

What Aggregate Civilian-Combatant Ratios Tell Us, And What They Don't: A Case Study from the Gaza Conflict

lieber.westpoint.edu/what-aggregate-civilian-combatant-ratios-tell-us-what-they-dont-case-study-gaza-conflict

March 20, 2026

by [Oded Hen](#) | Mar 20, 2026



Public debate about contemporary armed conflict increasingly relies on aggregate civilian-to-combatant casualty ratios as indicators of legality and moral responsibility. Some use these ratios to argue that a military campaign is either exceptionally restrained or exceptionally destructive. However, international humanitarian law (IHL) requires assessing proportionality and distinction not through aggregate outcomes across an entire conflict, but at the level of specific attacks, based on the information reasonably available to decision-makers at the time they decide to use lethal force.

A recent analysis published by Action on Armed Violence (AOAV) titled “[Why Israeli Claims of Low Civilian-to-Combatant Harm in Gaza Do Not Hold Up](#)” offers a detailed statistical framework for estimating civilian-to-combatant ratios and interpreting their implications. The article engages a current and important public debate and reflects a serious effort to measure harm under conditions of limited access and contested reporting. However, the analysis raises a broader methodological and legal question that extends well beyond the Gaza context: can aggregate civilian-combatant ratios, particularly those derived from statistical modelling, determine whether a belligerent has violated IHL?

In conflicts lacking unrestricted battlefield access and transparency in combatant status recordkeeping, aggregate casualty estimates necessarily rely on layered inferential assumptions. These may include demographic classification rules, correction factors for undercounting, and modelling of indirect mortality. Each additional inferential step introduces further uncertainty, even when the resulting estimates appear internally coherent. Under such conditions, statistical modelling does not resolve evidentiary limitations but redistributes them across assumptions. The central question is therefore not only what aggregate ratios show, but how much confidence can reasonably be assigned to conclusions derived from modelling in the absence of verifiable incident-level classification.

This post argues that modelling-based conclusions about legality are unsustainable under such conditions, and that any legal determination grounded primarily in modelling rather than verifiable incident-level evidence would be insufficient. Aggregate casualty ratios may illuminate patterns, identify areas of concern, or justify further investigation. But, in the absence of incident-level analysis, they cannot establish proportionality, demonstrate unlawful targeting, or determine whether a belligerent has met its IHL obligations. The AOAV analysis, which I use as a case study to demonstrate this argument, clearly illustrates the limits of aggregate casualty modelling in assessing the legality of military conduct in modern urban warfare and how it may lead to false conclusions.

Proportionality under International Humanitarian Law (IHL)

IHL regulates the conduct of hostilities through the principles of distinction, proportionality, and feasible precautions. Codified in Additional Protocol I (AP I) to the Geneva Conventions (GC), proportionality obliges belligerents to refrain from attacks likely to cause civilian losses “excessive in relation to the concrete and direct military advantage anticipated” ([AP I, art. 51](#)). The rule is inherently case-specific and forward-looking, with an attacking force determining the lawfulness for each individual attack, based on the circumstances known or reasonably knowable to the attacking force at the time.

IHL, in both treaty and customary form, does not establish a conflict-wide numerical threshold that signals whether something is legal or illegal. Some State manuals and commentary indicate that precautionary measures are implemented at higher levels of operational command, and that the anticipated “military advantage” may, in certain circumstances, be assessed in relation to an attack considered as part of a broader military operation rather than confined to the immediate tactical gain of a single strike ([International Committee of the Red Cross 2024, Rule 14](#)). These interpretations, however, do not alter the structure of the rule. A military campaign is not proportionate or disproportionate solely by reference to an overall ratio of civilian-to-combatant deaths. Significant civilian losses may still be lawful if justified by critical military gains, whereas even minor losses may be unlawful if linked to trivial objectives.

