

Dirty Wars, Databases, and Indices

C. Maria Keet

Abstract

Hicks, M.H.-R. & Spagat, M. (2008). The Dirty War Index: A public health and human rights tool for examining and monitoring armed conflict outcomes. *PLoS Medicine* 5(12), e243. <http://dx.doi.org/10.1371/journal.pmed.0050243>

Peace research looks at the topics of what makes a society peaceful as well as how to go from a (post-) war situation to a peaceful society *sensu* positive peace. Concerning the latter, it is important to have an understanding about the (preceding) conflict, which includes the kind of atrocities committed and the extent of these acts. It is known to be difficult to find and interpret such data in line with the reality. Hicks and Spagat (2008) tried it nevertheless, not only with a few examples, but devising a specific measure, the *Dirty War Index* (DWI), to synthesise into one value the dirtiness of any given war, which is intended as “a data-driven public health tool that identifies rates of particularly undesirable or prohibited, i.e., “dirty,” outcomes inflicted on populations during war (e.g., civilian death, child injury, or torture)” (p1). To qualify the intention of what is “dirty”, the authors link the DWIs explicitly “to international humanitarian law to make public health outcomes directly relevant to prevention, monitoring, and humanitarian intervention for the moderation of war’s effects” (p1). The principal calculation is [Number of “dirty,” i.e., undesirable or prohibited cases/Total number of cases] x 100, where the dirty/normal cases can follow along the lines of existing conventions and laws, such as with respect to child soldiers, torture, prohibited weapons, and so forth.

The idea of multiple DWIs can indeed be interesting—an opinion amplified by the editorial of *PLoS’s* flagship journal (Sondorp, 2008)—but only to give a *rough indication* and the ratio should be taken in conjunction with the absolute numbers, because there is the real danger of misuse due to unfair sanitation of data. It can easily fall victim to, or even stimulate intended, misinterpretation by showing some neat aggregated numbers without having to assess the source data and by brushing over the reality on the ground that a bean-counting person may not be aware of and more readily can set aside in favour of the aggregated numbers. Taback (2008), in the same journal issue, already voiced reservations about the feasibility of the DWI by focusing on issues with data collection and statistical issues, and on selection bias, missing data, and censoring in particular; similar problems have been observed and documented elsewhere as well (Weinstein, 2007). I would like to add two more issues to these reservations. One is to demonstrate the easiness of “colouring” a DWI to support one’s political preference through two different approaches, the other is to expose the hidden issue of the *databases themselves* on which tools such as the DWI rely.

Let us first take at face value the database(s) on which the data for calculating a DWI relies. Hicks and Spagat provide a relatively extensive example for Colombia (text and table 1, p2, in the paper), which supposedly illustrate the usefulness of the DWI. One of their DWIs is the measure of [nr. of civilians killed] / [total nr. of civilians killed + combatants killed] * 100. The “guerillas” (presumably FARC) have a DWI of 2498/5444*100=46, the “government forces” 593/659*100=45, and the “illegal paramilitaries” 6944/6985*100=99, where these numbers as reported in the article are taken from the Colombia conflict database (CERAC 2009), which the authors are involved in.^[1] Hicks and Spagat explain that “Guerrillas rank 2nd in killing absolute numbers of civilians”, as if the government forces deserve a laurel for having the best (closest to 0) DWI—with a mere 1-point margin—and as if paramilitaries are

independent of the government whereas it is the norm, rather than the exception, that governments tend to arrange for a third party to do the dirty work for them (with or without external funding) so as to look comparatively good in the international spotlights. For the sake of illustration, let us aggregate by “opponents of FARC”, then we obtain a DWI of $[593+6944]/[659+6985] * 100 = 98.6$, which is substantially more dirty than FARC, which cannot be explained away anymore by data collection biases. Put differently, FARC is, in both DWI and absolute numbers, the “cleaner” one and one could use this DWI to argue FARC has good reason to be annoyed with the current violent governance in the country. Such a difference in results of calculating the same DWI that is entirely dependent on preferred (dis)aggregations do not provide a useful service to devise policies for achieving a lasting peace, whichever calculation one prefers.

