# Supplemental Information for "Crowd Cohesion and Protest Outcomes" (for online publication only)

April 5, 2022

## Contents

# A Details on Implementing Natural Language Processing

There are different tools for implementing natural language processing, which are built on different mathematical models. The tool I use, word2vec (meaning "word to vector"), has a particularly good track record in terms of mimicking human coding (Altszyler et al. 2016; Naili, Chaibi and Ben Ghezala 2017). This is "a single-hidden-layer, fully-connected, feed-forward neural network that has been shown to learn the meanings of words given their contexts in natural language texts. Word2vec produces word vectors, in the form of real-valued vectors, from raw text input in a process called embedding" (Radford 2021, 161). Word2vec solves analogy tasks by trying all words in a vocabulary until it maximizes an equation representing the similarity between words' meanings. For further details, see Mikolov, Chen, Corrado and Dean (2013), Mikolov, Sutskever, Chen, Corrado and Dean (2013), Rong (2014), and Church (2016). The word2vec model I use was trained on GenSim (Řehůřek and Sojka 2010) with Skip-Gram negative sampling, a window of five words, and 300 dimensions.

Word2vec is imperfect, particularly when handling negation words, as in the following examples: "Raise taxes!" and "Do *not* raise taxes!" These two sentences are semantically similar, though they convey opposite meanings. Such noise is rather common among natural language processing techniques (Hu et al. 2018). However, this may not be a major concern. Pang, Lee and Vaithyanathan (2002) modeled the effects of negation words on the performance of machine learning algorithms for classifying documents by general sentiment (positive or negative). Under one condition of their experiment, they tagged every word between a negation word (e.g., "not," "isn't") and the first punctuation mark following the negation word. Preliminary findings indicated that removing the negation tag had a slightly harmful, but overall negligible, effect on accuracy. Kiritchenko, Zhu and Mohammad (2014) further note that negation words do not always reverse the meaning of a word; negations tend to change the *meaning* of positive words but only the *intensity* of negative words. For example, "not good" conveys that something is bad, whereas "not terrible" still expresses that something is bad, but less bad than simply "terrible." Additionally, a scan of the CCC survey responses revealed very few negation words. In short, I expect any negation words in the survey responses to have some effect on the performance of word2vec but not to the degree that they severely undermine my estimates of cohesion.

To efficiently calculate the cohesion of all survey responses from each protest in the CCC database (not just two responses at a time, as in Table 2), I employed the Tool for Automatic Analysis of Cohesion (TAACO), which converts word vectors into document vectors. See Crossley, Kyle and McNamara (2016) for details. This is an online platform that let me batch-process text files and generate a cohesion score per event, ranging from a possible 0 (least cohesive) to 1 (most cohesive).

## A.1 Checking Robustness to an Alternative Processing Technique

I tested the robustness of my results to using an alternative technique for natural language processing, Latent Semantic Analysis (LSA). This model employs a slightly different way of calculating textual cohesion vis-à-vis word2vec (Altszyler et al. 2016; Naili, Chaibi and Ben Ghezala 2017).

Rather than relying on a neural network, LSA analyzes a document term matrix through singular value decomposition. I used TAACO to calculate both LSA and word2vec scores for each protest. See Table B.1 for summary statistics. Figure A.1 reproduces the results from the main analysis using word2vec as the cohesion measure, and Figure A.2 shows results using LSA instead of word2vec. Results are robust in the first two models, though not after adding all controls.

Figure A.3 translates the marginal effects into predicted probabilities of a concession at different levels of word2vec and LSA respectively. The pattern is similar in both panels, although LSA predicts slightly higher probabilities than word2vec.

