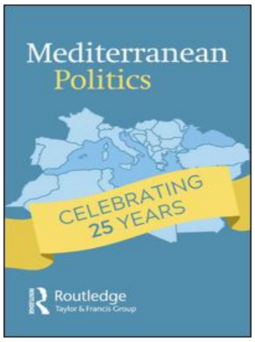

Routledge
Taylor & Francis Group

Mediterranean Politics

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/fmed20

Which protests count? Coverage bias in Middle East event datasets

Killian Clarke

To cite this article: Killian Clarke (2021): Which protests count? Coverage bias in Middle East event datasets, Mediterranean Politics, DOI: 10.1080/13629395.2021.1957577

To link to this article: https://doi.org/10.1080/13629395.2021.1957577

MEDITERRANEAN POLITICS
https://doi.org/10.1080/13629395.2021.1957577

Routledge
Taylor & Francis Group

Check for updates

Which protests count? Coverage bias in Middle East event datasets

Killian Clarke

School of Foreign Service, Georgetown University, Washington, DC, USA

ABSTRACT

Since the 2011 Arab Spring revolutions many scholars of the Middle East have built and analyzed locally-sourced protest event datasets, which have been hailed for providing superior coverage to various off-the-shelf datasets that rely primarily on English-language sources. This paper assesses the extent of these coverage improvements. It shows that across five different MENA countries, locally-sourced datasets identify considerably more events than most off-the-shelf datasets. It then compares one locally-sourced dataset of protests in Egypt from January 2012 to July 2013 to two prominent off-the-shelf datasets: ACLED and SCAD. These comparisons reveal that both ACLED and SCAD significantly overcount large, urban, violent, and political events. Next the paper compares the Egypt dataset to data compiled by two Egyptian activist groups, and finds that the locally-sourced dataset is also biased in key respects, undercounting small labor events outside the capital. Finally, the paper demonstrates the implications of these biases by showing how statistical models of protest repression differ when using the locally-sourced dataset versus SCAD. Scholars of Mediterranean politics analyzing within-case and sub-national mobilization dynamics should use locally-sourced datasets whenever possible, but should also be aware that using local sources does not entirely eliminate certain forms of bias.

KEYWORDS Protest; contentious politics; event analysis; Middle East; Egypt

I. Introduction

Over the last decade the field of Middle East political science has witnessed a marked increase in the use of event datasets to study protest and mobilization. The embrace of this technique was precipitated by the Arab Spring revolutions of 2011, which swept away or destabilized a host of long-standing dictatorships. Scholars interested in making sense of these momentous political changes saw event data as an obvious way of capturing and analysing the mobilization dynamics that had brought them about. The technique of event analysis has long been used by social scientists to study conflict and mobilization, but the method has grown more prominent

CONTACT Killian Clarke Killian.Clarke@georgetown.edu

Supplemental data for this article can be accessed here.

in recent years with the emergence of a variety of data collection projects that cover a range of countries and regions using mostly English-language sources. While these off-the-shelf datasets can be useful for studying protest comparatively or cross-nationally, after the Arab Spring scholars of the Middle East argued that the coverage gaps in these datasets were considerable, making them inadequate for the purposes of analysing dynamics of contention within a single case, particularly one experiencing revolutionary mobilization. They therefore opted to build their own datasets, relying mostly on locally-based news sources with reporters who speak and write in Arabic.

These were not the first locally-sourced datasets focusing on countries of the Mediterranean region. Before 2011, scholars had collected and analysed event data in studying protest cycles in Italy (Tarrow, 1989), labour unrest in Egypt (Beinin & Duboc, 2011), anti-colonial mobilization in Morocco (Lawrence, 2013), and anti-austerity protests in Spain (Portos, 2016). But the number and range of event data collection projects that emerged after 2011 was something new. These projects covered many of the countries that had experienced major mobilization during the Arab Spring, including: Egypt (Barrie & Ketchley, 2018; Clarke, 2020; Gunning & Baron, 2014; Ketchley, 2017; Ketchley & Barrie, 2019; Lachapelle, n.d.), Tunisia (Barrie, 2018; Berman, 2019), Morocco (Berman, 2020), Syria (Mazur, 2019, 2021), and Lebanon (Majed, 2020). And with a new wave of revolutions sweeping the region in 2019 – in Sudan, Algeria, Iraq, and Lebanon – some scholars who cut their teeth with event datasets on the Arab Spring have already begun re-deploying the method to study this more recent wave of revolutions (e.g., Berman et al., 2020).

While these datasets have been hailed for providing superior coverage to the widely-used off-the-shelf datasets, rarely has this claim been systematically examined or validated. Moreover, these locally-sourced datasets have themselves never been assessed for possible coverage bias. This paper evaluates and compares various protest event datasets to identify these forms and sources of bias. It first confirms that locally-sourced datasets across a range of MENA countries capture far more events than most off-the-shelf datasets do. It then shows that these coverage gaps are biased in important ways: the off-the-shelf datasets tend to include a higher proportion of large, urban, violent, and political events than do the locally-sourced datasets. However, the paper also finds that these locally-sourced data are themselves not devoid of bias: they tend to undercount small labour protests outside the capital. Finally, the paper discusses the implications of these biases for the purposes of social science research on protest, showing that results from quantitative analyses are likely to differ in meaningful ways depending on whether a locally-sourced or off-the-shelf dataset is used.

The main dataset in these analyses and comparisons covers Egypt from 1 January 2012 to 3 July 2013. These data were sourced from the major Arabic-language national newspaper *al-Masry al-Youm*. I built the dataset with a team of five research assistants as part of a broader project analysing dynamics of counterrevolution in Egypt and cross-nationally (Clarke, 2020). I first compare the protest counts in this and several other locally-sourced MENA datasets to protest counts in five widely-used off-the-shelf datasets: the Armed Conflict Location & Event Data Project or ACLED (Raleigh et al., 2010), the Social Conflict Analysis Database or SCAD (Salehyan et al., 2012), the GDELT Project, the Non-Violent and Violent Campaigns and Outcomes (NAVCO) 3.0 data (Chenoweth et al., 2018), and the Mass Mobilization in Autocracies Dataset or MMAD (Weidmann & Rød, 2019). After showing that these off-the-shelf datasets tend to undercount protests, I conduct a more focused comparison of distributions in event types in my Egypt dataset versus in ACLED and SCAD. Finally, I compare my Egypt dataset to data collected by two local activists groups in Egypt after the revolution.

Ultimately, this paper demonstrates that there are clear advantages to building original event datasets using local sources. Not only are off-the-shelf datasets biased in important ways, but these biases are likely non-trivial for the purposes of social science research – they almost surely affect empirical results and findings in meaningful ways. For scholars interested in studying long-term protest trends, conducting region-wide or global analyses, or making cross-country comparisons there may be no alternative to using these datasets – but awareness of their biases is crucial for avoiding erroneous inferences. For scholars of the Mediterranean region interested in conducting within-country analyses, like those that many MENA scholars undertook after the Arab Spring, these datasets are likely to be problematic, and locally-sourced datasets should be examined instead. Even then, scholars should be aware that the use of local sources improves, but does not eliminate, coverage bias; they should conduct their own assessments of these biases and think carefully about how it might shape or change their empirical findings and theoretical conclusions.

II. Coverage bias in event datasets

There is a long-standing tradition in both sociology and political science of collecting and analysing data on contentious 'events' using reporting in newspapers and other media sources.[1] The basic parameters of the approach are fairly well-defined: a predetermined corpus of media sources covering a specific period of time (or a sample thereof) are reviewed and assessed to identify events that meet particular criteria (e.g., protests, strikes, battles, terrorist attacks).[2] Researchers may identify these events using indexes, word searches, human readers/coders, or automated text analysis (Croicu &

Weidmann, 2015). Based on the information included in the media source's reporting, key variables are then coded into a standard or uniform format (e.g., location, timing, number of participants, number of casualties). The result is a cross-sectional 'catalogue' of events that can be assigned to particular time periods and/or physical locations.