The issue of aggregations scratches on the surface of more fundamental problems: the database itself. As the authors already note, “DWI analysis can use any data source (media reports, epidemiological surveys, coroners’ reports) *as long as the data are adequately valid, accurate, and comprehensive*” (p4, emphasis added), and even those qualifications are rather stretchable notions; e.g., who decides what, and when, the data collection is sufficiently comprehensive, and what may be sufficiently comprehensive for one purpose may not be for another. The front-end of a database may have an easy form to be filled in by the data entry clerk, but that hides what data can be stored, how it is stored, and how feasible it is to analyse the data through querying the database. Although one should not assume a database designer intends to impose his or her political preference though the design of the database, it is, in fact, very easy to do—also unintentionally. For instance, one can aggregate creatively (e.g., lumping together all Sunni and Shia organisations in Lebanon, be it during data entry or as “permissible access” for querying the data) or for specific sections report much more detail (e.g., the function of the building bombed, such as “water purification plant”, “electricity plant”, “medicine factory” and so forth, instead of just “industrial facility”). A notorious aggregation is that of “civilian”, in particular because there is no uniformly agreed-upon definition of civilian and where the fine line lies, if any, with, to name a few categories and terms: combatants, soldiers, paramilitaries, ex-army officials, police, and those who provide the logistics for the operations.

Related to aggregation, but more subtle, is under- or over-representing geographical information: one can play with granularity (by city, region, country, continent), with the categorization criteria of administrative or otherwise coherent regions (e.g., state borders, language, parish, “ethnicity”), and include or exclude notions such as “occupied territory” that are related to the actors, or that a particular area may, or may not, be categorized or named differently at the same time and be stored as such in the database. In addition, conflicts change through, among other factors: changing alliances, new types of arms used, regions affected, and types of violent acts. It then depends on how the database is implemented if it can adequately deal with such changes, be it changing values, changing attributes (e.g., adding a finer-grained categorization of arms), or new relations among the entities (later on deciding one also wishes to record, say, the coalitions of actors throughout the conflict) and how to record such temporal aspects of the structure of the database and its contents, if at all.

Further, one can facilitate pollution of the data in the database by using no or very limited database integrity constraints, which leaves the database inconsistent. For instance, one can enforce that for each occurrence of rape added to the database, there must be a corresponding police record; or one can precisely omit this constraint to obtain an estimate of reported incidences versus officially recorded incidences in the crime statistics, but then a query for “retrieve the amount of rape cases” most likely will not yield the same total number. Such issues and pitfalls as outlined here are, from an ICT perspective, first and foremost dealt with during the *requirements analysis* stage (informally: what is the purpose of the database, what is it supposed to do) and the *conceptual modelling* stage, where a diagram is made of what kind of data should go in the database. Going through these stages will not solve all problems of “war databases”, but can certainly ameliorate it. Moreover, if a given database has a clearly stated purpose *x*, then using it for purpose *y* should at least alert the data analyst to be careful before drawing conclusions.

With these database issues in mind, let us now return to Hicks and Spagat’s other main illustration: the protracted conflict in Northern Ireland. The source data are taken from the fine-grained CAIN web service (2009) and is used to calculate two complementary DWIs (table 4 and related text, p5): “aggressive acts (killing civilians) and endangerment to civilians (by not wearing uniforms)”. [\[2\]](#) The “British Security Forces” (BSF) have a “Civilian

mortality DWI” of 52, the “Irish Republican Paramilitaries” (IRP) 36, and the “Loyalist paramilitaries” (LP) 86. One first has to observe the chosen naming and aggregations; e.g., does this title refer to the IRA, or are they lumping the IRA together with the Real-IRA and Continuity-IRA, and likewise for all UFF, LVF, and so forth? Consulting the CAIN web service to ascertain what is going on, it, interestingly, lists 29 groups; that is, the details are there, but the article authors preferred this particular aggregation for the DWI calculation. In addition, Hicks and Spagat’s “number of civilian + civilian political activist” total to, respectively for each aggregate, $190+738+873=1801$, but the CAIN source data has a total of $1797\text{ civilian}+58\text{ civilian political activists}=1855$, and then a series of statuses such as “ex-British army”, “ex-IRA” and so forth, who, while being “ex-” are not real civilians according to the CAIN database; either way, the two totals should be the same, but they are not. The other one, “Attacks without uniform DWI” are, according to the table, “approaches 0” (BSF), “approaches 100” (IRP) and “approaches 100” (LP) without actual values to do the calculation with; despite the vagaries, for the IRP the authors prefer the adjective “extremely high rate” but for the LP it is only “very high rate” even though the given values are the same. They offer a comparatively long explanation for the nastiness of the IRP, but it is obvious that the BSF and LP have the dirtiest civilian DWI and that LP killed most civilians, no matter how one wants to explain it away and dress it up with DWIs (maybe not so coincidentally, the authors are affiliated with UK institutions).