Figure A.1: Average Marginal Effects of Crowd Cohesion on Probability of Concession (word2vec Models)

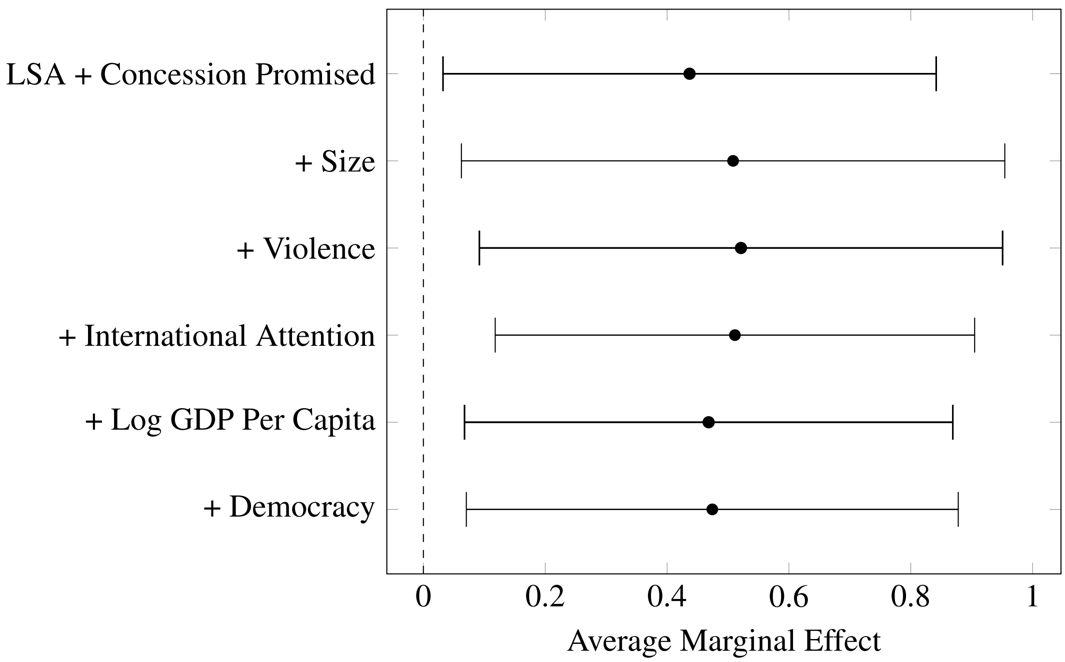


Points show marginal effects of semantic cohesion of protester demands (as measured by word2vec) on the probability of a concession. Bars show 95% confidence intervals. Each row adds different covariates to the preceding row's model. Models include country fixed effects.

Figure A.2: Average Marginal Effects of Crowd Cohesion on Probability of Concession (LSA Models)

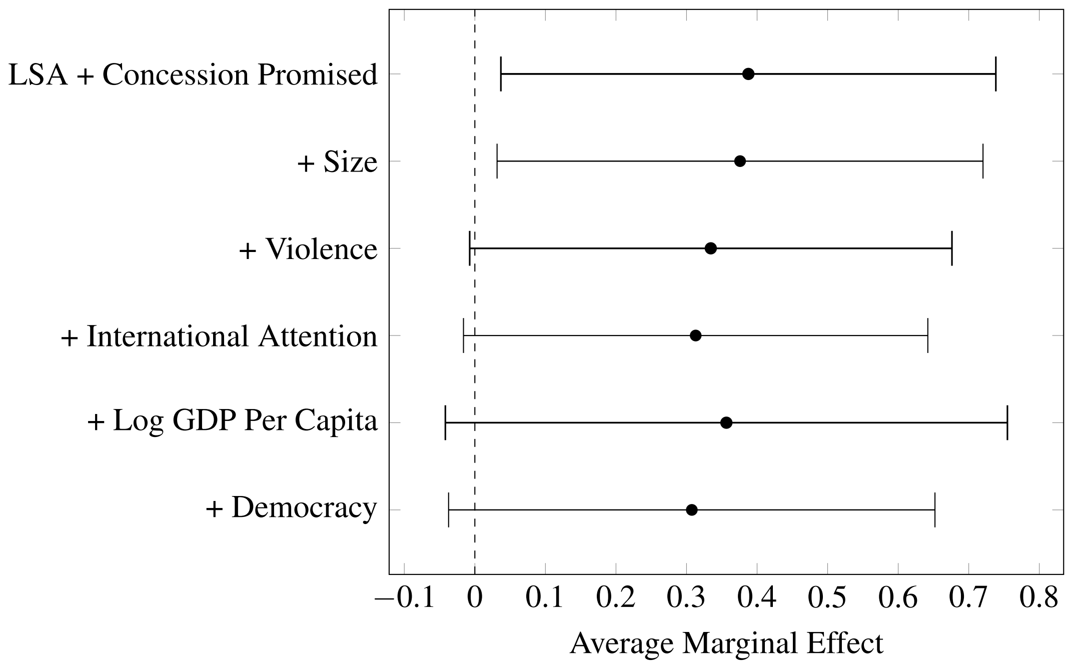


Points show marginal effects of semantic cohesion of protester demands (as measured by LSA) on the probability of a concession. Bars show 95% confidence intervals. Each row adds different covariates to the preceding row's model. Models include country fixed effects.

