0	7	33
71	Flooding/heavy rain	0.63	1.26	0	7	127
72	Large forest or village fire	0.53	1.63	0	13	108

(continued)

Table 1. (continued)

Code	Event	Mean	SD	Minimum	Maximum	Total
73	Earthquake	0.20	0.65	0	5	40
74	Drought	0.18	0.54	0	3	36
75	Crop failure/plague	1.86	3.24	0	23	376
80	Development activities/NGOs	0.69	1.37	0	7	140
81	Complaint against NGO	0.21	0.62	0	4	42
85	Construction, reparation, or rehabilitation of a school or health center	0.11	0.41	0	4	22
86	Construction, reparation or rehabilitation of a church or mosque	0.10	0.39	0	3	21
87	Other construction, reparation or rehabilitation	0.52	1.76	0	15	105
88	Organization of security patrols	0.09	0.51	0	4	19
89	Work to improve agricultural productivity	0.32	0.83	0	4	65
90	Social	0.06	0.27	0	2	13
91	Funeral	0.61	1.66	0	11	123
92	Wedding/Other celebrations	0.49	1.17	0	7	98
93	Visit or meeting organized by national or provincial authorities	0.13	0.53	0	5	26
94	Visit or meeting organized by territory authorities	0.04	0.24	0	2	8
95	Visit or meeting organized by the chiefdom or locality authorities	0.18	0.59	0	5	37
96	Visit or meeting organized by the representative of a political party	0.27	1.05	0	9	55
97	Visit or meeting organized by the King	0.18	0.92	0	9	36
0	Nothing to report	0.97	2.15	0	13	196
82	Practice message	0.42	1.74	0	16	85
98	Unclassifiable issue (followed by text)	1.65	2.56	0	14	334
99	Security alert	0.42	1.04	0	5	85

Note: IDP = internally displaced persons. MONUSCO = United Nations Organization Stabilization Mission in the DR Congo; FARDC = Armed Forces of the Democratic Republic of Congo; NGOs = nongovernmental organizations. Mean, standard deviation, minimum and maximum are by village-month: that is, a mean of 1.37 for code 11 indicates the number of times MONUSCO presence was reported during an average village-month. Based on a total of 202 observations of village-month events. Events in the column "Total" add up to 5,081. When a text message could not be hand-coded to an event (e.g., "33"), it would be assigned to a category (e.g., "30"). Codes 84 through 89 and 91 through 97 were introduced only during the expansion in August 2010. Before the expansion codes 98 and 99 were 83 and 84, respectively.

Reporters automatically received weekly phone credit that they could use freely and were reimbursed for the number of messages sent. Remaining in the system required sending at least one message a week, but this message could be blank. Thus, sending messages to the system was free, but it was also voluntary—while users did not have to pay for each message, they did not get any marginal financial rewards for sending content either. On the receiving side was a standard cell phone linked to a laptop computer. With freely available software (FrontlineSMS and R), messages received were automatically filtered, coded for content, cleaned to remove duplicates, and merged into a database. Graphs and tables were generated automatically and then mounted into bulletins with different levels of source identifiability. Noncoded text messages (often from Swahili or one of the local languages into French and English) were translated manually.[11]

Our first goal in implementing the pilot was to assess the feasibility of data collection of this form. Establishing feasibility was important in light of a number of obvious challenges.

A first concern was that human capacity would be too weak for users to implement the project. Indeed, nineteen of our fifty-four reporters had only primary-level education, and only two had received education beyond secondary school. This concern is likely to be present in many hard-to-access areas where conflict takes place. We mitigated this concern through a (minimum) two-day training by the field coordinator and the use of relatively simple codebooks with precoded events. In fact, over the eighteen months of operation, *Voix des Kivus* received—in addition to SMS with codes—1,144 text messages in Swahili or one of the local languages; this suggests that the codebook, while useful, was not necessary for many reporters.

A second concern was technical capacity. On the receiving side, a cheap netbook, a phone, and free software were sufficient. However, things were different on the sending side. Phone coverage was not a problem—of the eighteen villages, no village had to be replaced after selection due to lack of coverage. Electricity, however, was a problem, with reporters sometimes walking up to three hours to charge their phone. The problem was solved by purchasing US$25 solar chargers and donating one to each *Voix des Kivus* village. Providing consistent access to electricity through these low-cost solar chargers was important in ensuring regular information from the villages in our sample.

A third concern, however, is specific to crowdseeding systems: Would participation in the system produce security risks for reporters? Over the course of eighteen months, *Voix des Kivus* received thousands of messages, many of which were sensitive, including information on various types of abuses perpetrated by different actors. Because of the close connection between individual reporters and the project, we were concerned that there could be reprisals against reporters for sharing information. Beyond regular monitoring, we used three strategies to address this concern. First, we operated in just four villages at first, and subsequent villages were added only after one year of careful monitoring. Second, the weekly bulletins were generated in two versions. One version contained sensitive information (with village identifiers); the other did not. The sensitive bulletins were only disseminated to a very restricted set of organizations. Public versions contained no village identifiers. Third, the system allowed reporters to include a code (1–4) to a message to indicate the event's level of sensitivity and with whom the information was to be shared ("4" only with Voix des Kivus, "3" also with close partners, "2" also with United Nations Organization Stabilization Mission in the DR Congo [MONUSCO], and "1" everybody).[12] We found that at no point did any user indicate any security concern of any form arising from their participation in the project. However, although no concerns arose in our case, there are still general grounds for concerns for the security of the reporters in a crowdseeding system because of the direct link between reporters and the project. We return to this concern in the concluding section.

The Data

In this section, we describe the data and assess data quality.[13] Specifically, we explore reporting over time (is there evidence of fatigue?), reporting across monitors (do different reporters report differently? Are results robust to weights based on the number of reporters reporting the event?), and reporting by type of event (are conflict events more likely to be underreported when the potential for sanctioning is higher?). Moreover, we compare the *Voix des Kivus* data to two existing conflict data sets collected by traditional survey methods and by ACLED's primarily media-based system.

Voix des Kivus was launched in August 2009 and operated during the first twelve months in only four villages. Then from August 2010 onward, the pilot project was expanded to eighteen villages. The expansion over time is illustrated in the top left panel of Figure 2.[14] The regular system continued through to January 2011 (the vertical line in Figure 2) but remained open for users to send messages—though without incentives for reporters—through to July 2011.[15]

The top right panel of Figure 2 shows the distribution of total event reports (dashed line) and conflict event reports (solid line) over time, per village. Overall, the system generated a relatively constant stream of messages over this period. As shown in the lower left panel of Figure 2, reporters exhibited relatively little fatigue—while different reporters reported at different rates, they nevertheless reported at relatively constant rates over time.[16]

In this short period, the project received a total of 4,783 non-empty SMS messages. The reporters sent messages about a total of 5,081 events of which 4,623 were unique—many village events were thus reported by more than one reporter. Of these non-empty SMS messages, 1,144 were text messages.[17] Table 1 provides summary information per event (where the unit of analysis is the village-month). Identical messages received from the same reporter within thirty minutes of each other, and identical messages received from different representatives in the same village within twenty-four hours of each other, are treated as single events. The table also reports the total number of messages sent per event. We find a large number of messages reporting the presence of troops (685). The most common conflict events include looting (233), violence between villagers (167), and

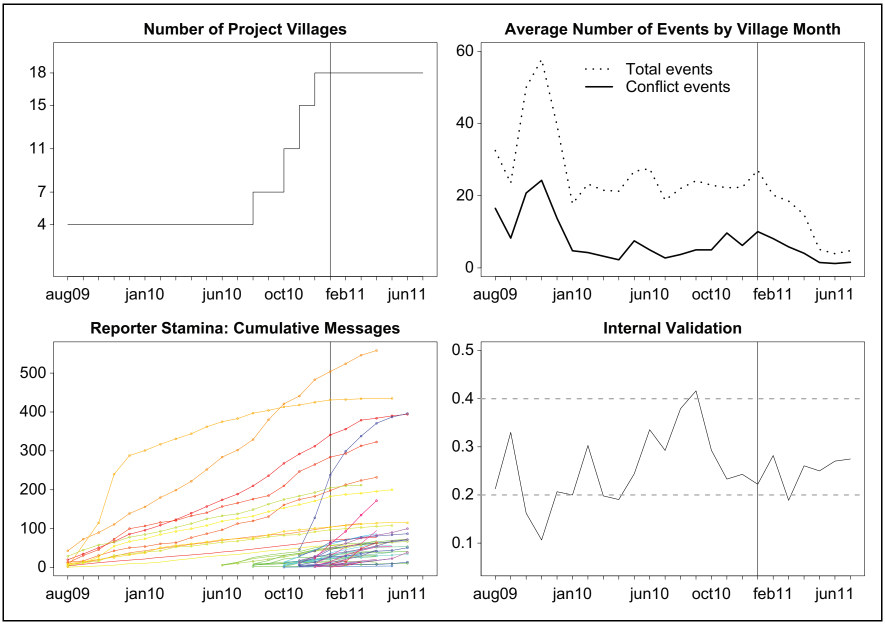

Figure 2. Top left panel shows the number of villages enrolled in the system over time; top right panel shows the average number of events per month by event type. Lower left panel shows reporter stamina—the cumulative number of events sent per reporter; this panel also indicates when reporters started and stopped sending messages. The lower right panel shows the level of "Internal Validation"—that is, the share of messages sent that reported on events that were also reported by other reporters.

kidnappings (135)—which in many cases involved forcing villagers to transport loads. We also find frequent reports of conflict-related deaths (292). Reports of sexual violence (99) indicate approximately equal numbers of men and women victims. Other major categories include outbreaks of diseases (277) and deaths due to disease (470), and crop failures (376).