At the same time, aggregate outcomes reflect a wide range of variables unrelated to targeting decisions, including force structure, operational environment, civilian movement patterns, adversary tactics, and reporting practices. For this reason, legal evaluation requires examination of specific engagements rather than statistical summaries of conflict-level harm. This doctrinal structure places a clear limit on what aggregate casualty ratios can establish. Even highly precise calculations do not capture what proportionality legally assesses: the relationship between anticipated military advantage and expected incidental harm in particular attacks.

The AOA Modelling Framework and Its Inferential Structure

The AOA analysis seeks to estimate the scale and structure of civilian harm by modelling casualty patterns, demographic distributions, and potential undercounting to demonstrate that the Israeli prime minister's claim about the ratio is false. However, the article replaces one contested estimate with another derived from a modelling framework that rests on a series of substantive assumptions. This effort reflects a legitimate concern: in conflicts with restricted access to the battlefield and obscured combatant status metrics, casualty classification is inherently uncertain, and statistical modelling may therefore appear to provide a way to reconstruct patterns that are otherwise imperceptible through direct observation.

Such modelling may indeed provide descriptive insights by highlighting demographic trends, suggesting possible reporting gaps, and generating hypotheses about patterns of harm, but it cannot support the conclusion "that Israel's conduct in Gaza involves widespread and wholly disproportionate civilian loss ... placing the campaign among the most civilian-destructive of recent decades."

Within this broader modelling approach, the civilian-to-combatant ratio emerges through a series of inferential stages. First, modelers categorize individuals using demographic indicators, often treating women, children, and the elderly as presumptive civilians, and inferring men's civilian status through demographic balancing. These estimates may undergo adjustment later using historical correction factors that purportedly account for undercounting or misclassification in prior conflicts. Finally, modelers may incorporate indirect mortality predictions to produce an expanded estimate of total harm. Each stage depends on assumptions that are themselves contestable, and uncertainty at earlier stages carries through the model as a whole. The resulting ratio is therefore not a direct observation but the cumulative product of multiple modelling decisions.

At the same time, although the model may serve as a humanitarian indicator by highlighting possible suffering in a conflict zone, the central difficulty lies not in measuring harm, but in the legal and normative conclusions drawn from aggregate modelling. The relevant question is not whether statistical modelling can produce internally coherent estimates, but whether those estimates can establish that a belligerent violated IHL.

Answering that question requires attention to the structure of the data on which aggregate ratios depend. In the specific case of Gaza, scholars and journalists who regard the Gaza Health Ministry as a broadly credible source of casualty data emphasize that the Ministry does not record affiliation or participation in hostilities, and that wartime conditions inevitably produce gaps, duplications, and reporting delays ([Al-Mughrabi & Farge, 2025](#); [Fox, 2024](#); [NY Post, 2025](#); [Spencer, 2025](#)).

The Claims Advanced by the AOA Analysis

The AOA analysis advances several distinct claims that together support a broader conclusion. It argues, first, that official claims regarding relatively low civilian-to-combatant ratios are unfounded; second, that the available evidence consistently points to substantially higher ratios; and third, that demographic adjustments, undercount corrections, and indirect mortality estimates are used to produce an expanded estimate of civilian harm. It further maintains that available casualty lists are sufficiently reliable to support ratio-based inference, and that restricted access to the battlefield makes statistical modelling not only useful, but necessary.

To be sure, each of these claims is debatable. The central analytical issue, however, is that none of them, individually or cumulatively, can establish whether a belligerent satisfied or violated its proportionality obligations. Even internally coherent modelling relies on layered inference, including demographic classification, correction factors, assumptions about participation in hostilities, and estimates of indirect mortality.

When modelers combine multiple inferential steps, uncertainty tends to increase rather than diminish. Taken together, these claims operate at analytically distinct levels: the reliability of casualty measurement; the inferential assumptions used to classify individuals; and the normative or legal interpretation of the resulting ratios. Treating these levels as equivalent risks obscures the difference between uncertainty in empirical classification and conclusions about legal responsibility.