Figure A.3: Predicted Probabilities of Concession at Different Levels of Crowd Cohesion: word2vec vs. LSA Results

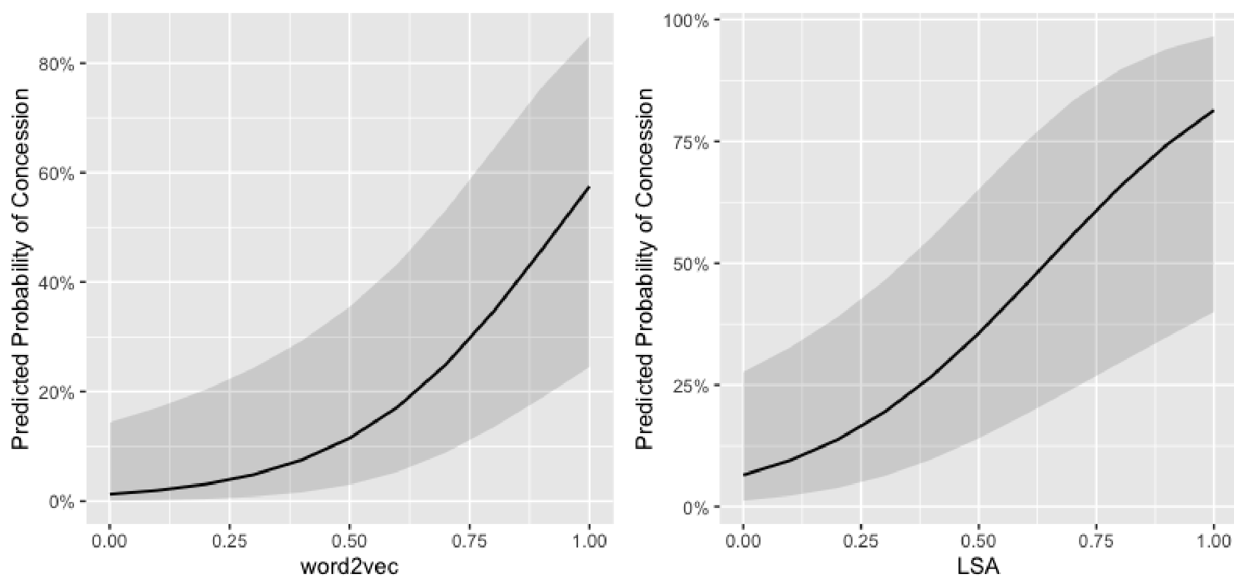


Predicted probabilities, with 80% confidence intervals, of a concession at increasing levels of semantic cohesion of protester demands. Semantic cohesion is measured by word2vec (left panel) and LSA (right panel). Probability of *Concession Promised* held at its mean; country set to Belgium.

# B Summary Statistics

| Variable | Min | Max | Mean | Median | Std. Dev. |
|---|---|---|---|---|---|
| Concession | 0.000 | 1.000 | 0.206 | 0.000 | 0.407 |
| Semantic Cohesion (Word2Vec) | 0.000 | 1.000 | 0.626 | 0.642 | 0.209 |
| Semantic Cohesion (LSA) | 0.000 | 1.000 | 0.396 | 0.368 | 0.193 |
| Concession Promised | 0.000 | 1.000 | 0.206 | 0.000 | 0.407 |
| Size (0-1000) | 0.000 | 1.000 | 0.247 | 0.000 | 0.434 |
| Size (1,001-10,000) | 0.000 | 1.000 | 0.381 | 0.000 | 0.488 |
| Size (10,001-100,000) | 0.000 | 1.000 | 0.320 | 0.000 | 0.469 |
| Size (100,001-million) | 0.000 | 1.000 | 0.052 | 0.000 | 0.222 |
| Violence | 0.000 | 1.000 | 0.113 | 0.000 | 0.319 |
| International Attention | 0.000 | 1.000 | 0.125 | 0.000 | 0.332 |
| Log GDP per Capita | 9.630 | 10.980 | 10.440 | 10.460 | 0.320 |
| Electoral Democracy Index | 0.691 | 0.919 | 0.863 | 0.873 | 0.070 |