Tables 2 and 3 show conditional correlations between these reported events. Specifically, each table shows the bivariate relationships between the column event and the row event as estimated by ordinary least squares. We control for village fixed effects to capture the correlation of different types of events within a village. Though we provide no causal interpretations, a number of patterns seen in Table 2 are worth highlighting. First, we see the presence of troops is not itself a good predictor of conflict-related deaths, though it is strongly related to other types of violence (a category that includes looting and forced labor). Reports of attacks are associated with all other conflict and violent events, especially conflict deaths and displacement but are only weakly associated with health events and social events. Adverse health events are not directly associated with attacks, but they are

Table 2. Correlations between Families of Events Reported by the *Voix des Kivus* Pilot.

	Presence	Attacks/ Actions	Conflict Deaths	Local Violence	Displacement	Health	Disasters	Development	Social
Presence		0.20 (0.45)	0.04 (0.61)	0.17 (0.00)	0.22 (0.16)	0.45 (0.00)	0.17 (0.12)	0.29 (0.04)	0.31 (0.04)
Attacks	0.03 (0.45)		0.16 (0.00)	0.07 (0.00)	0.15 (0.01)	0.04 (0.23)	0.10 (0.01)	0.25 (0.00)	0.04 (0.41)
Conflict deaths	0.06 (0.61)	1.99 (0.00)		0.20 (0.00)	0.43 (0.03)	0.05 (0.70)	0.24 (0.09)	0.86 (0.00)	−0.09 (0.62)
Local violence	0.74 (0.00)	2.43 (0.00)	0.53 (0.00)		1.26 (0.00)	0.66 (0.00)	1.82 (0.00)	1.82 (0.00)	1.61 (0.00)
Displacement	0.08 (0.16)	0.43 (0.01)	0.10 (0.03)	0.11 (0.00)		0.24 (0.00)	0.12 (0.07)	0.49 (0.00)	0.21 (0.02)
Health	0.44 (0.00)	0.31 (0.23)	0.03 (0.7)	0.14 (0.00)	0.63 (0.00)		0.21 (0.05)	0.27 (0.05)	0.48 (0.00)
Disasters	0.13 (0.12)	0.58 (0.01)	0.11 (0.09)	0.32 (0.00)	0.25 (0.07)	0.17 (0.05)		0.42 (0.00)	0.38 (0.00)
Development	0.13 (0.04)	0.89 (0.00)	0.24 (0.00)	0.19 (0.00)	0.61 (0.00)	0.13 (0.05)	0.25 (0.00)		0.24 (0.02)
Social	0.13 (0.04)	0.14 (0.41)	−0.02 (0.62)	0.15 (0.00)	0.23 (0.02)	0.20 (0.00)	0.20 (0.00)	0.21 (0.02)	

Note: Unit is the village-month; Numbers are estimated “marginal effects” of column reports on row reports, given village fixed effect; *p* values in parentheses. Data from month 16 to end of period. To simplify comparisons, the “Development” Category in this table excludes patrols (88) and complaints (81); the social category excludes funerals (91). Sexual violence (64 and 65) is classed under “local violence,” rather than “health.”

Table 3. Correlations between Conflict Related Events Reported by the *Voix des Kivus* Pilot.

	MONUSCO present	FARDC present	Rebels present	MONUSCO actions	FARDC actions	Rebel actions	Male deaths	Female deaths	Child deaths
MONUSCO present		0.58 (0.00)	0.24 (0.52)	0.43 (0.62)	0.06 (0.76)	−0.03 (0.87)	0.07 (0.58)	0.09 (0.32)	0.08 (0.62)
FARDC present	0.78 (0.00)		0.08 (0.86)	−0.57 (0.56)	0.04 (0.88)	0.20 (0.43)	0.00 (1.00)	0.02 (0.88)	−0.06 (0.75)
Rebels present	0.02 (0.52)	0.00 (0.86)		1.43 (0.00)	0.17 (0.00)	−0.02 (0.70)	0.04 (0.22)	0.04 (0.08)	0.03 (0.47)
MONUSCO actions	0.01 (0.62)	−0.01 (0.56)	0.27 (0.00)		0.03 (0.16)	−0.02 (0.45)	0.00 (0.83)	0.00 (0.65)	0.00 (1.00)
FARDC actions	0.01 (0.76)	0.01 (0.88)	0.61 (0.00)	0.57 (0.16)		0.08 (0.43)	0.34 (0.00)	0.28 (0.00)	0.43 (0.00)
Rebels actions	−0.01 (0.87)	0.03 (0.43)	−0.06 (0.70)	−0.29 (0.45)	0.07 (0.43)		0.09 (0.11)	0.07 (0.08)	0.10 (0.16)
Male death	0.04 (0.58)	0.00 (1.00)	0.35 (0.22)	0.14 (0.83)	0.91 (0.00)	0.27 (0.11)		0.47 (0.00)	0.49 (0.00)
Female death	0.1 (0.32)	0.01 (0.88)	0.71 (0.08)	0.43 (0.65)	1.42 (0.00)	0.42 (0.08)	0.91 (0.00)		1.34 (0.00)
Child death	0.03 (0.62)	−0.02 (0.75)	0.17 (0.47)	0.00 (1.00)	0.71 (0.00)	0.19 (0.16)	0.3 (0.00)	0.43 (0.00)	

Note: MONUSCO = United Nations Organization Stabilization Mission in the DR Congo; FARDC = Armed Forces of the Democratic Republic of Congo. Unit is the village-month; numbers are estimated "marginal effects" of column reports on row reports, given village fixed effect; *p* values in parentheses. Data from month 16 to end of period.

associated with displacement, disasters, and local violence. General development actions are associated with almost all events; and social activities (such as visits and meetings) are associated with all events except attacks and conflict deaths. Overall, there is a clear clustering of bad things going together over time.

Table 3 focuses on the more disaggregated measures and reveals a number of striking correlations. First, the presence of UN peacekeepers (MONUSCO) and government soldiers (Armed Forces of the Democratic Republic of Congo [FARDC]) tends to be highly correlated. Second, we see that rebel group presence is associated with actions by MONUSCO and FARDC. Third, we see that presence in itself is generally not strongly associated with civilian deaths—the exception being the correlation between rebel presence and female deaths. Actions by FARDC are very strongly associated with deaths, particularly women's deaths. Actions by MONUSCO are not associated with deaths and actions by rebels or unidentified groups are strongly associated with women's deaths only.

Data Quality

Voix des Kivus was implemented to assess the feasibility of using crowdseeding to gather high-quality data. In the absence of a gold standard for conflict event data, it is difficult to formally validate the approach employed here. Nevertheless, five considerations on data quality are worth emphasizing.

First, in its first year, *Voix des Kivus* employed a field coordinator to visit each of the reporters at least once every two weeks to assess the quality of the messages sent.[18] The coordinator would assess whether reporters interpreted codes in the same way as the researchers, whether respondents employed the correct codes, given the events they wished to report, and so on. Moreover, the field coordinator would verify whether events had actually taken place and whether there were other events that took place that were not reported. Throughout the pilot, the coordinator found no instances of erroneous reports of major conflict events (incursions, assaults, or killings) but did find numerous instances where events were not communicated by one or all reporters, suggesting a vulnerability to type II errors.