Of course, even the most comprehensive set of media sources does not cover *every* event that occurs. As Earl et al. (2004) point out, these coverage gaps are not in and of themselves problematic unless the sample of events identified is in some way unrepresentative of the total population.[3] Scholars have identified a host of factors that can lead to biased media coverage. For example, news agencies tend to cover events that are closer in proximity to their headquarters, or, in the case of wire services, where they have international offices (Danzger, 1975; McCarthy et al., 1996; Mueller, 1997; Woolley, 2000). Relatedly, there is generally an urban bias in event coverage, and this bias can be particularly acute during periods of conflict, when rural locations are difficult or dangerous to reach (Kalyvas, 2004). There is also a well-documented tendency to cover larger events (Barranco & Wisler, 1999; Hendrix & Salehyan, 2015; Hug & Wisler, 1998; McCarthy et al., 1996; Oliver & Maney, 2000; Oliver & Myers, 1999). And events where there is more violence, either committed by protesters or because of police repression, tend to receive more coverage (Barranco & Wisler, 1999; Hendrix & Salehyan, 2015; Oliver & Myers, 1999).

With these findings as guidance, below I evaluate a variety of potential axes of bias in my and other event datasets covering Egypt:

- **Urban bias**: are urban events covered more than rural ones?
- **Location bias**: are certain cities or governorates (like the capital city of Cairo) covered more than others?
- **Event type bias**: are certain types of events (like political events) covered more than others?
- **Size bias**: are larger events covered more than smaller ones?
- **Violence bias**: are more violent or repressive events covered more than less-violent ones?

In the next section I provide a fuller introduction to my locally-sourced Egypt dataset and explain my strategy for comparing it to other relevant datasets along these axes.

III. Data and method

The main locally-sourced dataset I evaluate captures contentious events in Egypt from 1 January 2012 to 3 July 2013. This period represents the final eighteen months of Egypt's post-revolutionary transition, including the final six months of rule by the military's Supreme Council of the Armed Forces

(SCAF) and the full year of Mohamed Morsi's presidency, before he was ousted in a military coup on 3 July 2013. The dataset was built by me and a team of five research assistants using the daily Arabic-language newspaper *al-Masry al-Youm*. We used this newspaper because at the time of the transition it was among the most professional and independent national newspapers in Egypt, and was far less partisan than many comparable national newspapers. Moreover, in contrast to certain smaller independent newspapers, *al-Masry al-Youm* had the staff and the budget to sustain a truly national news operation, with coverage in every governorate of Egypt. Furthermore, the decision to consult *al-Masry al-Youm* follows other social scientists who have built event datasets on adjacent periods in Egyptian history, nearly all of whom used this newspaper as their sole or primary source (e.g., Barrie & Ketchley, 2018; Gunning & Baron, 2014; Ketchley, 2017; Ketchley & Barrie, 2019; Lachapelle, n.d.).[4]

The research assistants read every issue of this newspaper from 1 January 2012 to 3 July 2013, and coded all incidents that met the criteria of a 'contentious event' into an event catalogue. A contentious event was defined as a public, collective, and voluntary endeavour involving a group of people in a specific place trying to influence the actions or policies of some authority.[5] These events included protests, demonstrations, strikes, marches, sit-ins or occupations, roadblocks or blockades, boycotts, petitions, and mass attacks. They then used the information in the articles to code more than eighty variables associated with each event, including information on timing, location, turnout, demands, organizers and participants, tactics, use of violence, and repression. They identified 7,522 discrete events during this eighteen-month period, occurring over 12,576 event-days (because some events spanned multiple days).

The analyses below conduct three sets of comparisons between this dataset and other event data. First, I compare the protest counts in the Egypt dataset to protest counts in five off-the-shelf datasets that cover all or part of this period in Egyptian history: ACLED, SCAD, GDELT, NAVCO, and MMAD. Further, to demonstrate that the findings in the rest of the paper are likely to apply to other MENA cases, in this section I also compare protest counts from locally-sourced event datasets covering Morocco, Tunisia, Syria, and Iraq to the same five off-the-shelf datasets. Second, I compare distributions in the types of protests captured in my dataset to the distributions in ACLED and SCAD, which both cover the full post-revolution period in Egypt and include a number of variables that allow for direct comparison to my dataset. Finally, I compare my Egypt data to two datasets compiled, respectively, by the Egyptian activist organization Demometer and the labour advocacy NGO the Egyptian Center for Economic and Social Rights (ECESR), whose researchers relied on a broader range of local sources than *al-Masry al-Youm*. These data were collected in the aftermath of the revolution by selftaught activists and researchers learning as they went; for these reasons the data present some challenges for social science research, though they are still quite helpful for the purposes of assessing biases in my data.[6]

IV. Locally-sourced MENA data compared to off-the-shelf datasets

As a first step to showing differences between locally-sourced event datasets and existing off-the-shelf datasets we can look simply at raw event counts over time. Hendrix and Salehyan (2015) have estimated that their SCAD dataset captures 76 per cent of all events in Africa. As we shall see, the true extent of under-counting in most of these off-the-shelf datasets is considerably more than that.

I compare annual or monthly protest counts in five MENA countries: Tunisia, Morocco, Syria, Iraq, and Egypt. These countries were selected based on where I was able to gain access to locally-sourced datasets, however they also represent a fairly diverse cross-section of the region. The Egypt data is my own, described above. The Tunisia and Morocco data are from Chantal Berman's (2019) PhD dissertation on state responses to social protest before and after the Arab Spring. The Syria data are from Kevin Mazur's (2019, 2021) research on the Syrian revolution. And the Iraq data are from an ongoing data collection project by me, Chantal Berman, and Rima Majed (Berman et al., 2020). Like my Egypt data, all of these datasets focused on contentious events and relied primarily on local Arabic-language sources.

For off-the-shelf datasets I selected ACLED, SCAD, GDELT, NAVCO, and MMAD because they are among the most widely used datasets for social science research on protest and conflict in the Global South. These datasets are somewhat distinct in their source base and method of data collection. GDELT web-scrapes its data from vast amounts of English language media sources. This technique has the advantage of achieving very broad coverage, but it also results in quality issues and a tendency to over- or double-count events (see Berman, 2020 for a discussion). ACLED uses human coders, and relies on a range of international and local news sources, wire services, social media accounts, and NGO research, though its source base differs considerably from country to country (in some countries it uses local language sources, but in others it does not). SCAD, NAVCO, and MMAD also use human coders and rely on some combination of the international wire services AFP, AP, and BBC Monitoring. The datasets also provide different temporal and geographic coverage. GDELT has the broadest coverage: the whole globe dating back to 1979. ACLED covers a large number of countries, though most only for the last decade; however its coverage of Africa (where the project originated) dates back to 1997. SCAD covers Africa, Central America, and the Caribbean from 1990 to 2015. NAVCO covers 26 countries across world regions from 1991 to 2012 (in some cases this coverage is selective). And MMAD covers protest in all autocracies from 1990 to 2018. Finally, each dataset captures a range of event types, only some of which would be considered contentious events (for example, some include armed conflict events and incidents of state violence). For the comparisons below, I exclude all events that do not meet the criteria of contention used in my and the other locally-sourced datasets.[7]

Figure 1 shows annual or monthly protest counts for the following countries and periods: Egypt from January 2012 to June 2013; Tunisia from 2006 to 2016; Morocco from 2006 to 2016[8]; Syria from February 2011 to July 2012; and Iraq for two separate periods, July 2010 to June 2012 and September 2019 to March 2020. Each figure shows the locally-sourced data in dark blue, followed by the five off-the-shelf datasets. These datasets do not cover all countries and all periods, which explains why they are not represented in all figures.

Figure 1. Event counts, locally-sourced data vs. off-the-shelf data.

The main takeaway from these figures is fairly straightforward: the locally-sourced datasets include more events than the off-the-shelf datasets, sometimes by significant orders of magnitude. The datasets that rely on wire services – SCAD, NAVCO, and MMAD – have similarly sized reporting gaps. At the low end, they record zero events during periods when local sources pick up several hundreds (note, for example, Tunisia from 2006 to 2010). Rarely do they capture more than 15 per cent of events in a given period. ACLED's coverage gaps are more varied. In some countries and periods, as in Morocco and Tunisia before the Arab Spring and in Egypt during 2012, they record less than 5 per cent the number of events in the locally-sourced datasets. But for other countries, like Iraq during the 2019–2020 Tishreen Thawra, they capture more than half the number of events in the locally-sourced dataset. GDELT has the largest protest counts of all the off-the-shelf datasets and in some cases, like Syria in 2012, its counts actually exceed those in the locally-sourced datasets. But given the quality issues in GDELT, and its tendency to double-count events, such findings should be treated with caution. Another clear takeaway from these figures is that the extent of undercounting is not uniform over time. The off-the-shelf datasets tend to capture a higher proportion of events around watershed political events and/or heightened periods of mobilization – e.g., the Arab Spring period in Morocco, Tunisia, and Iraq and the run-up to the July 3 coup in Egypt. This finding makes sense when we remember that these datasets all rely primarily on wire services and/or English-language sources, which likely devote much more attention and resources to covering protests during these politically charged moments.