The Problem of Combatant Classification

Any civilian-to-combatant ratio depends fundamentally on a belligerent's means of classifying individuals. In contemporary urban warfare, especially in conflicts involving non-State violent actors, such classification is often highly uncertain and, in many cases, practically impossible. Fighters operate without uniforms, and civilians may intermittently participate in hostilities, blurring the line between civilian and combatant status. Others perform supporting or logistical roles that enable violence, including by providing operational assistance or knowingly remaining within military assets to deter potential attack. These challenges raise a more basic question: which people count as civilians, and what role do they play in the conflict? Is a person who shelters a fighter still immune from attack? What

about an individual who stores ammunition in his home or secretly holds hostages? And if such a person is targeted, how should he be classified after death, as a civilian or as a combatant?

In the absence of reliable incident-level identification, analysts often rely on demographic inference, statistical modelling, or other indirect indicators. For example, demographic patterns such as skewed sex or age distributions among casualties, as reported by AOA, can arise from multiple mechanisms unrelated to targeting decisions, including differential exposure to combat zones, mobility patterns, support function performance, or varying levels of participation in hostilities. These methods, therefore, generate probabilistic classifications rather than determinations of legal status. They can produce numerical estimates, but they cannot confidently determine whether particular individuals were civilians or combatants. Aggregate ratios are also highly sensitive to marginal reclassification, so even minor adjustments in classification assumptions, such as how irregular fighters, support personnel, or civilians engaged in hostilities are categorized, can substantially alter overall ratios and produce large shifts in the resulting estimates without any change in the underlying events.

These limitations have an additional implication. In the absence of transparent, incident-level classification, disputes over ratios typically involve competing inferential frameworks rather than demonstrable, concrete evidence. Replacing one model-dependent estimate with another does not conclusively establish that an alternative ratio is incorrect; it simply produces a different estimate shaped by different assumptions.

Indirect Deaths and Problems of Attribution

AOA's analysis claims that indirect deaths from starvation and disease are "effectively 100% civilians," almost entirely children and the elderly. This claim, however, rests on modelling assumptions and extends beyond what such methods can reliably establish. Aggregate casualty assessments often incorporate indirect deaths arising from displacement, infrastructure collapse, disease, or deprivation. The moral gravity of these harms is not in dispute, yet they do not affect civilians exclusively. Indirect mortality may also include combatants who die from injury, deprivation, or disease.

Additionally, including indirect mortality within civilian-to-combatant ratios risks conflating two analytically distinct phenomena: the humanitarian consequences of war; and the legal assessment of targeting decisions in specific attacks. Indirect deaths may result from infrastructure collapse, displacement dynamics, governance breakdown, disease transmission, or failures in aid distribution, all factors that do not correspond directly to proportionality assessments governing uses of force. Even when indirect mortality is substantial, assigning it overwhelmingly to one legal category, such as civilians, introduces a level of precision that mortality modelling cannot sustain and risks misrepresenting what the ratio is intended to measure.

The Operational Reality of Urban Warfare

Modern urban warfare creates conditions in which civilians and military objectives are inseparable, structurally increasing the risk of incidental harm. In Gaza, military assets and personnel have been embedded within dense civilian environments, including through the use of voluntary and involuntary human shields ([Israel Defense Force, 2026](#)). Where civilians are deliberately co-located with military objectives, the likelihood of civilian harm rises even when belligerents attack lawful targets and take precautionary measures.

This operational reality is largely absent from AOA's ratio-based analysis. Aggregate civilian-to-combatant figures cannot capture the conditions under which such outcomes arise, including the forced or voluntary presence of civilians as shields near military infrastructure, constraints on evacuation, or the legal responsibility of the party that employs human shields. Elevated civilian harm in dense urban combat, therefore, cannot, by itself, establish unlawful targeting or disproportionality. Legal assessment remains tied to specific attacks and their operational context, and it cannot be resolved through aggregate casualty ratios alone.

Comparisons, Legal Responsibility, and the Limits of Ratio-Based Inference

Some frame ratio-based analysis in terms of implicit or explicit normative thresholds that suggest that particular civilian-to-combatant ratios themselves reflect a grave disregard for civilian life. Yet neither IHL nor the Rome Statute of the International Criminal Court (ICC) recognizes any numerical ratio as establishing illegality, culpability, or moral violation. Aggregate figures may warrant scrutiny or prompt further investigation, but they do not function as legal or ethical benchmarks in their own right.