Second, because *Voix des Kivus* distributed phones to three people per village, we have a measure of internal validation. The bottom right panel of Figure 2 shows the share of events that were reported by at least two reporters out of all events reported. We find that internal validation was around 30 percent at the outset but increased after the project start until August 2010—the moment the project expanded to more villages.[19]

Third, we have a limited ability to compare event information received from the *Voix des Kivus* system with data received from a survey we conducted in the same region in a similar period. In the latter, we asked each village chief, how many events from a list of potential events had taken place in the village in the preceding month. Our survey villages overlapped with just three *Voix des Kivus* villages during the life span of the project.[20] We consider three ways of assessing the congruence of the survey and the *Voix des Kivus* data. First, we compare the average number of each type of event across villages under each approach. On average, we find nearly twice as many reported events according to the survey. In particular, there were higher numbers of children's deaths due to conflict and citizen deaths due to natural causes, disease outbreaks, and funerals reported in the survey. Second, despite the difference in means, the correlation between these two measures of average incidence is 46 percent which is substantively large and statistically significant. This suggests that similar types of events get reported to the same relative extent under the two systems. Third, we generate measures of the average incidence of conflict and nonconflict events for each of the three villages using each of the two data sources, yielding twelve data points.[21] The correlation between the *Voix des Kivus* and survey estimates is low largely because of one outlying observation in which a village reported a much higher incidence of army visits under the *Voix des Kivus* system. If this outlier is dropped, the correlation between the *Voix des Kivus* and the survey-based measure is 95 percent. These correlations are encouraging, although we note that survey measures are also subject to measurement error and bias, and so on their own, the correlations do not establish the reliability of the data.[22]

Fourth, a problem for any data collection method in conflict areas, including surveys, is censorship. For example, informants might report fewer conflict events precisely because of the presence of armed groups and high levels of conflict. In practice, our reporters provided detailed data on events even when troops were present or in the vicinity. Messages such as "FARDC passed by [*Village Omitted*] and shot with guns at 5:30 am" and "A FARDC soldier shot a woman" are not uncommon. And messages such as "FARDC kidnapped people who were then forced to transport things" and "FARDC kidnapped people to build a shelter for the soldiers" were received almost weekly by *Voix des Kivus*. In addition, we observe many instances in which there is simultaneous reporting of the presence of armed actors and adverse events of various kinds. Moreover, at no point did reporters indicate to us that they felt a need to censor reports.

Finally, we note a concern of uneven reporter participation. The total quantity of messages sent varies substantially across villages; from month 16 onward, total reporting varies from a low of 14 from one village to a high of 522 for another. Alongside this broad variation is a positive correlation between reports of events of many types (as seen in Tables 2 and 3). These patterns may reflect true variation, but they are also consistent with variation in the overall activity of reporters.[23] As we demonstrate in the next section, we find mixed evidence on differential behavior by reporter type. Nevertheless, we believe that heterogeneity in propensities to report could be a fruitful focus of scrutiny in future applications.

Source Dependence

A crowdsourcing system receives information from an unidentified, anonymous public. A crowdseeding system, in contrast, makes use of identifiable users.

Although the system did not gather information about the reporters themselves, the ability to analyze the "crowd" in more detail and learn whether particular types of reporters provide different information is a significant benefit.

In the case of *Voix des Kivus*, a reporter's position (chief, women's representative, or elected individual) is a natural starting point to differentiate the crowd. We find that village chiefs sent substantially more messages than the women's representative or the elected reporter: respectively, 11.62, 8.30, and 8.90 messages per person per month on average.[24]

Subdividing the data by conflict and nonconflict events, we find that around 30 percent of the messages sent concern conflict events—with elected reporters sending relatively the most.[25] Regressing the number of conflict events reported per month on the type of reporter and the number of nonconflict events reported in that month—with the chief as baseline and standard errors clustered at the village level—partially confirms this. We obtain estimates equal to 0.23 (0.63) and 1.49 (1.33) for the women's representative and the elected reporter, respectively (standard errors in parentheses). These differences in messaging by reporter position are not statistically significant however.

Reporters may nevertheless differ in the *types* of events they report. A major task of the chief is to settle disputes among villagers, particularly those over land. Another is to be the village's first point of contact for visitors such as migrants, government officials, and so on. In contrast, the head of the women's association may be the go-to person for women who experience sexual violence or domestic violence. On the other hand, our three reporters are likely to be similarly aware of events such as funerals and weddings, and activities such as road repairs and construction of the school, because Congolese villages are small. The table in the Online Appendix presents the average number of messages sent per village-month by reporter type, where we also test the difference between them. On average, chiefs send more messages regarding land conflict (code 44) than do women representatives, but fewer than elected reporters do—only the first difference is slightly significant. Chiefs—compared to elected reporters—are statistically more likely to report the arrival of refugees and internally displaced persons (code 53) and visits or meetings organized by the chiefdom or locality authorities (code 95) and by representative of a political party (code 96). We do not find that the women representatives report more domestic violence (code 45) or more sexual violence against women (code 64).

The crowd is not a random subset of the total population and in general what one learns likely depends on whom one asks. Overall, however, our analysis suggests that in these data at least, differences in reporting across the three categories of reporters in the *Voix des Kivus* villages are relatively muted.

Relation with ACLED Data

We close our discussion of data quality with a comparison of data from the system and data from what we believe is the best available alternative: the Armed Conflict

Location and Events data set (ACLED), which gathers together data on the specific dates and locations of conflict events, as well as ancillary information such as the types of event and the groups involved (Raleigh et al. 2010).[26] The data set is global in reach, and for the DRC, the data set contains 6,926 conflict events since January 1, 1997.

Because conflict clusters in isolated areas, one possible benefit of a crowdseeding system is that it records conflict events that otherwise would have gone unnoticed. To test this idea, we compare the number of conflict events reported to *Voix des Kivus* with the number of conflict events reported by ACLED for the same region and time period.[27] The ACLED database reports twenty-seven conflict events divided over the following event types: battle with no change of territory control, battle with rebel control location, battle with government regaining territory, headquarters or base establishment, nonviolent activity by a conflict actor, rioting/protesting, violence against civilians, and nonviolent transfer of location control. During this same period, *Voix des Kivus*, on the other hand, received a total of 1,439 conflict event messages, and this increases to 2,118 messages if events related to the presence of conflict actors are included.[28] This difference in number of conflict events reported is particularly large considering that *Voix des Kivus* reports events for only eighteen villages, while ACLED reports conflict events for the whole area.

Alongside these quantitative differences, there may be qualitative differences reflecting, for example, the differences in the sources upon which these data are based. The ACLED data set derives information from a variety of sources that include news reports, humanitarian agencies, and research publications.[29] *Voix des Kivus*, on the other hand, relies on reporting by community-based reporters. To examine differences, we focus on measures of *who* perpetrates violence against civilians.[30] This is a sensitive topic and different reports might under- or over-report the importance of one actor or another. Both ACLED and *Voix des Kivus* separate out violence against civilians by perpetrator. The *Voix des Kivus* data indicate that FARDC is an important perpetrator of violence against civilians, being responsible for almost 50 percent of the attacks against civilians. ACLED, on the other hand, indicates that most violence against civilians is perpetrated by rebel groups, and less than 20 percent of the violence is initiated by government soldiers.[31] This number drops to 10 percent when analyzing the ACLED data for the whole period (totaling 239 attacks against civilians).[32] The two data sets thus paint a very different picture when it comes to who perpetrates conflict against civilians.

Research Applications

A system like *Voix des Kivus* provides rich data to learn about conflict areas. Relationships between different event types can be analyzed as we did above. Particular use can be made of the data's panel component to learn about the dynamics of different events over time.[33] Another interesting application of a phone-based system is to combine it with external data sources and use the data collected by phone as an outcome measure. The local UN offices could provide historic information about patrol routes which can then be used to learn whether these patrols pacify a region or simply move violence by a rebel group from one area to another.[34] The data could be combined with satellite information about weather patterns to learn about the impact of, for example, rainfall on different types and levels of conflict. To learn about election violence, data could be collected from Congo's National Election Committee on the location of voting booths (in November 2011, the presidential elections took place in Congo and was in some areas characterized by high levels of violence). Another possibility is to use a *Voix des Kivus*-style system as a treatment in order to examine the effects of transparency on conflict outcomes of interest.

In this section, we provide two applications, where *Voix des Kivus* data provide outcome data, to illustrate the utility of this kind of data for research. First, we try to assess the conflict impact of development aid. Second, we use the data to learn about the temporal–spatial nature of conflict. In both cases, we highlight ways in which the conclusions we draw might differ from conclusions that we might draw using more traditional approaches.

Application 1: The Conflict Impact of Development Programs

There is broad recognition that development and conflict are closely related. Major development interventions focus on countries emerging from war and include reintegration, reconstruction, capacity building, and other initiatives. It remains unclear, however, if and how these investments are effective. It is expected that by improving the level of development, these projects could reduce risks of violence. Miguel, Satyanath, and Sergenti (2004), for example, find that economic growth is strongly negatively related to civil conflict. And De Ree and Nillesen (2009) argue that foreign aid has a direct negative impact on the probability of an ongoing civil conflict to continue.[35] On the other hand, the very introduction of development actors, financing, and projects could also increase violence. Nunn and Qian (2014), for example, suggest that US food aid increases the incidence, onset, and duration of civil conflicts in recipient countries and related arguments have been made by Anderson (1999), De Waal (2009), and Polman (2011). There are multiple reasons why these relations—or opposite relations—may obtain. These include changing the resources available for looting, changing the capacity of local communities to respond, changing the incentives of local communities to participate in violence, and increasing the extent to which communities are—or are believed to be—monitored by external actors.

Much of this development conflict literature is at the macro, often the national, level. Conflict research at the micro-level, and in particular, the impact of development aid on local levels of violence, is sparse and what does exist has up to recently been largely of a descriptive nature.[36] This is not surprising because attempts to find out whether these aid programs actually have an impact on the level of violence face major identification and measurement challenges (for recent studies that seek to address the identification challenge at the micro level, see Berman, Shapiro, and Felter 2011; Crost, Felter, and Johnston 2012; Beath, Christia, and Enikolopov 2011).