V. Comparison of Egypt event data to ACLED and SCAD

In this section I move beyond raw event count comparisons, and look at differences in the distribution of event *types*in my Egypt dataset versus ACLED's and SCAD's. Again, here, I am interested in probing for certain forms of bias that have been well-documented in the contentious politics and conflict scholarships, e.g., urban bias, size bias, etc. I select ACLED and SCAD for further analysis (versus the other three datasets above) for several reasons. First, SCAD codes a larger number of relevant protest variables than do other datasets, allowing for direct comparison to my data on a variety of key metrics. Moreover, because it relies on a very similar set of sources to MMAD and NAVCO (i.e., wire services) the findings from these comparisons are likely to apply to those datasets as well. I also examine ACLED because it uses a more diverse range of sources than just wire services, which means that it includes a larger number of events with a somewhat different set of characteristics.[9] GDELT's data is poorly suited for comparing distributions of event types, both because of quality problems and the limited number of relevant variables.

My dataset captures 7,522 contentious events in Egypt from 1 January 2012 to 3 July 2013, compared to 1,014 (13 per cent) events in ACLED and 593 (8 per cent) in SCAD. I begin by assessing urban bias; Figure 2 shows the distribution of events in all three datasets according to whether they occurred in an urban or a rural location.[10] All three datasets include a significant share of urban events (which does not necessarily reflect bias, given that protests tend to occur disproportionately in cities). However, whereas 24 per cent of events in my dataset occur in rural locations, rural events make up only 4 per cent of ACLED's data and 14 per cent of SCAD's data, suggesting significant urban bias in both datasets.

Another way to explore geographic bias is to compare the number and share of events that occur in each governorate of Egypt. In the appendix, I include a figure that shows these statistics for each of the three datasets. The main takeaway from these comparisons is that both ACLED and SCAD are over-counting events in Cairo. Whereas Cairo-based events represent only 27 per cent of my sample, they represent 42 per cent of ACLED's sample and 53 per cent of SCAD's. The analysis also reinforces the finding in Figure 2 that there is urban bias in these two datasets, as they both appear to undercount

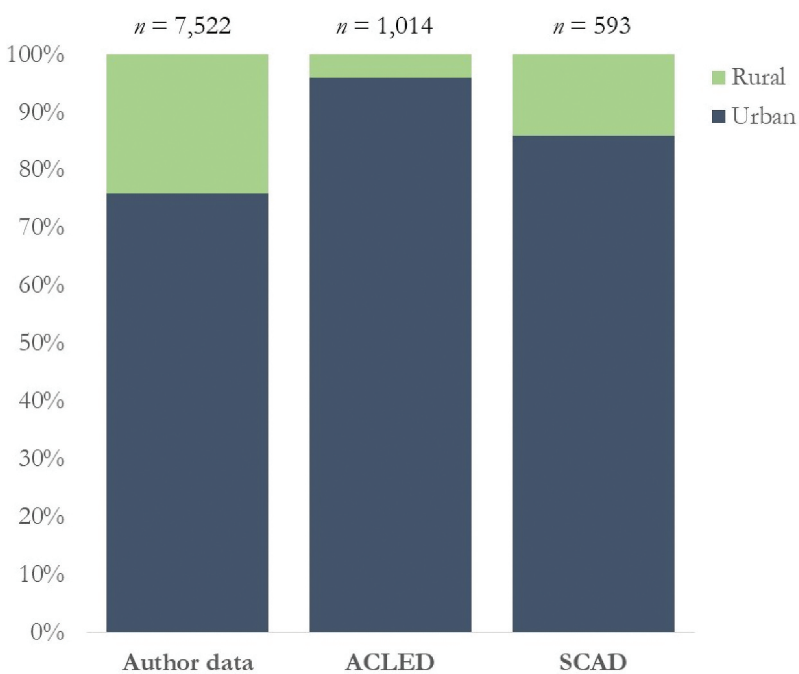

Figure 2. Urban share of events (Author data, ACLED, and SCAD).

events in more rural governorates like Dakahlia, Menoufia, Kafr el-Sheikh, and Menia.

Next, I examine several other forms of bias, this time through comparison of my dataset only to SCAD, which includes a number of event-level variables that ACLED does not code, including the number of participants, repression levels, and demands. Figure 3 plots the distribution of event sizes in the two datasets according to five categories: events of more than 1,000 participants, events with 100–999 participants, events with 10–99 participants, events with less than 10 participants, and events where the reporting did not indicate a participation number. The figure reveals that SCAD has a bias towards larger events: only 10 per cent of its events include less than 100 participants, versus 44 per cent in my dataset. In addition, SCAD has a higher share of events (37 per cent versus 25 per cent in my dataset) for which the number of participants was not reported, presumably because wire services include less rich and detailed information on events.

Next I assess the degree to which the SCAD data are biased towards events involving more violence. Figure 4 plots the distribution of events according to three main tactics: strikes, demonstrations, and riots. Figure 5 plots the distribution of events according to the level of repression: lethal repression, non-lethal repression, or no repression. We see from both these figures that SCAD's dataset has a higher proportion of violent events. More of its events (29 per cent) are riots than in the my dataset (12 per cent); it also seems to undercount strikes more than demonstrations. And the SCAD dataset also

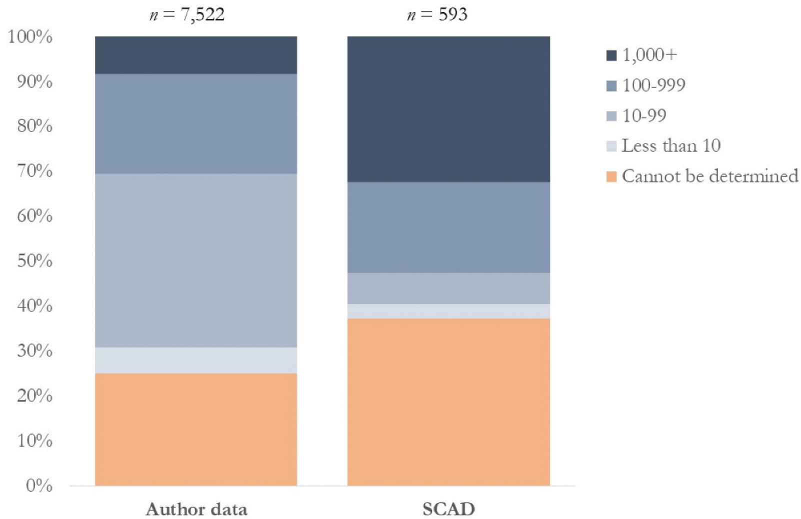

Figure 3. Distribution of events by number of participants (Author data and SCAD).

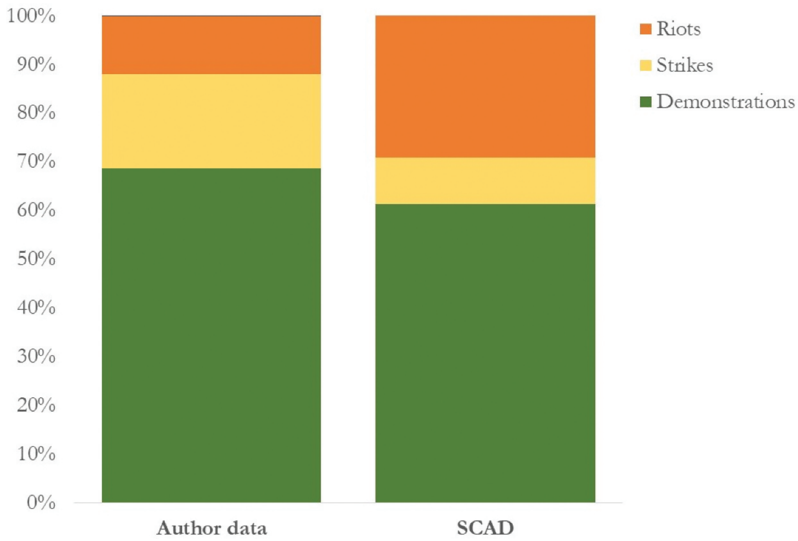

Figure 4. Distribution of events by protest tactic (Author data and SCAD).