Table B.1: Summary Statistics

# C Theorizing the Temporal Window of Observation

In the statistical analysis of CCC data, I chose three years as the cutoff for observing protest concessions, because this was the longest period after the most recent protests that provided adequate news coverage for coding concessions (the most recent protest occurred in 2014, and coding occurred in 2018). Some scholars have used smaller temporal windows for observing protest outcomes. For instance, Chenoweth and Stephan (2012, 42) consider a protest campaign a success if it achieves its stated goals within a year of peak activities. Franklin (2015, 70) measures concessions two months after a contentious challenge. There is always some element of arbitrariness in setting a cutoff (Gamson 2015, 385), but the general consensus in the literature is that larger windows are better. Gamson (2015, 385) argues that extending the period of observation allows analysts to include concessions that they might miss by using a premature termination date. Tarrow (1998, 217) concurs, noting that protests sometimes put reforms on the policy agenda that only later manifest as concrete changes. In the very long term, a given protest can leave organizational and cultural "residues" that eventually mobilize entirely new "cycles of contention" (Tarrow 1998, 217) and reinvent a society's political culture (Allam 2018, 6). While my empirical analysis would ideally capture such downstream effects, it was limited by the likelihood that recent protests had not yet actualized their full consequences. Indeed, as Allam (2018, 6) maintains, protests may shape society almost interminably, in ways that are difficult to perceive, let alone measure. The more temporally distant a protest, the less eager journalists become to report on how power holders responded (Huet-Vaughn 2013, 11). To consistently code all events in my dataset with adequate confidence in the quality of news sources, I therefore chose the longest observable span of time, which was three years. Echoing Allam (2018), I do not claim that three years marks the definitive endpoint of a cycle of resistance; I merely propose this window as long enough to record medium-term concessions across all cases in the sample.

# D Experimental Sample Details

The survey firm Dynata administered the survey experiment online to South African residents. Participants accessed the survey through email invitations, saved links, or Dynata's panel member web pages. 1080 people started the survey, and 99 percent qualified based on Dynata's quotas for balancing the survey against the national population. Ultimately, 1051 respondents completed the survey. Table D.1 compares the experimental sample to the national population. National statistics on race are from the 2011 South African Census; other national statistics are from the the more recent 2018 Afrobarometer Survey, which was designed to be nationally representative. The table shows that the sample is comparable to South Africa's population, although white and urban people are somewhat over-represented.

Table D.1: National Representativeness of Experimental Sample

| | Sample Percent | National Percent |
|---|---|---|
| **Gender** | | |
| Male | 46 | 49 |
| Female | 54 | 50 |
| **Age** | | |
| 15-25 | 16 | 23 |
| 26-35 | 33 | 30 |
| 36-45 | 24 | 19 |
| 46-55 | 12 | 14 |
| 56+ | 10 | 14 |
| **Race** | | |
| Black | 56 | 76 |
| White | 23 | 9 |
| Coloured | 13 | 9 |
| Asian | 7 | 3 |
| Other | 2 | 1 |
| **Location** | | |
| Urban | 85 | 69 |

# E Balance Tests

I test for balance between treatment and control groups by regressing treatment status on demographic covariates. The results in Table E.1 show no significant differences across groups. *Education* was a scale from 0 (no formal schooling) to 8 (post graduate education). *Income* was a respondent's self-reported household income decile relative to other households.