This section shows how crowdseeded data can be coupled with a downstream experiment to address these problems at the micro level. Between 2007 and 2011, a randomized development intervention funded by the UK's Department for International Development (DfID) was implemented by the International Rescue Committee and CARE International throughout Eastern Congo—including the province of South Kivu. The villages were selected into the development program by a public lottery—thus in expectation treatment and control villages are alike in all respects except for the intervention itself. The development project provided communities with financing of up to US$70,000 to construct local projects such as schoolrooms or clinics—part of this money was received and managed by the communities directly. The intervention followed a "do no harm" approach—that is, the projects were implemented in a way that sought to ensure that they did not spark local conflicts. During the selection of the *Voix des Kivus* villages, we took account of the random assignment of this development intervention, stratifying our sample upon the villages' treatment status in order to obtain balance between the number of *Voix des Kivus* villages that had the development program and those that did not.

Even with a small sample such as we have, random assignment ensures that simple difference in means provides an unbiased estimate of the treatment effect. Of course, with a small sample, our power is weak and we can expect our estimates to be noisy, though this does not threaten unbiasedness (e.g., Mutz and Pemantle 2011; Imai, King, and Stuart 2008).

The top line in top left panel of Figure 3 shows the estimated impact of development aid on the incidence of conflict using simple differences in means. The points are estimates calculated for each month based upon village-month data for that month and the preceding two months. The gray area indicates the 95 percent confidence interval, where the standard errors are clustered at the village level. The bottom line shows the corresponding p values.[37] Conflict events are measured as the message codes: 21 through 56, 64 through 65, and 99. Note that this construction separates the troop presence measures from the violent event measures. Two results stand out. First, we see that the magnitudes vary over time, with strong estimates of effects in early periods which then decline and reemerge, though never to initial levels. Second, the data provide evidence of a positive impact of development aid: that is, villages with the development program have a lower level of conflict incidence. Not only are the point estimates of the development impact consistently below the zero line, so are most of the 95 percent confidence intervals, which is reflected in the fact that the majority of estimated p values (thirteen of the twenty-two) are at or below .1.

A set of robustness checks, presented in the other three panels in Figure 3, give further confidence in these results. First, we provide results using a randomization

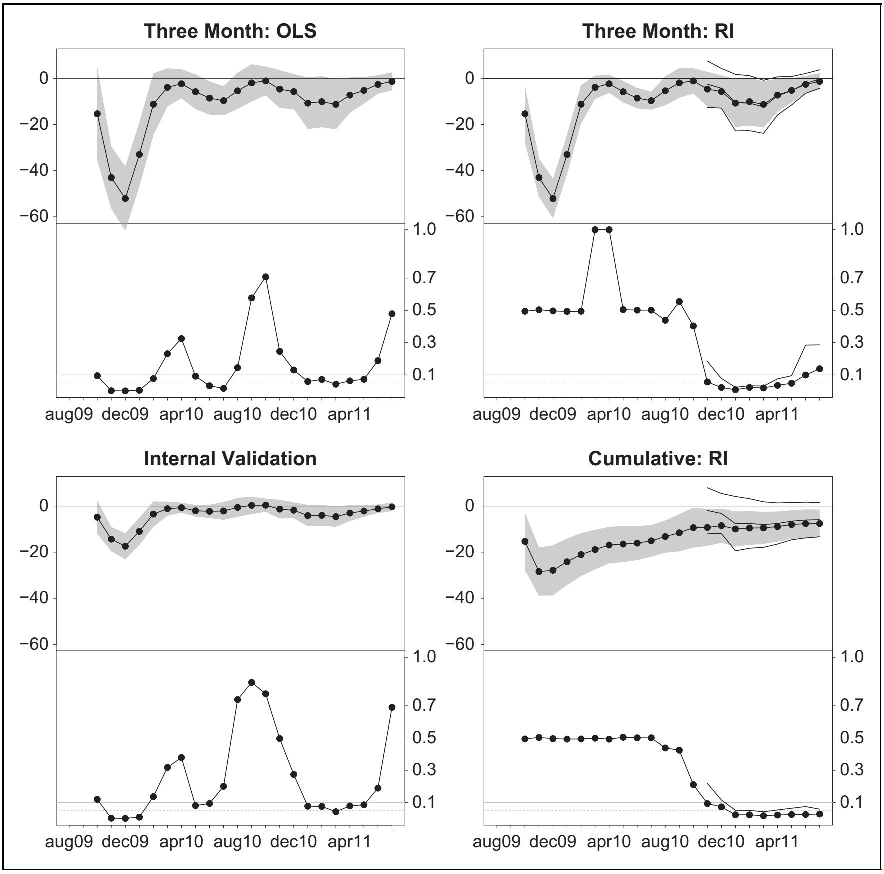

Figure 3. Conflict impact of development aid.

Note: The top left panel shows estimated effects (top black line) based on data for that month and the preceding two months, including 95 percent confidence intervals (gray area) and the corresponding *p* values (line in lower subpanel). The top right panel presents results of the randomization inference for the full sample and the restricted sample (short, thin lines). The restricted sample drops the initial four villages. The bottom left panel presents results where conflict events reported by one or two individuals are weighted one-third, or one half as much, respectively, as those reports reported by three individuals. Based on accumulated data over time, the bottom right panel shows results of randomization inference for the full sample and the restricted sample.

inference procedure (Fisher 1935).[38] We repeatedly randomly reassigned our eighteen villages to the treatment and for each of these new reassignments to treatment, we calculated the distribution of possible estimated effect under the null of no effect. We then estimated how likely we would have obtained results as strong or stronger than we did under this null. These results are presented in the top right panel of Figure 3. Using this approach, we find seven of the twenty-two of the estimated *p* values at or below .1, and five of the twenty-two below .05. Note that in the early periods in which there are only four villages, the *p* values from randomization inference can never drop below .25, correctly reflecting the severely restricted set of permutations that are possible with this small set of cases. In later periods, where it is technically possible to have lower *p* values, these values drop to below standard thresholds.

In this figure, we also include shorter, thin lines that are estimated when we restrict our analysis to exclude the first four purposefully selected villages that entered the system.[39] We find that excluding the first four *Voix des Kivus* villages makes little difference to our estimates. The estimates, confidence intervals, and *p* values are almost identical.

In a second robustness check, we leverage the possibility for internal validation of crowdseeded data. Recall that each event can be reported by one, two, or all three of our reporters. Because we are more confident that events reported by three individuals have taken place (or have been substantially important), we regard those conflict events reported by one or two individuals one-third (one-half) as much as those events reported by three individuals. Results are presented in the bottom left panel in Figure 3. The overall results are very similar to those from the main analysis although the magnitudes (*p* values) are generally somewhat lower (higher). Finally, the lower right panel of the figure shows comparable results for the cumulative effects; here again *p* values fall as the sample expands; using the full sample, the effects are strongest though this in part reflects the effects induced by the first set of villages. As shown by the thinner lines on the graph, however, the results are not very sensitive to dropping these units.

Figure 3 reports results for each month. Given this collection of findings, can we conclude that the project led to less conflict *overall*? That is, while in most periods the estimates are negative and statistically significant, there are periods in which the estimate is not significantly different from zero. This gives rise to a multiple comparisons problem. We again use randomization inference to address this problem.[40] Specifically, we repeatedly randomly reassigned our eighteen villages to the treatment to generate a distribution of estimated effects *for the entire period* under the sharp null of no effect for any unit.[41] We then assess the likelihood that we would have obtained results as strong or stronger than we did. We calculate a *p* value of .02 for the full period and .04 for the restricted period. This cumulative evidence suggests that the positive effect we estimate for the development project is unlikely due to chance even for this small sample.[42]

This analysis has been implemented using a very small number of cases from a relatively small part of Congo, and we therefore do not pretend to make a general claim here for the relation between development aid and conflict. Our goal here is to illustrate the promise of a real-time system such as *Voix des Kivus*. Our analysis points, we think, to an important advantage of real-time data for the assessment of causal effects. Existing studies of the effects of development interventions often produce contradictory results. Yet, these studies differ on many dimensions of design, including the temporal lag between intervention and measurement. An analysis using real-time data allows one to measure the extent to which *the variance in estimated effects depends on the timing of measurement*. For example, measures of conflict events that took place in the month prior to data collection might yield an estimated effect with low variance. However, while this within-period variance is low, the across-period variance might be high: very different estimates might have been produced if data were gathered a month earlier or a month later. The top left panel in Figure 3 suggests the timing of measurement contributes a large share of the true variance in estimates of treatment effects.[43] For example, a researcher collecting data in October 2010 (for that month and the preceding two months) would conclude that there is no impact of development aid ($\beta = -.9$ and $p = .74$). On the other hand, if that same researcher had conducted her study three months earlier or three months later, she would have concluded that there is strong evidence for a positive effect of development aid ($\beta = -8.7$, $p = .02$, and $\beta = -8.8$, $p = .07$, respectively).