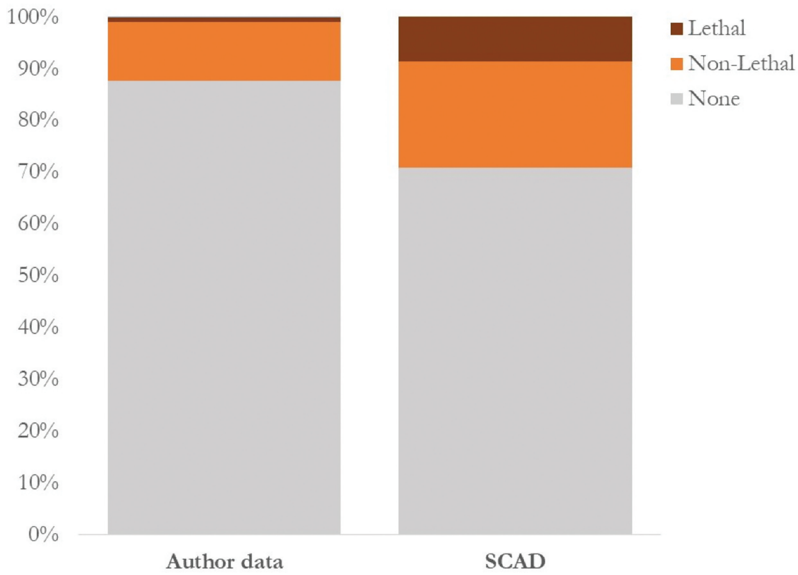

Figure 5. Distribution of events by repression level (Author data and SCAD).

disproportionately includes events that experienced repression (29 per cent versus 12 per cent in my dataset), especially those involving lethal repression (9 per cent versus 1 per cent in my dataset).

Finally, we can compare the events in these two datasets according to the types of demands they raised. Here there are some issues of commensurability, as the demand categories in SCAD are not identical to the ones in my dataset. I therefore group their demand types into five broad categories, which roughly align with mine: politics & human rights, labour, social & corruption, religion, and security.[11] As Figure 6 reveals, SCAD disproportionately includes events with demands related to religion and politics & human rights, and it tends to undercount events involving labour or social demands. This bias makes sense given the reporting priorities of international wire services, which are writing for foreign audiences that are likely to be more interested in political, human rights, and religious issues than in labour strikes or social protests over issues like electricity provision, education, and corruption.

Overall, then, these comparisons suggest that relying on a single local-language national newspaper effectively reduces a number of forms of bias that the existing literature cautions scholars to be aware of. Both ACLED and SCAD appear to overcount events in cities, especially in Cairo. Moreover, SCAD seems to be disproportionately capturing larger and more violent protests, focusing on issues of politics and human rights. However, while my dataset does seem to offer an improvement in coverage to these two datasets, there is no way of knowing from the analyses above *how much* of the bias has been reduced by using *al-Masry al-Youm*. Certainly this one newspaper does not report on all the events that occur in Egypt. In the next section I therefore turn to an assessment of just how many events – and of what type – it might be missing.

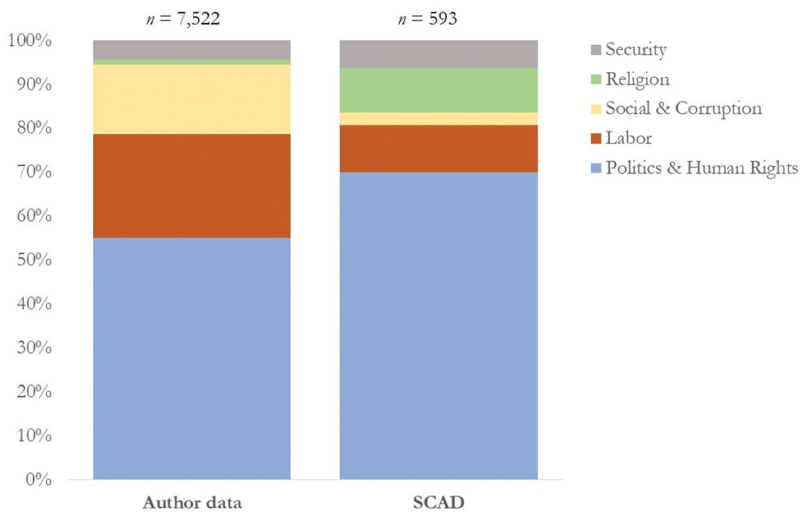

Figure 6. Distribution of events by demand type (Author data and SCAD).

VI. Comparison to two activist datasets

To assess the coverage gaps and potential biases in my dataset I compare it two datasets compiled by local Egyptian activists and researchers in the aftermath of the Egyptian revolution. The first is a research group called Demometer, which was formed in 2011 and run by a loose collective of journalists and human rights activists. They would release periodic protest reports with charts showing monthly and regional distributions of event counts, as well as breakdowns by demands and sectors (e.g., Demometer, 2013). The second is the Egyptian Center for Economic and Social Rights (ECESR), a labour rights research and advocacy NGO that was established before the 2011 revolution (for an overview of this NGO's role in the revolution see Clarke, 2014). During the post-revolution period it was active in helping workers organize, and supported the effort to build an independent labour union federation. It also launched a major data collection initiative focused on tracking labour and social protest across the country.

There are advantages and disadvantages to these groups' datasets, which I discuss in more detail below. The main advantages are that the data were collected nearly in real-time, with researchers collecting data on events shortly after their occurrence. This also allowed both research groups to leverage a more diverse set of sources than is often possible when retrospectively collecting protest data. The disadvantages are that the data were not collected using all of the rigorous and systematic standards that most social scientists would expect; these were, after all, activists learning how to collect and collate data on-the-fly, motivated by a desire to document a particularly tumultuous and important moment in Egyptian history. The ECESR data has fewer issues in this regard – its main shortcoming for the purposes of this comparison is that it only covers social and labour protests (whereas my dataset covers all protests). But the Demometer data does have problems with the consistency of its sourcing and the manner in which events were recorded. Still, both datasets are helpful for the specific purpose of diagnosing coverage biases in my *al-Masry al-Youm*-based dataset.

I begin with a comparison to Demometer, whose dataset covers broadly the same type of events as mine (i.e., all contentious evets in Egypt).[12] Demometer's data collection strategy involved having several researchers review the local news for reporting on protests that occurred the day before. In general, they would review most of the major Arabic-language national newspapers in Egypt, e.g., *al-Masry al-Youm, al-Wafd, al-Ahram, Youm Sabea, ONews, al-Shorouk, al-Tahrir, al-Badil, al-Dostour*, and *al-Watan*. However, they appear to have relied on certain sources more than others, and some sources are only used for certain months, which means that their dataset is not based on a uniform sample across a consistent source base.[13] While this issue in their sourcing strategy likely renders their data problematic for purposes of direct social science research, their dataset is still helpful for pointing towards possible coverage biases in my dataset, which relied on only one of the newspapers that they would regularly consult. After reviewing the day's newspapers, their researchers would paste the links from the relevant articles into Microsoft Word documents, along with text excerpts describing the protest events. They kept track of event counts in a spreadsheet, which they then used in their reports, but, crucially for my purposes, they never created an event catalogue with each event listed in its own row. This meant that in order to compare my data to theirs across key metrics I needed to extract information on their events from the links and text summaries in their Microsoft Word documents. While it would have been exceedingly difficult to go through all of these links (which numbered in the tens of thousands) and recreate their dataset, I used these raw data documents to conduct an assessment of event distributions for a random subset of days in both datasets.

I began by randomly selecting twelve days during the Morsi year.[14] These days are enumerated in Figure 7, with the number of events in my dataset that started on each of these days represented by the blue bars (222 total). I then turned to the Demometer documents and worked with a research assistant to assess all the events they had collected for these twelve days, using their original sources. In total they had collected 533 events across these days. We determined that 189 of these events did not qualify as contentious events by our definition. Another 135 were events that we had coded in our data (Demometer had not captured 87 of our events). This left 212 events that we had missed. These missed events are represented in

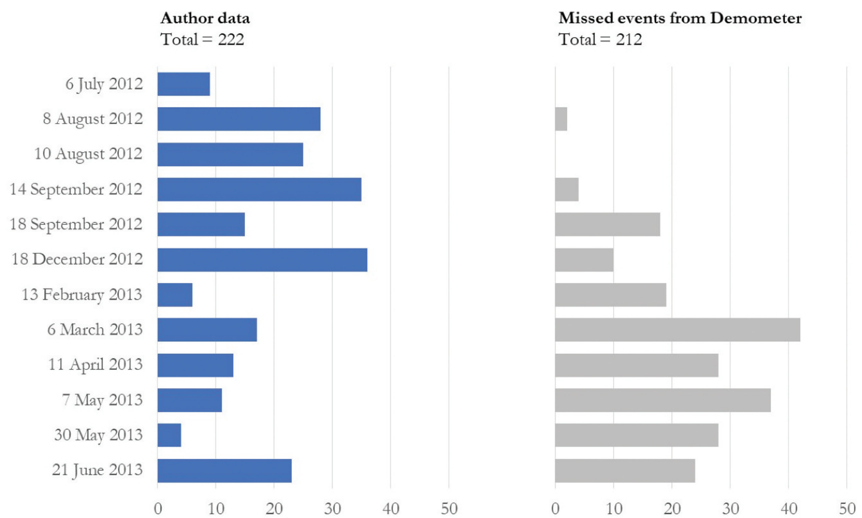

Figure 7. Daily Event Counts (Author data and Demometer).