I will leave Hicks and Spagat’s “female mortality DWI” of the Arab-Israeli conflict and the “child casualty DWI” of Chechnya for the interested reader to analyse (including the term “unexploded ordnance” that injured or killed children—by exploding). More generally than the examples of “multiple-interpretation” in favour of one actor or another, is that a DWI perhaps will do more harm to conflicts than contributing to resolution precisely due to its oversimplification and blind (uninformed?) reliance on the underlying databases it uses; misuse by politicians and headline-media is all too easy to imagine.

However, an easy way out to the problems with data collection, database design, and subsequent data analysis does not exist. Partial solutions are proper requirements analysis, database development, and, above all, *transparency* of the data analysis that should not be ripped from its context and simplified into a uni-dimensional index. Another partial solution may be to develop a generic categorization of types of victims, types of violent acts, types of arms, etc. and how they relate, so that each empirically-oriented peace and conflict researcher can take part or all of such an openly available application-independent categorization (also called a *domain ontology*) and use that for developing the database. In the meantime, one is strongly advised to check the source data carefully before calculating or propagating any DWI.

Allow me to make a final remark in that my comment “stated purpose x, then using it for purpose y should at least alert the data analyst” might be applied to this review, because the DWI was presented as a “public health tool” whereas I cast it in terms of peace & conflict research and ICT. Public health facilities are a scarce good in conflict zones and allocating limited resources based on DWI values may well prolong the conflict, depending on how the DWI value is constructed and who is therewith perceived to be the most virtuous or less vicious or most in need of health care and deserve to be served first. Health care policies are not necessarily neutral and should not be seen in isolation; hence, despite – or, perhaps, thanks to – its flaws, Hicks and Spagat’s paper can be useful in education to sensitize students to this topic and for researchers to delve deeper into the influence of health care as well as better addressing the problematic of dealing with raw data and quantifying and synthesising it.

References

CAIN Web Service (2009). <http://cain.ulst.ac.uk/sutton/index.html> (last accessed on 4 June 2009);

http://cain.ulst.ac.uk/sutton/tables/Organisation_Responsible.html (4 June 2009)

<http://cain.ulst.ac.uk/sutton/tables/Status.html> (4 June 2009)

CERAC (2009). http://www.cerac.org.co/home_english.htm (last accessed on 4 June 2009).

Hicks, M.H.-R. & Spagat, M. (2008). The Dirty War Index: A public health and human rights tool for examining and monitoring armed conflict outcomes. *PLoS Medicine*, 5(12), e243. <http://dx.doi.org/10.1371/journal.pmed.0050243>

Sondorp, E. (2008). A new tool for measuring the brutality of war. *PLoS Medicine*, 5(12), e249.
<http://dx.doi:10.1371/journal.pmed.0050249>

Taback, N. (2008). The Dirty War Index: Statistical issues, feasibility, and interpretation. *PLoS Medicine*, 5(12), e248.
<http://dx.doi.org/10.1371/journal.pmed.0050248>

Weinstein, J.M. (2007). *Inside rebellion—the politics of insurgent violence*. Cambridge: Cambridge University Press.

Footnotes

[1] Both Hicks and Spagat are research associates of CERAC; the database is available after registration, which has substantially less types of information and less explanation than the CAIN Web Service database.

[2] As mentioned earlier, the latter DWI is theoretically problematic, because the distinction between actors who use violence and their supporters in the population is often not that clear, and off-duty soldiers are not necessarily automatically civilians; but the argument is long and I shall not digress further on the topic here.

About the Author

Dr. C. Maria Keet is an Assistant Professor (ricercatore a tempo determinato) at the KRDB Research Center, Faculty of Computer Science, Free University of Bozen-Bolzano, Italy. She holds a PhD in Computer Science (2008), a BSc(hons) 1st class in Computing (2004), an MA in Peace & Development Studies (2003), and an MSc in Microbiology (1998). She has worked in industry for 3.5 years, volunteered in several capacities in various organisations and is currently chairperson and webmaster of the NGO Professors without Borders.

A Publication of:



University for Peace