Application 2: Spatial Clustering and Diffusion of Conflict

Conflict is often geographically clustered. One reason might be that geographically proximate units are similar to each other; another is that conflict might "spill over" from one location to the next. Population flows may be one driver for diffusion by facilitating the spread of arms, combatants, and ideologies conducive to conflict (Kuran 1998), altering an area's ethnic composition (Buhaug 2008) or exacerbating economic competition (Salehyan and Gleditsch 2006; Van Acker 2005; Claessens, Mudinga, and Ansoms 2013; Vlassenroot and Huggins 2005; Autesserre 2009, 2010).[44]

While initially limited to research on international conflict (Starr and Most 1983; O'Loughlin 1986), diffusion-based explanations have since permeated cross-national studies of civil war (Braithwaite 2010; Buhaug 2008). Disaggregated research on the local dynamics of conflict, however, is more sparse. One reason for this is the difficulty of obtaining high-quality, dynamic data on conflict at the micro level. This is unfortunate because there is reason to believe that conflict diffusion can operate at a very micro level. For example, in Eastern Congo, most population movements take place within chiefdoms, with migrants moving to nearby villages often only several kilometers away. In this second application, we illustrate the usefulness of crowdseeding data to learn about conflict diffusion at this level. Moreover, we will also illustrate the usefulness of temporal fine grain by showing how data collected following a survey-based approach can bias conclusions about conflict diffusion.

Diffusion has been an important force in the spread and duration of conflict in Congo's South Kivu province. The First Congo War (1996–1997) and the Second Congo War (1998–2003) were plausibly a direct result of the 1994 Rwandan genocide, with an estimated influx of nearly 1.5 million Rwandans.[45] Internal migration rates are also high in the region. A survey conducted in 2007 in over 600 randomly selected villages throughout Eastern Congo finds that a full 71 percent of individuals in South Kivu had fled at least once at some point during the 1996–2007 period due to armed activities by organized armed groups or militias.[46] Table 1 corroborates the presence of forced movements also for our *Voix des Kivus* villages (codes 53, 54, and 56).

Does conflict diffuse in the South Kivu province? We investigate the question by analyzing whether conflict in a village in a given period is more likely if conflict was higher in the previous period in neighboring villages. To estimate the magnitude and scope of such spillover effects, we pool our village-period data and model the diffusion process with a simple spatial autoregressive model:

$$y_{it} = \alpha + \rho \sum_{j \in N(i)} y_{j,t-1} + \beta y_{i,t-1} + \varepsilon_{it}, \tag{1}$$

where $\sum_{j \in N(i)} y_{j,t-1}$ is the spatial lag, which can be interpreted as a sum of the number of conflict events in villages in i's neighborhood, $N(i)$, in the preceding period.[47] The term $y_{i,t-1}$ is the lagged dependent variable for village i, and ε_{it} is an error term—assumed to be independent and identically distributed.[48] Note that this equation does not include village fixed effects, though these are examined subsequently.[49]

Equation (1) requires a choice about how a neighbor is defined. Are only villages less than ten kilometers away neighbors or are villages forty kilometers away also neighbors? This decision can exert a significant influence on the magnitude and scope of estimated diffusion effects (Anselin 2002; Zhukov and Stewart 2012).

In South Kivu, those individuals who migrate due to conflict often only move a few kilometers from their village of origin—moving to the nearest safe village or to those villages with family. In 2012, one of the authors collected information about the migration histories of 8,199 adults in the Buhavu chiefdom—one of the two chiefdoms in South Kivu's territory of Kalehe.[50] The data show that 92 percent of all movements that originated in the Buhavu chiefdom and were caused by conflict—which totaled 4,431 movements—ended in a different village, but in the same chiefdom. Conflict-induced migration is thus very local in South Kivu, and we might expect that conflict diffusion is as well. As a result, to test whether conflict diffusion operates at such a fine-grained level, we estimate the spatial autoregressive model multiple times, each time with a different definition of neighbor. Specifically, equation (1) is estimated with a neighbor consecutively defined as those villages inside the intervals 0 to $\{6–50\}$ kilometers.

We also test directly the benefit of a crowdseeding system to collect very precise data on the timing of conflict events. To do so, we estimate equation (1) twice: once where we do not leverage the temporal variation of the data and once where we do. Specifically, to learn about conflict diffusion, a researcher might conduct a survey twice in our *Voix des Kivus* villages, collecting in each survey the number of conflict events that took place in the previous period. With both surveys taking

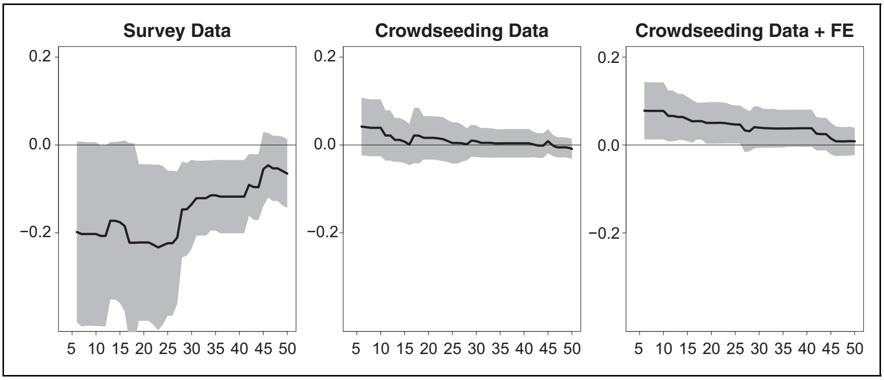

Figure 4. Conflict Diffusion in South Kivu: by type of data collected.

Note: The two left panels show results from estimating equation (1) on the survey-based data and the crowdseeded data, respectively, with errors clustered at the village level. The solid black lines indicate point estimates at one kilometer intervals. Gray areas indicate 90 percent confidence intervals. The right panel presents results from estimating equation (1) on crowdseeded data but now also controlling for village-level fixed effects. For all three panels values at the *x*-axis are the number of kilometers below which a village is defined as neighbor.

place at different points of time, the conflict events recorded in the first survey would constitute the baseline levels on the outcome variable for the subsequent period ($y_{i,t-1}$ in equation 1).[51] In contrast to such survey-style data, which collect data for two periods only, a crowdseeded data set is more fine-grained with information about the exact timing of events. To leverage this information, we undertake the same analysis but using the village-week as the unit of analysis.[52] We now compare the results based upon these different data sets.

The first two panels in Figure 4 present the results for each type of analysis. The black, solid line indicates the estimated impact of conflict in neighboring villages in the previous period on a village's incidence of conflict in this period, for a given definition of neighbor; the gray areas indicate the 90 percent confidence interval.

The left panel presents results from estimating equation (1) based on the "survey-style" data—that is, data without the temporal fine grain, clustering the error at the village level. Two results stand out. First, conflict in neighboring villages in the previous period is associated with *fewer* conflict events in this period. Second, conflict diffusion appears to operate at the micro level. An additional conflict event in a village less than twenty-five kilometers away decreases the number of conflict events by 0.2; a result that is not only statistically but also substantially significant. As we broaden the definition of neighbor, and thus include more villages, the estimated magnitudes decrease. Given these results, a researcher might conclude that migrant spillovers have a positive impact on neighboring villages. Perhaps in anticipation of responses by peacekeepers or other armed groups, rebel groups shun an area after attacking a village during the previous period.

Results are very different, however, when the temporal variation provided by crowdseeded data is taken into account. The center panel presents results from estimating exactly the same model but on the more fine-grained data. We now find evidence for conflict diffusion: that is, the point estimates are consistently *above* the zero line. We again find that previous conflict in villages nearby have a larger impact than those villages further away. An additional conflict event in a village less than ten kilometers away increases the number of conflict events by 0.05. The figure illustrates clearly how this impact decreases as we broaden the definition of neighbor. The point estimate for villages within fifty kilometers is zero.

A weakness of both of these analyses is that they fail to distinguish between true diffusion and common causes of conflict among neighbors. While we assumed ε_{it} to be independent and identically distributed, variables such as the presence of a rebel group or a bad harvest might cluster and this might drive our results. Hegre et al. (2001), for example, find that the apparent clustering of civil war is fully explained by the clustering of domestic factors (mainly gross domestic product per capita and regime type). We do not have information about such factors for our *Voix des Kivus* villages. However, a major benefit of data that has variation over time is that it allows the researcher to estimate equation (1) controlling for village-level fixed effects.[53] Results from doing so are reported in the right panel of Figure 4. We find that our point estimates almost double and that the confidence interval hits the zero bound only at 25 kilometers.

This application illustrates two points. First, it highlights the importance of getting the geographic level right. Conflict dynamics might operate at the very micro level geographically (see also Schutte and Weidmann 2011). Diffusion that can be observed at a very fine resolution may be invisible at a lower resolution. For this, crowdseeding data share strengths with data sets such as ACLED that provide fine geographic detail (a.o. Raleigh et al. 2010; Hegre, Ostby, and Raleigh 2009; Buhaug et al. 2011; Weidmann and Ward 2010), though it benefits from a greater density of reported events. Second, the analysis illustrates the importance of temporal fine grain. It illustrated how conclusions based on temporally coarse data can produce conclusions diametrically opposed to what is found with more fine-grained data. It is possible that effects work at different registers—a more sophisticated model than that provided in equation (1) might find, for example, that in the short run, there is local diffusion, but in the longer run, conflict moves farther afield. Assessing such dynamics, however, requires exactly the fine grain illustrated previously.