Figure 7 with the grey bars; they imply that *al-Masry al-Youm* reported on just over 50 per cent of the contentious events that were covered by the Egyptian media during this period. Note that there appear to be more missed events in 2013 than in 2012; we believe this is because Demometer introduced new sources into its corpus (e.g., *ONews* and *al-Ahram*). The size of the reporting gap may therefore be most accurate for the 2013 dates, though there is no way to be sure based on the data we have. To allow for direct comparison to the distribution of event types in our dataset, we coded these 212 missed events according to our coding schema. The rest of the figures in this section compare the distribution of events in my data (for the randomly selected subset of days) and the distributions when the missed events from Demometer are added to my data.[15]

First I examined the geographic distribution of events by comparing the proportion of events in each of Egypt's governorates (again, this full figure is available in the appendix). This analysis revealed a bias towards Cairo-based events in my data: 19 per cent of the events for these selected days occurred in Cairo, whereas when the missed events from Demometer are added that share goes down to 14 per cent. After Cairo, the governorates where my dataset seems to be overcounting events most are Giza (which covers the western half of Cairo's urban area), Damietta, Dakahlia, and Menia. Giza is a relatively urban governorate, but the latter three are not, suggesting that it is not just urban areas that are overrepresented. This is also confirmed when we look at the governorates that are most underrepresented in my data: Gharbia, Sharkia, Fayoum, Assiut, and Alexandria. The first four are rural governorates, but Alexandria is Egypt's second largest city with a 99 per cent urban population. Overall, then, though there may be a Cairo bias in *al-Masry al-Youm*'s coverage, there does not appear to be any strong urban bias beyond that.

Figure 8 shows the distribution of my events and the combined events according to the number of participants. Here, again, I am interested in comparing the distributions for the events in my dataset that started on the randomly selected set of days and the distribution once those events are combined with the missed events from Demometer. Also, for reference, on the left I include the distribution of event sizes for my full dataset. As the figure makes clear, the events that *al-Masry al-Youm* did not report on tended to be small, mostly less than 100 people. There are also a considerable share of events where participation sizes could not be determined, which can partly be attributed to the cursory reporting in many of Demometer's sources. Overall, then, *al-Masry al-Youm* seems to cover larger events more than smaller ones.

Figure 9 examines the possibility of a violence bias. In my dataset (both the full version and the randomly sampled days) roughly 12–15 per cent of protests suffered some form of repression. But in the combined data this

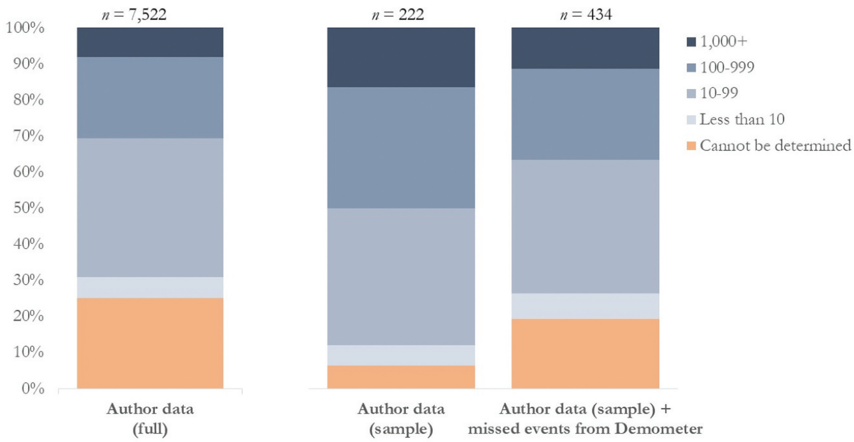

Figure 8. Distribution of events by number of participants (Author data and Demometer).

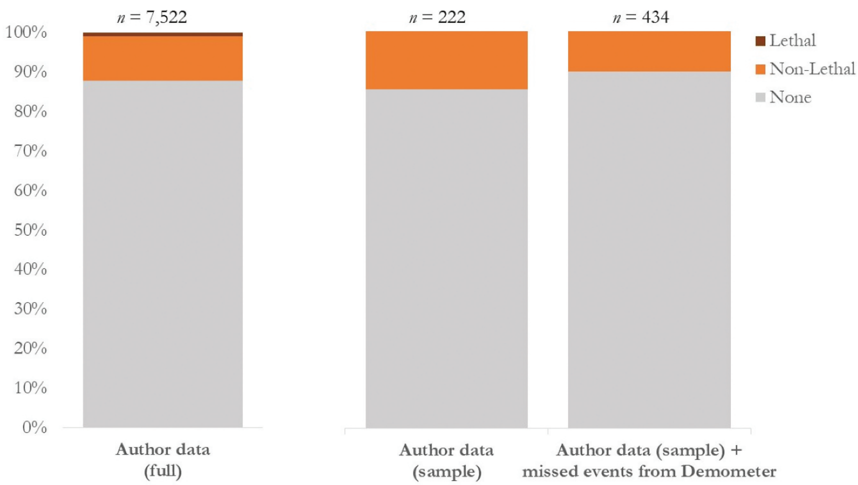

Figure 9. Distribution of events by repression level (Author data and Demometer).

proportion shrinks to 10 per cent. The events *al-Masry al-Youm* failed to report on were, in other words, mostly non-violent.

Figure 10 looks at the distribution of demands across the two datasets. The randomly sampled set of events in my dataset have roughly the same distribution of demands as the full dataset, though there are a higher share of events related to security concerns (driven by a series of protests on August 8 and August 10 following a terrorist attack in the Sinai peninsula). But when we add the missed events from Demometer the distribution shifts:

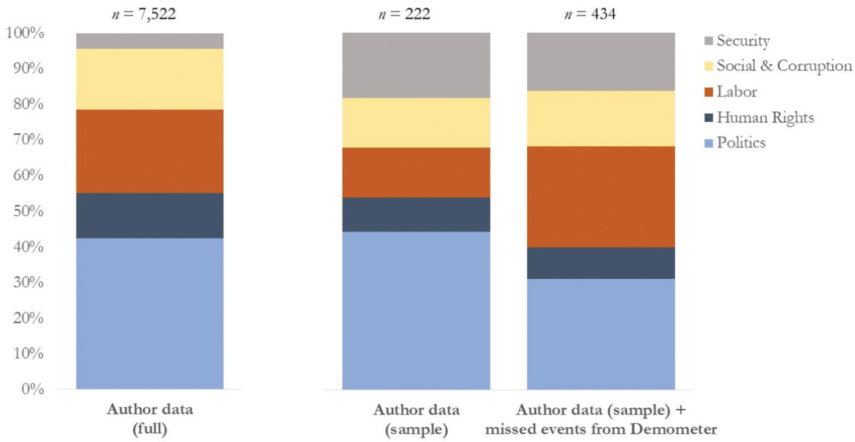

Figure 10. Distribution of events by demand type (Author data and Demometer).

the proportion of events airing labour demands increases from 14 per cent to 28 per cent and the proportion airing political demands shrinks from 44 per cent to 31 per cent. *Al-Masry al-Youm* appears to overcount political events and undercount labour events.