Table E.1: Relationship between Treatment Status and Covariates

| | DV = Cohesive Crowd Treatment |
|---|---|
| Age | $-0.00$ |
| | $(0.01)$ |
| Male | $0.02$ |
| | $(0.13)$ |
| Asian | $-0.88$ |
| | $(0.59)$ |
| Coloured | $-1.07$ |
| | $(0.57)$ |
| White | $-0.75$ |
| | $(0.56)$ |
| Black | $-0.89$ |
| | $(0.55)$ |
| Education | $-0.09$ |
| | $(0.06)$ |
| Income | $-0.00$ |
| | $(0.03)$ |
| Intercept | $1.42^{*}$ |
| | $(0.68)$ |
| Num. obs. | 1051 |

Standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

# F Experimental Results with and without Covariates

Table F.1: Experimental Results Are Robust to Omitting Covariates

| | DV = Vote for Tax | | DV = Recall Demands | |
|---|---|---|---|---|
| Cohesive Crowd Treatment | $0.33^{**}$ | $0.38^{**}$ | $0.02$ | $0.04$ |
| | $(0.16)$ | $(0.16)$ | $(0.14)$ | $(0.15)$ |
| Age | | $-0.01$ | | $0.02^{**}$ |
| | | $(0.01)$ | | $(0.01)$ |
| Male | | $0.14$ | | $-0.11$ |
| | | $(0.16)$ | | $(0.15)$ |
| Asian | | $0.99$ | | $0.26$ |
| | | $(1.09)$ | | $(0.65)$ |
| Coloured | | $1.11$ | | $-0.12$ |
| | | $(1.06)$ | | $(0.62)$ |
| White | | $0.52$ | | $0.31$ |
| | | $(1.06)$ | | $(0.61)$ |
| Black | | $1.66$ | | $0.03$ |
| | | $(1.04)$ | | $(0.60)$ |
| Education | | $0.04$ | | $0.27^{***}$ |
| | | $(0.07)$ | | $(0.07)$ |
| Income | | $0.09^{**}$ | | $-0.06^{*}$ |
| | | $(0.04)$ | | $(0.04)$ |
| Intercept | $-1.63^{***}$ | $-3.40^{***}$ | $1.11^{***}$ | $-0.84$ |
| | $(0.11)$ | $(1.17)$ | $(0.10)$ | $(0.77)$ |
| Num. obs. | 1087 | 1051 | 1049 | 1049 |

Standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

# G Ethical Considerations for the Experiment

My experiment was approved by the Institutional Review Board of my home institution (approval #022101) on February 4, 2021. Here, I affirm the ways in which my research practices relate to the principles of consent, deception, confidentiality, harm, and impact.

## G.1 Enrollment Criteria

Any South African who opted into the Dynata survey panel was eligible for the study, while anyone who did not opt in was ineligible. The panel was self-selected, and then Dynata balanced the sample on key demographics (age, race, gender, location) to the national population using a sampling frame. The opt-in nature of the panel ensured that all respondents were consenting adults (Dynata screened out underage respondents). I did not deliberately sample any vulnerable persons (prisoners, high-risk pregnant people, etc.).

## G.2 Power Differentials and Compensation

Although I did not deliberately sample vulnerable persons, there is an inherent power differential between me, a researcher based at a U.S. institution, and the South African residents who answered the survey. South Africa is an upper-middle income country but has higher income inequality than most countries in the world (Maluleke 2019). Therefore, my survey reached some economically disadvantaged people. Approximately 30 percent of the respondents earn less than 50,000 South African Rand (3,355 U.S. Dollars) per year.

Dynata does not directly compensate respondents monetarily, but rather offers points that respondents can redeem for prizes such as cash rewards and gift cards. Dynata is a global firm that works with local partners in different countries, including South Africa. All of my transactions during survey enumeration were with Dynata and not the local partner, who remained anonymous.

While South Africa is a democracy where political debate and market research are common, I ensured that none of the survey questions were politically or socially sensitive. For example, I asked only whether a respondent would support hypothetical protesters, not whether they personally participated in the kind of protest depicted in the experimental treatment. Although Dynata collected contact information to issue rewards for participating in the survey, I had no access to such identifying information.

## G.3 Consent and Deception

Informed consent was behavioral, consisting of clicking "yes" in response to the following question: "This survey will ask you about public affairs and basic information about yourself. It will take approximately five minutes to complete. There are no anticipated risks to participating, your responses will remain confidential, and you are free to opt out at any time. Do you agree to participate?" All participants gave consent in the same way. I did not request signed consent, to protect respondent anonymity.