Conclusion

The past decade has seen much effort to better understand the causes and consequences of violent conflict. Much of the empirical analysis, however, makes use of data at high levels of aggregation (often the country and year level). It is therefore not surprising that in their survey of the civil war literature, Blattman and Miguel (2010, 46) argue that "a major goal of civil war researchers within both economics and political science in the coming years should be the collection of new data, especially extended panel micro-data sets." Original micro-level data are normally collected by conducting surveys or interviews. The collection of conflict data in particular faces a set of challenges because of the nature of the areas in which conflict events take place. These events take place in insecure areas and are therefore often off-limits to researchers, and if data are gathered long after the fact, it may suffer from various forms of recall and selection bias.

In this article, we investigated whether an SMS-based "crowdseeding" system can be used to obtain high-quality, micro-level panel data on conflict. To address the question, we piloted the *Voix des Kivus* system between 2009 and 2011 in the war-torn South Kivu province of the DRC. Reporters were selected from randomly sampled areas, provided cell phones and credit, and invited to provide regular reports to the system. This approach holds out the promise of multiple benefits for data quality, which includes a claim to representativeness and the possibility of collecting conflict event data in real time. By implementing such a crowdseeding system, we hope to overcome the problems that come with current methods to collect data on conflict.

A first objective in this research was to probe the feasibility of employing a crowdseeding approach to gathering high-quality data on conflict events in real time. We found that users had the capacity and willingness to engage at high levels, the technical implementation both at the village level and the processing level were smooth. Also the costs required to collect data—especially compared to more traditional approaches—were relatively modest.

A second goal was to illustrate the use and benefits of crowdseeding data. In a first application, we took advantage of exogenous variation resulting from the random implementation of a development project to assess whether aid is associated with increased or reduced levels of violence. We implemented *Voix des Kivus* in such a way that some *Voix des Kivus* villages received the development project and others did not. Despite the small sample, we find evidence for a negative relationship between aid and conflict. In a second application, we exploit the fine grain of the data to examine conflict diffusion patterns. We find evidence that conflict events shift from site to site but find that evidence for diffusion disappears at a resolution of about fifty kilometers. Critically, by exploiting the continuous nature of our data, both applications highlight the sensitivity of estimates to the timing of measurement in a way that cannot be assessed from traditional cross-sectional data.

Ultimately, however, the value of the system depends on the quality of the data that it generates. Our initial probes, using validation by our agent in the field, through comparison of reports across reporters, and through comparison with survey data, suggest that the data are capturing major trends faithfully. However, the data differ both qualitatively and quantitatively from that from other sources. Comparisons with ACLED data, for example, suggest a greater volume of events and a higher attribution of conflict events to Congolese government troops. Moreover, high correlations in the reporting of different types of events as well as large variation in the number of events reported by different areas suggest that accuracy may be affected by uneven reporter engagement with the system. Given the promise of the system, we believe that there is need now for more formal validation of data from crowdseeding, perhaps in conjunction with a validation of crowdsourced data, that would seek to assess rates of type I and type II errors relative to data systematically collected through intensive application of traditional methods.

We close with reflections on the ethical implications of taking a project like this to scale. During the pilot project, we faced no incidents that threatened the safety of the reporters. However, things might be different if a project of this form was scaled up and attracted the interest of armed groups. For both humanitarian and research purposes, a project such as *Voix des Kivus* becomes truly useful only when it is taken to scale, but those are precisely the conditions that might create the greatest risks. We did not assess these risks because we could not bear them ourselves. But given the importance and utility of the data, these are risks that local and international groups operating in these regions might be prepared to bear.

Authors' Note

We would like to thank Junior Bulabula, Ruth Carlitz, Simon Collard-Wexler, Yair Ghitza, Stefan Lehmeier, Steven Livingston, Neelanjan Sircar, Patrick Meier, and Lauren Young. We also thank Tracy Longacre, Desire Cimerhe, and Alain Kabeya for outstanding in-country work. *Voix des Kivus* has been presented at a number of occasions: International Conference for Conflict Mapping, UNOCHA's Coordination Meeting, Folke Bernadotte Conference, World Bank Workshop (Bukavu), Free University Berlin, MPSA, the University of British Columbia, ISA, and Bristol University. We thank the participants at these meetings. The *Voix des Kivus* pilot project was funded by USAID and benefited from support from the group's learning network; we especially thank Borany Penh and John Berry. Thanks too to the IRC for their partnership on the larger study of the development interventions in Congo. Peter also thanks Wageningen University where part of this article was written. Most of all, we thank our fifty-four reporters who, by sending thousands of SMS, made it possible to learn about events that take place in Eastern Congo.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

Funding to implement this research was received from USAID.

Notes

1. Cohen and Arieli (2011) discuss how the populations in conflict areas are hidden from and hard to reach for researchers and argue how the "snowball surveying" method works to access such populations.
2. Various approaches to get at sensitive information in these environments have recently been developed. See, for example, Corstange (2008).

3. Luyendijk (2009), for example, discusses how in the Middle East the actors provide neatly packaged story lines to media reporters. A more subtle way to influence the data is for government or rebel groups to grant a researcher access to only certain areas and deny it to others.
4. Berinsky, Huber, and Lenz (2012), for example, find that information collected via Amazon's Mechanical Turk—a particularly popular crowdsourcing system to collect data for experimental research in political science—is not representative of the wider population.
5. A Ushahidi-based crowdsourcing platform was introduced in Congo in 2008 to provide a platform to report on conflict events. The system fell out of use with only a handful of messages being received, despite many ongoing events, thus illustrating the risks of a system that depends on possibly weak supply incentives.
6. The territory is the administrative unit below the province and above the chiefdom.
7. The sampling frame and the code used to select the villages are available upon request.
8. Note that to safeguard the safety of the reporters, the village names are omitted. Moreover, locations within chiefdoms have been scrambled and so the actual location of participating villages cannot be inferred from the location of points on this map. We provided information about village location only to precleared organizations. We thank United Nations Office of Coordination of Humanitarian Affairs for the shapefiles used to construct this map.
9. The idea behind these three reporters is that the chief is the go-to person for many affairs in a village ranging from land conflict to marital disputes. The head of the women's association, on the other hand, is often the go-to person for issues such as domestic violence. Finally, one person is elected to decrease concerns of capture. These reporters can also be selected randomly to obtain a representative sample within the village.
10. All documents can be found on the project's website: www.cu-csds.org/projects/event-mapping-in-congo/. This also includes the computer code to create bulletins and a "*Voix des Kivus* Implementation Guide" for organizations that want to set up a similar system.
11. While *Voix des Kivus* started as an academic exercise, such a system also has the potential to empower participating communities and provide a basis for response to practitioners and policy advocates (Van der Windt 2013). As a result, we set up a system in which each Monday a bulletin with event information was produced and disseminated to organizations that had received clearance from *Voix des Kivus* and its reporters. Many of these organizations are based in Bukavu and have the means to respond.
12. From the introduction of these codes in August 2010 onward, the project received 3,214 event messages. A total of 1,468 had a code that indicated the sensitivity of the message. Reporters were willing to distribute the data widely, with 1,092 messages (or 74 percent of those messages that included a code) being designated for distribution to "everybody." Codes 2, 3, and 4 were used, respectively, 108, 92, and 176 times. The majority of event messages (1,751 or 54 percent of all event messages), however, did not have a security code. Users suggested that this was not driven by a lack of understanding of the system, but by their preference to leave the decision of with whom to share what data with the project.
13. Data can be found at Van der Windt and Humphreys (2014).
14. Dates that *Voix des Kivus* started: August 1, 2009 (3×); August 17, 2009, August 10, 2010, August 13, 2010 (2×); October 9, 2010 (2×); October 26, 2010, October 30,

2010, November 1, 2010, November 2, 2010 (2×); November 4, 2010, December 6, 2010, December 8, 2010, and December 10, 2010. The village that started on August 17, 2009, started later than expected because not enough people were present at the first general assembly. We had as rule that at least 40 percent of the village had to be present, be informed of the pilot project, and give their consent.