The finding that *al-Masry al-Youm* tends to report less on labour events is one of the main motivations for comparing my data to ECESR's dataset, which focuses only on labour and social events. Comparisons to their data can help to assess what types of labour events *al-Masry al-Youm* tends to miss. Again, ECESR is a labour NGO and advocacy group that began collecting data on labour and social protests during the post-revolutionary transition. They provided me with their event catalogue for 2013, which formed the basis of an annual report on social and labour protest during that year (Egyptian Center for Economic and Social Rights, 2013). There are 5,232 events in this dataset, 87 per cent of which occurred during the first six months of 2013. Like Demometer, ECESR relied on a variety of news sources to identify events, though they do not state explicitly which ones (their methodology section says they rely on 'a number of newspapers'). They also generate data from 'direct communication with the protesters themselves, particularly through the labour unit of ECESR, and from information issued by the protesters' (Egyptian Center for Economic and Social Rights, 2013). ECESR has very strong networks with Egypt's independent labour movement, and they appear to have relied on these ties to collect their data. The dataset they shared with me had no information on the sources for specific events, so there was no way to assess what share came from media sources versus direct ties to labourers. Moreover, their dataset unfortunately had very few variables that would allow for direct comparison to my own dataset. As a result, I simply examine the geographic and temporal distribution of their events compared to the relevant subset of my events, i.e., those that aired social or labour demands.

Figure 11 plots the number of monthly social and labour events in my dataset compared to the number in the ECESR dataset for the first six months of 2013. It also plots a line showing the proportion of events in my dataset relative to ECESR for each month. The figure suggests that *al-Masry al-Youm* is capturing between 10 per cent to 25 per cent of the social and labour events in the country.

I then examined whether *al-Masry al-Youm*'s coverage of labour and social protest exhibits meaningful geographic biases. As with the analyses above, I looked at the distribution of events in each dataset according to their governorate location (full figure available in the appendix). In line with the finding from the Demometer analysis, *al-Masry al-Youm* seems to cover a disproportionate number of labour and social protests in Cairo: the share of Cairo-based labour and social events in my dataset is 20 per cent versus 14 per cent in the ECESR dataset. However, there is also little evidence of an urban bias beyond Cairo. For example, after Cairo the governorates where events are overreported most are Menia, Suhag, Menoufia, and Giza. The latter is a more urban governorate, but the first three are some of the least urbanized governorates in Egypt. Meanwhile, the governorates with the most severe underreporting include both the relatively rural governorates of Sharkia and Assiut but also urban governorates like Suez.

Ultimately these comparisons reveal that *al-Masry al-Youm* does report on certain types of contentious events more than others. It generally tends to cover larger and more violent events that occur in the capital and that raise political demands. And it undercounts small, non-violent labour events, particularly those that occur outside of Cairo. These biases make sense in light of what we know from existing literature: as a major national newspaper with a general readership *al-Masry al-Youm* has an interest in reporting on

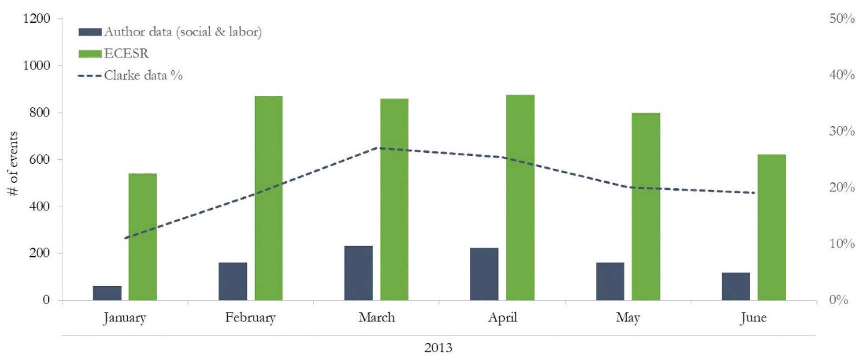

Figure 11. Monthly event count events (Author data and ECESR), Jan – Jun 2013.

more intense and political types of protests than small, parochial events outside the capital. The analysis also suggests that the biases in the off-the-shelf datasets reviewed above may be even more severe than originally suggested, since *al-Masry al-Youm* is itself biased in many of the same directions. Of course, what we cannot know from this analysis is what type of biases may exist in the sources consulted by Demometer and ECESR. For example, some analysts have accused newspapers like *Youm Sabea* or *al-Wafd*, which are important sources for the Demometer data, of deliberately inflating the number of protests at the end of Morsi's presidency in an effort to undermine him. If this accusation is correct then the *al-Masry al-Youm* data might actually be closer to 'the truth' than this analysis suggests. There is no way to know for sure, since of course there is no independent source for all the events that occurred in Egypt during this time. However, the findings from these comparisons are generally in line with the expectations in the methodological literature on event data – i.e., that coverage is weaker in national newspapers for smaller, less-violent, and less political events outside of capital cities.

VII. Implications

What types of lessons can we draw from the bias analyses above? On the one hand, we have found that a dataset of protests drawn from a major Arabic-language national newspaper in Egypt offers far superior coverage of a crucial recent period in Egypt's political history than two highly-regarded off-the-shelf datasets. Not only does it include 7 to 12 times the number of events, but it also appears to be far less biased on a range of important metrics, including event size, location, demands, and violence. On the other hand, we also learned that even this locally-sourced dataset continues to be biased across similar metrics when compared to datasets based on a larger number of local sources.

Whether and how scholars might respond to these findings depends on the types of questions they are asking and the types of analyses they hope to run. For example, Michael Biggs (2018) has recently called on scholars to examine the number of participants in protest, rather than the number of events, as their main dependent variable. Part of his argument is that focusing on turnout means that scholars need not worry as much about the under-reporting of small events, since those events do not contribute many protesters to overall turnout levels. Given this paper has confirmed that small events tend to be undercounted, scholars might choose to address the coverage biases identified here by modelling turnout rather than protest counts. However, this is, at best, only a partial solution. First, there are many questions for which protest count is the more relevant dependent variable, and most studies of protest continue to operationalize mobilization in this way. Second, as the figures above made clear, information on turnout is often unavailable for large numbers of events (25 per cent of events in my dataset, and 37 per cent in SCAD's). And, third, certain types of protest analyses cannot use size as the dependent variable – this would include, for example, event history analyses (e.g., Andrews & Biggs, 2006) and studies in which protests themselves are the unit of analysis (like studies of repression).

For scholars who either cannot or choose not to model protest size as their dependent variable, the comparisons above have suggested that the selection of which protest dataset to use is likely to be highly consequential for the types of inferences they end up drawing. To make this point, below I conduct a simple analysis of repression likelihood in my dataset and the SCAD dataset. The dependent variable is binary: whether a protest was repressed or not (lumping together lethal and non-lethal repression). The independent variables parallel those presented in Figures 2–7: whether the location of the protest was urban or rural (binary), whether the location was in Cairo or not (binary), the turnout size (categorical), the demands (categorical), and the tactics (categorical). I use logistic regressions to model this outcome across both datasets (full regression results are available in the appendix). Figure 12 shows the marginal effects from these regressions, with 95 per cent confidence intervals. For the size variable, the reference category is '1,000+'; for demands the reference category is 'Politics & Human Rights'; and for tactics the reference category is 'Demonstrations'.

Because my protest dataset contains about 12 times as many observations as SCAD's, the confidence intervals from the analysis on my data are much smaller. But not only does this model yield estimates with higher precision, but it also produces some results that are *substantively* different than SCAD's. For example, we would infer from SCAD's data that

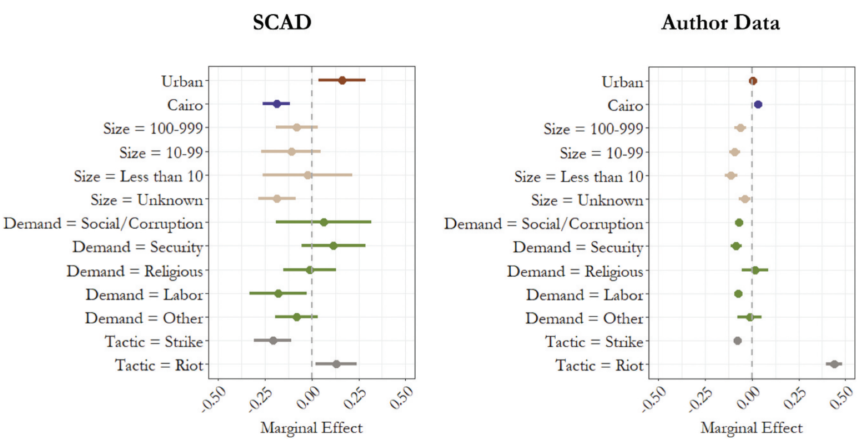

Figure 12. Marginal effects plots, likelihood of protest repression, author data vs SCAD.

urban protests are more likely to be repressed, but that protests in Cairo are less likely to be repressed. But the model on my data tells us that in fact urban protests are no more or less likely to be repressed, and that protests occurring in Cairo are actually slightly *more* likely to be repressed. These discrepant results are explicable based on what we know about SCAD's sources: the wire services it relies on generally cover more events in Cairo, but when they do report on events outside of Cairo they likely pick up more violent events occurring in other major cities. These biases lead us to the wrong set of conclusions about where repression is likely to be most severe. Other results are more similar across the two datasets, though much noisier in the SCAD analysis. For example, both models suggest that large events are repressed more, but the finding is much stronger and more consistent when using my dataset. And my data show that political events are more likely to be repressed than social, security, and labour events, whereas with SCAD's data there is only a meaningful difference with labour events. The results for tactics are most similar across the two datasets, though my data show a much stronger relationship between riots and repression.