I did not deceive respondents, but I used incomplete disclosure in order to avoid biasing answers. The above introduction of the survey stated the research goals only in the broadest terms, so as to not prime respondents.

## G.4 Impact

Beyond the incentives that Dynata offered to individual survey respondents, I aim for this research to benefit South Africans more generally. In keeping with the decolonial research practice of "reporting back" (Smith 2012), I intend to submit condensed versions of the results to popular media in South Africa, similar to an explainer in the "Monkey Cage" blog of the *Washington Post*.

# References


Allam, Nermin. 2018. *Women and the Egyptian Revolution: Engagement and Activism during the 2011 Arab Uprisings*. Cambridge: Cambridge University Press.

Altszyler, Edgar, Mariano Sigman, Sidarta Ribeiro and Diego Fernándo Slezak. 2016. "Comparative Study of LSA vs Word2Vec Embeddings in Small Corpora: A Case Study in Dreams Database." *Computing Research Repository* DOI: 10.1016/j.concog.2017.09.004.

Chenoweth, Erica and Maria J. Stephan. 2012. *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. New York: Columbia University Press.

Church, Kenneth Ward. 2016. "Emerging Trends: Word2Vec." *Natural Language Engineering* 23(1):155–162.

Crossley, Scott A., Kristopher Kyle and Danielle S. McNamara. 2016. "The Tool for the Automatic Analysis of Text Cohesion (TAACO): Automatic Assessment of Local, Global, and Text Cohesion." *Behavior Research Methods* 48:1227–1237.

Franklin, James C. 2015. "Persistent Challengers: Repression, Concessions, Challenger Strength, and Commitment in Latin America." *Mobilization* 20(1):61–80.

Gamson, William A. 2015. Defining Movement 'Success'. In *The Social Movements Reader: Cases and Concepts*, ed. Jeff Goodwin and James M. Jasper. Malden: Wiley Blackwell pp. 383–385.

Hu, Qinmin, Jie Zhou, Qin Chen and Liang He. 2018. "Promoting Word Sentiment and Negation in Neural Sentiment Classification." *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence* https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16505.

Huet-Vaughn, Emiliano. 2013. "Quiet Riot: The Causal Effects of Protest Violence." SSRN working paper https://ssrn.com/abstract=2331520.

Kiritchenko, Svetlana, Xiaodan Zhu and Saif M. Mohammad. 2014. "Sentiment Analysis of Short Informal Texts." *Journal of Artificial Intelligence Research* 50:723–762.

Maluleke, Risenga. 2019. "Inequality Trends in South Africa." Statistics South Africa Report No. 03-10-19. http://www.statssa.gov.za/?page_id=1854&PPN=Report-03-10-19&SCH=7680 .

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado and Jeffrey Dean. 2013. "Distributed Representations of Words and Phrases and their Compositionality." *https://arxiv.org/abs/1310.4546* .

Mikolov, Tomas, Kai Chen, Greg Corrado and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." *https://arxiv.org/abs/1301.3781* .

Naili, Marwa, Anja Habacha Chaibi and Henda Hajjami Ben Ghezala. 2017. "Comparative Study of Word Embedding Methods in Topic Segmentation." *Procedia Computer Science* 112:340–349.

Pang, Bo, Lillian Lee and Shivakumar Vaithyanathan. 2002. "Thumps Up? Sentiment Classification Using Machine Learning Techniques." *Proceedings of the Conference on Empirical Methods in Natural Language Processing* pp. 79–86.

Radford, Benjamin J. 2021. "Automated Dictionary Generation for Political Eventcoding." *Political Science Research and Methods* 9:157–171.

Řehůřek, Radim and Petr Sojka. 2010. "Software Framework fro Topic Modelling with Large Corpora." *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks* DOI:10.13140/2.1.2393.1847.

Rong, Xin. 2014. "Word2Vec Parameter Learning Explained." *Computing Research Repository* arXiv: 1411.2738.

Smith, Linda Tuhiwai. 2012. *Decolonizing Methodologies*. New York: Zed Books.

Tarrow, Sidney. 1998. *Power in Movement: Social Movements and Contentious Politics*. Cambridge: Cambridge University Press.