15. In fact, in January 2011, we sent a onetime US$20 of phone credit to the reporters with the message that this phone credit was for them to use in any way they want, that they would still be reimbursed for their messages, and that we would continue to collect short message service (SMS) messages and distribute the bulletins. However, we also told them that from that moment onward, they would no longer receive the weekly US$1.5, nor would they receive any feedback from the system. As can be seen from the top panel in Figure 2, the result was a gradual decline in the messages sent from January 2011 onward reaching zero by July 2011.
16. Reporting by one reporter (the upper line) is anomalous. We noticed after one month that messages by this reporter had been sent to the wrong phone number. The reporter had recorded his messages in a notebook however and subsequently uploaded them to the system (resulting in the steep slope in October 2009).
17. This is based on the Frontline SMS export of August 1, 2011.
18. To do so, *Voix des Kivus* would prepare a sheet with the messages sent by that reporter in the preceding two weeks. For security reasons, these were decoded in such a way that the information only made sense to our field coordinator.
19. The improvements and subsequent decline in validation may in part result from our communications with reporters. Consulting reporters suggested that the low internal validation early on was in part due to reporters informally adopting a division of labor where one would not send a message if they knew another had. In the early months with the first villages, we encouraged reporters not to do this and to report whether or not they thought others did also. This same messaging was not employed to the same extent with the later villages as they came online, and we see a subsequent drop in the level of internal validation.
20. And we note that even for these three *Voix des Kivus* villages, the survey took place during the final "voluntary" phase in which payments were no longer being made to reporters (see the third section).
21. This yields data of the form: $\{1, 34, 9, 10\}$, $\{1, 5, 2, 27\}$, and $\{0, 9, 7, 33\}$, denoting, respectively, the number of conflict events and nonconflict events measured by *Voix des Kivus* and the number of conflict and nonconflict events measured by the survey, by village.
22. We also note that in this case since the chief was also involved with *Voix des Kivus*, the comparison of systems involves comparing reports from the same actor, which may give rise to concerns that the two reports would be more similar than they would be if recorded from different subjects.
23. While positive correlations may be expected across many pairs of event types, we would expect a negative correlation at least for the incidence of flooding and of droughts. For this pair, we do indeed find a negative correlation, though it is small in magnitude (−.02) and there are three village-months in which both events were reported.
24. There is also a large variation over time for these three types of reporters, with standard deviations of, respectively, 14.10, 7.54, and 13.44.

25. The share of conflict events reported over total events reported is 0.28 (0.30), 0.32 (0.33), 0.34 (0.33), respectively, for the chief, the women's representative, and the elected reporter (standard deviation in parentheses). On average, they sent, respectively, 2.81 (5.89), 2.14 (2.48), 3.11 (7.31) messages concerning conflict per month.
26. Armed Conflict Location & Event Data (ACLED) can be downloaded freely at www .acleddata.com/data. The file used has been last updated on January 1, 2013.
27. That is, the period August 1, 2009, to December 31, 2010, and South Kivu's territories of Kalehe, Walungu, Mwenga, and Uvira.
28. The column "Total" in Table 1 presents this result disaggregated by event type.
29. The twenty-seven events were sourced from Africa Research Bulletin, Agence France Presse, All Africa, Associated Press Newswires, Associated Press Newswires, British Broadcasting Corporation, and Reuters.
30. ACLED collect information on the dyad characteristics of conflict events subdividing the twenty-seven events into sole military action (1×), military versus rebels (8×), military versus political militia (1×), military versus civilians (2×), military versus other (1×), sole rebel action (2×), rebels versus political militia (2×), rebels versus civilians (5×), political militia versus civilians (4×), and political militia versus others (1×). Violence against civilians corresponds closely to our definition of attack on the villages (codes 21–24). ACLED has a total of eleven such events, while *Voix des Kivus* has 103.
31. *Voix des Kivus*: United Nations Organization Stabilization Mission in the DR Congo (MONUSCO; 8, 6 percent), Armed Forces of the Democratic Republic of Congo (FARDC; 58, 46 percent), rebel (43, 34 percent), and unknown (18, 14 percent). ACLED: MONUSCO (0, 0 percent), FARDC (2, 18 percent), rebel (6, 55 percent), and unknown (3, 27 percent).
32. Specifically, this is the period February, 12, 1997, to December, 12, 2012. The breakdown by perpetrator is as follows: FARDC (25, 10 percent), Rebel (190, 80 percent), and unknown (24, 10 percent).
33. See, for example, Zeitzoff (2011) who draws data on hourly conflict intensity from Twitter and other social media to learn about the Gaza Conflict (2008–2009).
34. This is an important question that we should also ask in the "Application 1: The Conflict Impact of Development Programs" section where we analyze the conflict impact of a development program. Does this program decrease the level of violence or is violence simply moved to another area? We do not answer this question in the article. In the "Application 2: Spatial Clustering and Diffusion of Conflict" subsection, we do investigate the use of crowdseeding data to learn about conflict diffusion in more general.
35. See also Collier (2003) and Bates (2009).
36. In addition, little work exists that ties both levels together. See for a particular strong illustration of the disconnect between the micro and macro levels in the Democratic Republic of Congo: Autesserre (2010).
37. All p values in the analyses in this subsection are based upon two-tailed tests, given the uncertain expectations over the direction of the average effects of aid.
38. See also Barrios et al. (2012); Small, Have, and Rosenbaum (2008); and Ho and Imai (2006). The analyses take into account that not each village had the same probability to be selected into the Tuungane program.

39. These lines start at November 2010 for two reasons. first, as the bottom panel of Figure 2 illustrates, *Voix des Kivus'* expansion was only gradual: three villages were added in August 2010, four during October 2010, and so on. Second, the monthly estimates are based upon data from that month and the preceding two months.
40. Other approaches to deal with the multiple comparisons problem include using an index of measures (e.g., the approach adopted by Kling, Liebman, and Katz (2007)), statistical adjustments, such as the Bonferroni or Šidák corrections (see also Benjamini and Hochberg [1995], for methods to control the false discovery rate) or multilevel approaches (Gelman, Hill, and Yajima 2009).
41. Each regression is thus for the whole period under study.
42. We highlight that this test assesses the sharp null of no effect for any unit. The effects we observe are consistent with a spillover effect in which conflict shifted from treated units to control units and we cannot rule out this possibility in the current analysis.
43. Analyzing both sources of variance directly, we find that the variance in estimated effects depends strongly on the timing of measurement. Starting from month 4, the variance in estimates of the treatment effect over months 4 through 24 is 194; the average variance within each period, however, is just 12. Things settle down after month 6, but even still the variance in estimates over time (months 7–24) is 12 while the average variance is 11.
44. Another channel found to be important is the negative impact of conflict on regional economic growth, which lowers the opportunity costs of rebellion in neighboring areas (Murdoch and Sandler 2002).
45. For historical roots, see Prunier (1997, 2009).
46. See Humphreys (2008). Such conflict-induced migration has also been characteristic for the region in more recent years due to sustained rebel activity and military operations by the government such as the 2009 joint Congo-Rwanda military offensive Umoja Wetu against the Democratic Forces for the Liberation of Rwanda, followed by Kimia II and Amani Leo.
47. Note that including this summation term means the results are sensitive to the number of villages in the data set and thus controlling for temporal fixed effects is important. This is not an issue in our analysis below because the number of villages is fixed for the period chosen for analysis.
48. The interpretation is that there is the own-village lag effect and also the "neighbor" lag effect, and we are looking at the latter conditional on the former.
49. Note that there are biases from estimating fixed effects with lagged dependent variables that disappear when the number of time periods is large (Nickell 1981).
50. This area includes five of our *Voix des Kivus* villages: the upper five villages in Figure 1. The migration data can be found online: www.petervanderwindt.com/research.
51. Specifically, to simulate this setup in our data, we group six months for which we have information from all villages (November–April, 2011) into two periods each consisting of those conflict events that took place over three months.
52. Crowdseeded data can also be analyzed at a more fine-grained level such as by day or—if the data allow it—hour. We do not do so in this article.

53. Fixed effects in the presence of a lagged dependent variable runs a risk of bias, especially for short panels (Arellano and Bond 1991). An advantage of real-time data however is that panels can quickly become longer than wide.

References

Anderson, M. B. 1999. *Do No Harm: How Aid Can Support Peace—or War*. London: Lynne Rienner.

Anselin, L. 2002. "Under the Hood: Issues in the Specification and Interpretation of Spatial Regression Models." *Agricultural Economics* 27 (3): 247-67.

Arellano, M., and S. Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Employment Equations." *Review of Economic Studies* 58 (2): 277-97.

Autesserre, S. 2009. "Hobbes and the Congo: Frames, Local Violence, and International Intervention." *International Organization* 63 (2): 249-80.

Autesserre, S. 2010. *The Trouble with the Congo. Local Violence and the Failure of International Peacebuilding*. New York: Cambridge University Press.

Barrios, T., R. Diamond, G. W. Imbens, and M. Kolesár. 2012. "Clustering, Spatial Correlations and Randomization Inference." *Journal of the American Statistical Association* 107 (498): 578-91.

Bates, R. H. 2009. *Prosperity & Violence: The Political Economy of Development*. 2nd ed. New York: W. W. Norton & Company.

Beath, A., F. Christia, and R. Enikolopov. 2011. "Winning Hearts and Minds through Development Aid: Evidence from a Field Experiment in Afghanistan." Working paper.

Benjamini, Y., and Y. Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society. Series B (Methodological)* 57 (1): 289-300.

Berinsky, A. J., G. A. Huber, and G. S. Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20 (3): 351-68.

Berman, E., J. N. Shapiro, and J. H. Felter. 2011. "Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq." *Journal of Political Economy* 119 (4): 766-819.