Of course, this analysis is simple and meant only to be illustrative. But a review of some important recent protest studies that rely on off-the-shelf datasets reveal that the coverage biases found above could potentially be affecting real empirical results. For example, in a study that resembles the form and structure of the one above, Hendrix and Salehyan (2017) find that in Africa protests raising ethnic or religious claims are most likely to be repressed. But, as we learned from Figure 6, SCAD has a tendency to overcount exactly these types of protests (in my dataset they represent 1 per cent of the sample versus 10 per cent in SCAD). If the overreporting of ethnic and religious events was also due to those events' tendencies to be more violent, then the main finding in their article could be attributed at least partially to this reporting bias (though without replicating their analysis on a different dataset there is no way to know).

In another study that draws on SCAD, Dahlum and Wig (2019) argue that educated people are more likely to engage in protest, showing that protests in Africa tend to occur in areas with more educated populations. But, as we have learned, SCAD tends to oversample events in cities (where educated populations are concentrated) and political events (which tend to attract more educated, middle-class participants), raising important questions about their findings. Another example is a recent study by Baggott Carter and Carter (2020), which uses ACLED's data to show that protest levels subside following a pro-regime propaganda campaign. But such propaganda efforts are likely to have a stronger impact on political protests, and if ACLED's dataset is biased in the same way as SCAD's then political protests are likely overrepresented. The results in the paper might be considerably weaker if it included a more diverse sample of events, like labour and social protests, which would presumably be less responsive to propaganda campaigns. A final example is a study by Aidt and Leon (2016) which uses SCAD's data to argue that riots lead to moves towards democracy in Africa. However, given we know that SCAD captures more events during periods of heightened political turmoil (like democratic transitions), this result also could be driven, at least in part, by coverage bias.

VIII. Conclusion

Ultimately the findings from these comparisons and analyses provide a strong argument in favour of continuing to collect event data in the Mediterranean region from local news sources. These datasets are not devoid of their own coverage biases, as we have learned, but they do offer considerable improvements in coverage over existing off-the-shelf datasets. Of course, collecting locally-sourced data is time and resource intensive, and when conducting broad global or regional analyses there may be no alternative to datasets like ACLED and SCAD. None of the studies cited in the previous section could have been conducted on the same scale using locally-sourced data. But when scholars do choose to leverage these data they ought to be attentive to the ways in which coverage biases might be affecting their results. And for scholars interested in conducting sub-national analyses that speak to mobilization processes within one or a small number of cases, using off-the-shelf datasets is likely to lead to erroneous inferences and findings. Instead, wherever possible, these studies should be based on event datasets that draw from local sources – the payoff is clearly worth it.

Notes

1. For early examples of the method see Tilly (1995) and Kriesi et al. (1995).
2. For a more detailed overview of the method see Hutter (2014).
3. Coverage bias (or selection bias) is distinguished in the scholarship from description bias or researcher bias. Description bias refers to bias in the actual reporting of events (e.g., the use of false information), whereas researcher bias refers to biases in the process of coding or collecting events (Hutter, 2014).
4. In some MENA countries, where no single unbiased or politically neutral newspaper exists, scholars have draw on up to three local sources (e.g., Majed, 2020; Mazur, 2019), the idea being that one newspaper's biases can be counterbalanced by the others' (e.g., in Syria a pro-opposition and a pro-regime source). However, *al-Masry al-Youm* did not exhibit strong or overt editorial or political bias in its reporting during this period (though it began to do so after July 2013).

5. This operationalization is based on the understanding of contention enumerated by Doug McAdam et al. (2001).
6. Another option would be to compare my data to other locally-sourced Egypt event datasets built by social scientists, such as those cited above. However, these datasets all cover adjacent periods to mine (e.g., the year 2011 in the case of Barrie's and Ketchley's data, or the post-coup period in the case of Lachapelle's data).
7. In ACLED I use their 'protest' and 'riot' event types. In SCAD I include all event types except 'pro-government violence (repression)' and 'anti-government violence.' SCAD has two versions of its dataset: one in which events occurring in multiple locations are listed as a single event and one in which they are split out into separate events. I use the latter dataset in this and all analyses below, as this aligns with how I and the other MENA researchers conceptualized events in our datasets. I used all events in MMAD, drawing on their event-level dataset, rather than their report-level dataset. In NAVCO and GDELT, which rely on the CAMEO coding system, I used all event types with codes from 141–145. Because GDELT ends up with so many 'false positives' (e.g., whenever an article uses the verb 'demonstrate') I also follow Berman (2020) and subset the GDELT data by actor codes that indicate the presence of a civilian demonstrator.
8. In order to cover a longer period of time, Berman (2019) coded only the first five months of each year. The counts for the off-the-shelf datasets here are accordingly also only for the first five months.
9. In Egypt for the period under analysis more than 80 per cent of ACLED's events came from the following sources: the wire services AFP, AP, and Xinhua, the international news websites BBC and Africa News, and the Egyptian English-language news organizations Asway Masriya (English), Al Ahram (English), Daily News Egypt, Egypt Independent, and Nile News. They used no Arabic-language sources during this period.
10. This variable was constructed differently in each dataset. In SCAD I relied on a variable (locnum) that denotes the type of location (i.e., capital city, other major urban area). For most events, this variable indicated whether the event location was urban or rural. However, SCAD uses the location code 'nationwide' for events that occur simultaneously with other events in different locations. For these types of events, I relied on the location and event descriptions to determine whether the specific event location was urban or rural. For ACLED and my own dataset I used the district location of events to determine whether the location was urban or not. I coded all events in districts with an urban population above 50 per cent (according to the 2006 Egyptian census) as urban.
11. Politics & human rights includes the following SCAD demand categories: elections, foreign affairs/relations, human rights/democracy, pro-government. Labour includes the SCAD category economy/jobs. Social & corruption includes: food/water/subsistence, environmental degradation, education, economic resources/assets. Religion includes: ethnic discrimination/ethnic issues and religious discrimination/religious issues. And security includes: domestic war/violence/terrorism. Events with demands labelled other or unknown are excluded from the analysis.
12. There are some events in their data that I would not consider events, like planned protests that never took place. I remove these events from the analyses below.

13. For example, the vast majority of their events come from *Youm Sabea, al-Wafd*, and *al-Masry al-Youm*. But in 2013 they began using two new sources: *al-Ahram* and *ONews*. Another potential problem with their sourcing strategy is that many of the newspapers they used, like *Youm Sabea, ONews*, and *al-Wafd*, took a highly partisan, anti-government tone during the year of Morsi's presidency, and were accused of distorting and exaggerating levels of unrest in the country.
14. I used Excel's random number generator to select these days.
15. This is the most intuitive way to compare distributions, and aligns with the type of comparisons that I conducted in the previous section with ACLED and SCAD. Essentially we want to know what the distribution of event types *would* look like if we had a more complete sample (i.e., my events plus the missed events from Demometer). However, for reference, in the appendix I include the same figures using the distribution of event types across only the missed events from Demometer (i.e., before they are added into my dataset).

Acknowledgments

I thank Chantal Berman, Rima Majed, and Kevin Mazur for sharing versions of their event datasets with me. I further thank these three scholars, as well as Christopher Barrie, Mark Beissinger, Matt Buehler, Neil Ketchley, and three anonymous reviewers, for helpful feedback on this paper and for more general advice on sourcing strategies in event data collection. Finally, I thank Salah Bakr, Noah Darwish, Zahed Mukayed, Liam Sabbagh, and Firas Younes for excellent research assistance.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research was supported by a Political Science Doctoral Dissertation Research Improvement Grant from the National Science Foundation [AWD1005800] and the Social Science Research Council's International Dissertation Research Fellowship, with funds provided by the Andrew W. Mellon Foundation.