AOA's analysis seeks to strengthen its conclusions by comparing Gaza's estimated civilian-to-combatant ratio with conflicts associated with war crimes and genocide, and by referring to ongoing proceedings before international courts. Such comparisons, however, rest on a misunderstanding of international criminal culpability. Genocide and other core international crimes are not defined by casualty ratios or by comparative levels of civilian harm, but by specific legal elements, most importantly among them, intent, policy, and the character of particular acts.

In the case of genocide, the defining requirement is the intent to destroy, in whole or in part, a protected group as defined in Article II of the [Convention on the Prevention and Punishment of the Crime of Genocide \(1948\)](#) and in Article 6 of the [Rome Statute of the International Criminal Court \(1998\)](#). Neither the number of civilian deaths nor comparisons with other conflicts can establish or substitute for that element. A high civilian-to-combatant ratio, therefore, does not constitute evidence of genocide or of crimes legally analogous to it.

Conclusions

Aggregate civilian–combatant ratios can serve important analytical and humanitarian purposes. They may reveal patterns of harm, identify areas that warrant closer scrutiny, and bring large-scale suffering to public attention. In contexts of armed conflict, where access is limited and information is incomplete, such measures can help frame investigative priorities and contribute to a broader understanding of conflict dynamics. When used in this way, they are legitimate and valuable descriptive tools.

However, their evidentiary limits are clear. Aggregate ratios cannot establish proportionality, demonstrate unlawful targeting, or determine whether IHL obligations were satisfied. These legal judgments require incident-level analysis of specific attacks. Conflict-wide numerical outcomes, however carefully constructed, do not capture the elements required for legal assessment at the level of particular attacks. This limitation holds whether one casts civilian-to-combatant ratios as unusually low or unusually high. In both cases, aggregate statistical patterns purport to stand in for legal evaluation. The central problem is therefore not one of measurement accuracy, but of analytical mismatch: aggregate casualty metrics may describe patterns of harm across an entire conflict, whereas IHL governs the legality of decisions made in specific operational circumstances.

The AOAV analysis highlights this mismatch very clearly. It is a modelling framework built on layered inferential assumptions about classification, undercounting, and indirect mortality underlying conflict-wide estimates, which supposedly serve as a basis for far-reaching legal and moral conclusions. Increasing the sophistication of modelling does not resolve this difficulty, because the central issue is conceptual rather than technical. When classification is uncertain, and estimates depend on cumulative assumptions, model-generated ratios can support, at most, provisional humanitarian concern and bounded inference. They cannot sustain definitive claims about disproportionality, legal violation, or equivalence with legally defined atrocities. Assigning them that role misconstrues both the evidentiary capacity of the data and the requirements of the legal framework. In this sense, the AOAV analysis is not merely unconvincing but seriously misguided, as it attributes decisive normative weight to a form of evidence that cannot support it.

Aggregate casualty ratios cannot determine legal judgment. Recognizing this limit does not diminish the importance of highlighting large-scale suffering; rather, it clarifies the boundary between describing harm and assessing legality and prevents statistical patterns from shouldering a normative weight they cannot bear.

[Oded Hen](#) is a PhD candidate in the Department of Philosophy at Bar-Ilan University. His research focuses on moral philosophy, particularly the ethics of modern warfare and just war theory.

The views expressed are those of the author, and do not necessarily reflect the official position of the United States Military Academy, Department of the Army, or Department of Defense.

Articles of War is a forum for professionals to share opinions and cultivate ideas. Articles of War does not screen articles to fit a particular editorial agenda, nor endorse or advocate material that is published. Authorship does not indicate affiliation with Articles of War, the Lieber Institute, or the United States Military Academy West Point.

Photo credit: Jaber Jihad Badwan

[SUBSCRIBE](#)