Blattman, C., and E. Miguel. 2010. "Civil War." *Journal of Economic Literature* 48 (1): 3-57.

Braithwaite, A. 2010. "Resisting Infection: How State Capacity Conditions Conflict Contagion." *Journal of Peace Research* 47 (3): 311-19.

Buhaug, H. 2008. "Contagion or Confusion? Why Conflicts Cluster in Space." *International Studies Quarterly* 52 (2): 215-33.

Buhaug, H., K. S. Gleditsch, H. Holtermann, G. Ostby, and A. F. Tollefsen. 2011. "It's the Local Economy, Stupid! Geographic Wealth Dispersion and Conflict Outbreak Location." *Journal of Conflict Resolution* 55 (5): 814-40.

Claessens, K., E. Mudinga, and A. Ansoms. 2013. "Land Grabbing by Local Elites in the Territory of Kalehe," South Kivu, Eastern DRC. Working paper.

Cohen, N., and T. Arieli. 2011. "Field Research in Conflict Environments: Methodological Challenges and Snowball Sampling." *Journal of Peace Research* 48 (4): 423-35.

Collier, P. 2003. *Breaking the Conflict Trap: Civil War and Development Policy*. New York: Oxford University Press.

Corstange, D. 2008. "Sensitive Questions, Truthful Answers? Modeling the List Experiment with LISTIT." *Political Analysis* 17 (1): 45-63.

Coughlin, S. S. 1990. "Recall Bias in Epidemiologic Studies." *Journal of Clinical Epidemiology* 43 (1): 87-91.

Crost, B., J. H. Felter, and P. Johnston. 2012. "Aid under Fire: Development Projects and Civil Conflict." *American Economic Review* 104:1833-56.

De Nicola, F., and X. Giné. 2014. "How Accurate Are Recall Data? Evidence from Coastal India." *Journal of Development Economics* 106:52-65.

De Ree, J., and E. Nillesen. 2009. "Aiding Violence or Peace? The Impact of Foreign Aid on the Risk of Civil Conflict in Sub-Saharan Africa." *Journal of Development Economics* 88 (2): 301-13.

De Waal, A. 2009. *Famine Crimes: Politics & the Disaster Relief Industry in Africa*. Bloomington: Indiana University Press.

Fearon, J. D., and D. D. Laitin. 2003. "Ethnicity, Insurgency, and Civil War." *American Political Science Review* 97 (1): 75-90.

Fisher, R. A. 1935. *The Design of Experiments*. London: Oliver and Boyd.

Gelman, A., J. Hill, and M. Yajima. 2009. "Why We (Usually) Don't Have to Worry about Multiple Comparisons." Working paper.

Hegre, H., T. Ellingsen, S. Gates, and N. P. Gleditsch. 2001. "Toward a Democratic Civil Peace? Democracy, Political Change, and Civil War, 1816-1992." *American Political Science Review* 95 (1): 33-48.

Hegre, H., G. Ostby, and C. Raleigh. 2009. "Poverty and Civil War Events: A Disaggregated Study of Liberia Poverty and Civil War Events." *Journal of Conflict Resolution* 53 (4): 598-623.

Henderson, V., A. Storeygard, and D. N. Weil. 2011. "A Bright Idea for Measuring Economic Growth." *American Economic Review* 101 (3): 194-99.

Henderson, J. V., A. Storeygard, and D. N. Weil. 2012. "Measuring Economic Growth from Outer Space." *American Economic Review* 102 (2): 994-1028.

Ho, D. E., and K. Imai. 2006. "Randomization Inference with Natural Experiments: An Analysis of Ballot Effects in the 2003 California Recall Elections." *Journal of the American Statistical Association* 101 (475): 888-900.

Humphreys, M. 2008. *Community-driven Reconstruction in the Democratic Republic of Congo: Baseline Report*. Washington, DC: Columbia University and the International Rescue Committee.

Imai, K., G. King, and E. A. Stuart. 2008. "Misunderstandings between Experimentalists and Observationalists about Causal Inference." *Journal of the Royal Statistical Society: Series A* 171 (2): 481-502.

Kling, J. R., J. B. Liebman, and L. F. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1):83-119.

Kuran, T. 1998. "Ethnic Dissimilation and Its International Diffusion." In *The International Spread of Ethnic Conflict: Fear, Diffusion, and Escalation*, edited by D. A. Lake and D. Rotchild, 35-60. Princeton, NJ: Princeton University Press.

Lujala, P. 2010. "The Spoils of Nature: Armed Civil Conflict and Rebel Access to Natural Resources." *Journal of Peace Research* 47 (1): 15-28.

Lujala, P., N. P. Gleditsch, and E. Gilmore. 2005. "A Diamond Curse? Civil War and a Lootable Resource." *The Journal of Conflict Resolution* 49 (4): 538-62.

Luyendijk, J. 2009. *People Like Us: Misrepresenting the Middle East*. New York: Soft Skull Press.

Miguel, E., S. Satyanath, and E. Sergenti. 2004. "Economic Shocks and Civil Conflict: An Instrumental Variables Approach." *Journal of Political Economy* 112 (4): 725-53.

Murdoch, J. C., and T. Sandler. 2002. "Economic Growth, Civil Wars, and Spatial Spillovers." *Journal of Conflict Resolution* 46 (1): 91-110.

Mutz, D., and R. Pemantle. 2011. "The Perils of Randomization Checks in the Analysis of Experiments." Working paper.

Nickell, S. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49 (6): 1417-26.

Nunn, N., and N. Qian. 2014. "US Food Aid and Civil Conflict." *American Economic Review* 104 (6): 1630-66.

O'Loughlin, J. 1986. "Spatial Models of International Conflicts: Extending Current Theories of War Behavior." *Annals of the Association of American Geographers* 76 (1): 63-80.

Polman, L. 2011. *The Crisis Caravan: What's Wrong with Humanitarian Aid?* New York: Picador.

Prunier, G. 1997. *The Rwanda Crisis: History of a Genocide*. New York: Columbia University Press.

Prunier, G. 2009. *Africa's World War: Congo, the Rwandan Genocide, and the Making of a Continental Catastrophe*. Oxford, UK: Oxford University Press.

Raleigh, C., and H. Hegre. 2009. "Population Size, Concentration, and Civil War. A Geographically Disaggregated Analysis." *Political Geography* 28 (4): 224-38.

Raleigh, C., A. Linke, H. Hegre, and J. Karlsen. 2010. "Introducing ACLED: An Armed Conflict Location and Event Dataset." *Journal of Peace Research* 47 (5): 651-60.

Salehyan, I., and K. S. Gleditsch. 2006. "Refugees and the Spread of Civil War." *International Organization* 60 (2): 335-66.

Schutte, S., and N. B. Weidmann. 2011. "Diffusion Patterns of Violence in Civil Wars." *Political Geography* 30 (3): 143-52.

Small, D. S., T. R. Ten Have, and P. R. Rosenbaum. 2008. "Randomization Inference in a Group-randomized Trial of Treatments for Depression." *Journal of the American Statistical Association* 103 (481): 271-79.

Starr, H., and B. A. Most. 1983. "Contagion and Border Effects on Contemporary African Conflict." *Comparative Political Studies* 16 (1): 92-117.

Van Acker, F. 2005. "Where Did All the Land Go? Enclosure and Social Struggle in Kivu (D. R. Congo)." *Review of African Political Economy* 32 (103): 79-98.

Van der Windt, P. 2013. "From Crowdsourcing to Crowdseeding: The Cutting Edge of Empowerment?" In *Bits and Atoms: Information and Communication Technology in*

Areas of Limited Statehood, edited by S. Livingston and G. Walter-Drop, 144-56. New York: Oxford University Press.

Van der Windt, P., and M. Humphreys. 2014. "*Voix des Kivus* Data," V1, Harvard Dataverse Network. At: http://dx.doi.org/10.7910/DVN/27576.

Verwimp, P., P. Justino, and T. Bruck. 2009. "The Analysis of Conflict: A Micro-level Perspective." *Journal of Peace Research* 46 (3): 307-14.

Vlassenroot, K., and C. Huggins. 2005. "Land, Migration and Conflict in Eastern DRC." In *From the Ground Up: Land Rights, Conflict and Peace in Sub-Saharan Africa*, edited by Chris Huggins and Jenny Clover, 115-94. Pretoria, South Africa: Institute for Security Studies.

Weidmann, N. B., J. K. Rod, and L. E. Cederman. 2010. "Representing Ethnic Groups in Space: A New Dataset." *Journal of Peace Research* 47 (4): 491-99.

Weidmann, N. B., and M. D. Ward. 2010. "Predicting Conflict in Space and Time." *Journal of Conflict Resolution* 54 (6): 883-901.

Zeitzoff, T. 2011. "Using Social Media to Measure Conflict Dynamics: An Application to the 2008-2009 Gaza Conflict." *Journal of Conflict Resolution* 55 (6): 938-69.

Zhukov, Y. M., and B. M. Stewart. 2012. "Choosing Your Neighbors: Networks of Diffusion in International Relations." *International Studies Quarterly* 527 (2): 271-87.