ORCID

Killian Clarke http://orcid.org/0000-0002-0265-9170

References

Aidt, T. S., & Leon, G. (2016). The democratic window of opportunity: evidence from riots in Sub-Saharan Africa. *Journal of Conflict Resolution*, *60*(4), 694–717. https://doi.org/10.1177/0022002714564014

Andrews, K. T., & Biggs, M. (2006). The dynamics of protest diffusion: movement organizations, social networks, and news media in the 1960 sit-ins. *American Sociological Review*, *71*(5), 752–777. https://doi.org/10.1177/000312240607100503

Baggott Carter, E., & Carter, B. L. (2020). Propaganda and protest in autocracies. *Journal of Conflict Resolution, 65*(5), 919–949. https://doi.org/10.1177/0022002720975090

Barranco, J., & Wisler, D. (1999). Validity and systematicity of newspaper data in event analysis. *European Sociological Review*, *15*(3), 301–322. https://doi.org/10.1093/oxfordjournals.esr.a018265

Barrie, C. A. (2018). The process of revolutionary protest: development and democracy in the Tunisian revolution. Working paper, SocArXiv.

Barrie, C., & Ketchley, N. (2018). Opportunity without organization: Labor mobilization in Egypt after the 25th january revolution. *Mobilization: An International Quarterly*, *23*(2), 181–202. https://doi.org/10.17813/1086-671X-23-2-181

Beinin, J., & Duboc, M. (2011). A workers' social movement on the margin of the global neoliberal order, Egypt 2004–2009. In J. Beinin & F. Vairel (Eds.), *Social movements, mobilization, and contestation in the Middle East and North Africa* (pp. 181–201). Stanford University Press.

Berman, C. A. (2019). *Protest, social policy, and political regimes* [PhD dissertation]. Princeton University.

Berman, C. (2020). Policing the organizational threat in Morocco. *American Journal of Political Science*. Advance online publication. https://doi.org/10.1111/ajps.12565

Berman, C., Clarke, K., & Majed, R. (2020). Patterns of mobilization and repression in Iraq's Tishreen uprising. *Frozen Conflicts*. *POMEPS Studies 41*, pp. 21–28. https://pomeps.org/patterns-of-mobilization-and-repression-in-iraqs-tishreen-uprising

Biggs, M. (2018). Size matters: Quantifying protest by counting participants. *Sociological Methods & Research*, *47*(3), 351–383. https://doi.org/10.1177/0049124116629166

Chenoweth, E., Pinckney, J., & Lewis, O. (2018). Days of rage: Introducing the NAVCO 3.0 dataset. *Journal of Peace Research*, *55*(4), 524–534. https://doi.org/10.1177/0022343318759411

Clarke, K. (2014). Unexpected brokers of mobilization: Contingency and networks in the 2011 Egyptian uprising. *Comparative Politics*, *46*(4), 379–397. https://doi.org/10.5129/001041514812522770

Clarke, K. (2020). *Overthrowing revolution: The emergence and success of counterrevolution*. [PhD dissertation]. Princeton University.

Croicu, M., & Weidmann, N. B. (2015). Improving the selection of news reports for event coding using ensemble classification. *Research & Politics*, *2*(4), 205316801561559. https://doi.org/10.1177/2053168015615596

Dahlum, S., & Wig, T. (2019). Educating demonstrators: Education and mass protest in Africa. *Journal of Conflict Resolution*, *63*(1), 3–30. https://doi.org/10.1177/0022002717721394

Danzger, M. H. (1975). Validating conflict data. *American Sociological Review*, *40*(5), 570–584. https://doi.org/10.2307/2094196

Demometer. (2013). Mu'ashar al-Ida'a al-Ihtijaji Khelal al-'Aam al-Ru'asi al-Awl.

Earl, J., Martin, A., McCarthy, J. D., & Soule, S. A. (2004). The use of newspaper data in the study of collective action. *Annual Review of Sociology*, *30*(1), 65–80. https://doi.org/10.1146/annurev.soc.30.012703.110603

Egyptian Center for Economic and Social Rights. 2013. Taqrir Al-Ihtijajat al-Sanawiya 2013.

Gunning, J., & Baron, I. Z. (2014). *Why occupy a square? : People, protests and movements in the Egyptian revolution*. Hurst.

Hendrix, C. S., & Salehyan, I. (2015). No news is good news: Mark and recapture for event data when reporting probabilities are less than one. *International Interactions*, *41*(2), 392–406. https://doi.org/10.1080/03050629.2015.982117

Hendrix, C. S., & Salehyan, I. (2017). A house divided: Threat perception, military factionalism, and repression in Africa. *Journal of Conflict Resolution*, *61*(8), 1653–1681. https://doi.org/10.1177/0022002715620473

Hug, S., & Wisler, D. (1998). Correcting for selection bias in social movement research. *Mobilization: An International Quarterly*, *3*(2), 141–161. https://doi.org/10.17813/maiq.3.2.6ptv3133154x28n5

Hutter, S. (2014). Protest event analysis and its offspring. In D. D. Porta (Ed.), *Methodological practices in social movement research*(pp. 335–367). Oxford University Press.

Kalyvas, S. N. (2004). The urban bias in research on civil wars. *Security Studies*, *13*(3), 160–190. https://doi.org/10.1080/09636410490914022

Ketchley, N. (2017). *Egypt in a time of revolution: Contentious politics and the Arab spring*. Cambridge University Press.

Ketchley, N., & Barrie, C. (2019, December). Fridays of revolution: Focal days and mass protest in Egypt and Tunisia. *Political Research Quarterly*, *73*(2), 308–324. https://doi.org/10.1177/1065912919893463

Kriesi, H., Koopmans, R., Duyvendak, J. W., & Giugni, M. (1995). *New social movements in Western Europe: A comparative analysis*. U of Minnesota Press.

Lachapelle, J. (n.d.). *Rules of repression: Military regimes, coercion, and autocratic survival*. Unpublished manuscript.

Lawrence, A. K. (2013). *Imperial rule and the politics of nationalism: Anti-colonial protest in the French empire*. Cambridge University Press.

Majed, R. (2020, October 2). In defense of intra-sectarian divide: street mobilization, coalition formation, and rapid realignments of sectarian boundaries in Lebanon. *Social Forces, 99*(4), 1772–1798. https://doi.org/10.1093/sf/soaa076

Mazur, K. (2019). State networks and intra-ethnic group variation in the 2011 Syrian uprising. *Comparative Political Studies*, *52*(7), 995–1027. https://doi.org/10.1177/0010414018806536

Mazur, K. (2021). *Revolution in Syria: Identity, networks, and repression*. Cambridge University Press.

McAdam, D., Tarrow, S., & Tilly, C. (2001). *Dynamics of contention*. Cambridge University Press.

McCarthy, J. D., McPhail, C., & Smith, J. (1996). Images of protest: dimensions of selection bias in media coverage of Washington demonstrations, 1982 and 1991. *American Sociological Review*, *61*(3), 478–499. https://doi.org/10.2307/2096360

Mueller, C. (1997). International press coverage of East German protest events, 1989. *American Sociological Review*, *62*(5), 820–832. https://doi.org/10.2307/2657362

Oliver, P. E., & Maney, G. M. (2000). Political processes and local newspaper coverage of protest events: From selection bias to triadic interactions. *American Journal of Sociology*, *106*(2), 463–505. https://doi.org/10.1086/316964

Oliver, P. E., & Myers, D. J. (1999). How events enter the public sphere: Conflict, location, and sponsorship in local newspaper coverage of public events. *American Journal of Sociology*, *105*(1), 38–87. https://doi.org/10.1086/210267

Portos, M. (2016). Taking to the streets in the context of austerity: A chronology of the cycle of protests in Spain, 2007-2015. *Partecipazione E Conflitto*, *9*(1), 181–210. https://doi.org/10.1285/i20356609v9i1p181

Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research*, *47*(5), 651–660. https://doi.org/10.1177/0022343310378914

Salehyan, I., Hendrix, C. S., Hamner, J., Case, C., Linebarger, C., Stull, E., & Williams, J. (2012). Social conflict in Africa: A new database. *International Interactions*, *38*(4), 503–511. https://doi.org/10.1080/03050629.2012.697426

Tarrow, S. G. (1989). *Democracy and disorder: Protest and politics in Italy, 1965-1975*. Clarendon Press.

Tilly, C. (1995). *Popular contention in Great Britain, 1758-1834*. Harvard University Press.

Weidmann, N. B., & Rød, E. G. (2019). *The Internet and Political Protest in Autocracies*. Oxford University Press.

Woolley, J. T. (2000). Using media-based data in studies of politics. *American Journal of Political Science*, *44*(1), 156–173. https://doi.org/10.2307/2669301