



The Indian Journal for Research in Law and Management

Open Access Law Journal – Copyright © 2026

Editor-in-Chief – Dr. Muktai Deb Chavan; Publisher – Alden Vas; ISSN: 2583-9896

This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

Presented Compliance, Absent Accountability: Article 36 Weapons

Review as an Accountability Sink

Srijan Kumar Sharma

ABSTRACT

This paper asks whether Article 36 of Additional Protocol I functions as a genuine accountability safeguard for weapons review, or merely produces its appearance. Using doctrinal reconstruction of Article 36's architecture (via the ICRC's Guide and Commentary) combined with an interdisciplinary framework borrowed from Dan Davies' concept of the "accountability sink," the paper maps three structural features; delegation to process, broken feedback loops, and information-attenuating certification, onto State practice, supported by evidence of near-universal non-disclosure among States Parties. Autonomous weapon systems (AWS), illustrated through the 2020 Kargu-2 incident and analysed against the ICRC's "human control" standard and the Bemba command-responsibility jurisprudence, serve as the paper's aggravating case rather than its cause. The central argument departs from the dominant "accountability gap" literature: rather than diagnosing an absence awaiting correction through harmonisation or added process, the paper contends Article 36 is a functioning mechanism that converts the demand for accountability into its own discharge, such that AWS have not broken the mechanism but merely exposed its pre-existing design.

INTRODUCTION

In March 2020, during the closing stages of the Libyan civil conflict, a Kargu-2 loitering munition was reported to have engaged retreating Haftar-affiliated forces without a human

operator in the loop at the moment of attack.¹ The UN Panel of Experts' account of the incident, cautious and heavily hedged as it was, nonetheless did something unusual wherein it forced open a question that international humanitarian law had spent over four decades assuming was already closed, under which somewhere between the manufacture of that weapon and its deployment on the battlefield, a State was legally required to have determined; as per Article 36 of Additional Protocol I;² whether its employment would, in some or all circumstances, violate international law. So far, no public records exist of the determination of this question in this specific instance. This is not a gap in the available evidence but instead is the ordinary condition of Article 36 practice.

This very condition that moves beyond the mere textual framing of the said article is the central question of this study. The dominant scholarly response to episodes like Kargu-2 has been to diagnose an "accountability gap",³ a space in which autonomous weapon systems (AWS) diffuse legal responsibility so thoroughly across programmers, manufacturers, commanders and operators that no individual can be shown to possess the knowledge or intent international criminal law requires. The "gap" diagnosis is not wrong in itself but provides an incomplete picture, and this very incomplete picture has significant practical consequences such that a gap implies an absence as something IHL forgot to build, and could in principle still build, through better weapons review procedures, harmonised state practice, or clearer rules on human control. However, this work argues that Article 36 is not an absence but a mechanism, and that the mechanism does not fail to produce accountability so much as it actively produces its appearance while dissipating its substance. Borrowing a term recently developed outside legal scholarship; Dan Davies' account of the "accountability sink",⁴ a structure that absorbs institutional pressure to explain a failure without any actor within it

¹ the UN Panel of Experts on Libya report, S/2021/229 (2021).

² Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts (Protocol I), art. 36, June 8, 1977, 1125 U.N.T.S. 3

³ Tim McFarland & Tim McCormack, *Mind the Gap: Can Developers of Autonomous Weapons Systems Be Liable for War Crimes?*, 90 Int'l L. Stud. 361 (2014); Thompson Chengeta's work; Verdiesen et al.)

⁴ *The Unaccountability Machine: Why Big Systems Make Terrible Decisions and How the World Lost Its Mind* (2024)

being able, or required, to supply that explanation; this paper contends that Article 36 weapons review functions, in State practice as it actually exists, as precisely this kind of sink wherein, It presents mere compliance but does not produce accountability.

Autonomous weapon systems supply the clearest illustration of this argument, and they run through this paper as its central case, but they are not its cause. This work reconstructs the doctrinal architecture of Article 36 to show how much of what passes for a binding review obligation is, on the International Committee of the Red Cross's own account, a matter left to State discretion; it further imports the accountability-sink framework and maps its structural features onto that architecture. We shall also turn to AWS specifically, taking note of its prevalence in modern-day conflicts, and reading the "human control" standard developed by the ICRC and others not as a safeguard under strain but as the sink's clearest operating mechanism. The work also attempts to argue against the available evidence of State practice; especially, the near-total absence of published reviews; and further argues that this opacity is functional rather than incidental; while considering what follows for reform proposals in case the diagnosis offered here is correct.

THE THEORETICAL UNDERTAKING OF ARTICLE 36

Article 36 of Additional Protocol I is often described, including by the ICRC itself, as one of the few genuinely preventive rules in international humanitarian law; a provision that operates before the battlefield rather than on it. According to it, A High Contracting Party is under an obligation, "in the study, development, acquisition or adoption of a new weapon, means or method of warfare," to determine whether its employment "would, in some or all circumstances, be prohibited"⁵ by the Protocol or by any other applicable rule of international law. When read *prima facie*, this looks like a demanding standard wherein a State must certify not merely that a weapon is lawful in its intended use, but that it does not fall foul of IHL in any circumstance of foreseeable employment; but when understood against the

⁵ Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts (Protocol I), art. 36, June 8, 1977, 1125 U.N.T.S. 3.

ICRC's own interpretive guidance, the obligation is considerably thinner than it first appears, and it is this gap between the apparent stringency of the standard and the actual thinness of what it requires that supplies the raw material for the argument developed in this work.

Two features of the material scope are worth isolating before turning to the interpretive point; First, the review attaches to "means and methods" together, that is, the weapon cannot be assessed apart from the way it is expected to be used, so that a weapon lawful in one mode of employment may fail review in another. Second, and more consequentially, the ICRC's own Commentary confines the obligation to a weapon's "normal or expected use" meaning, meaning a State need not "foresee or analyse all possible misuses",⁶ since almost any weapon can be misused in a manner IHL would prohibit. This is a sensible limiting principle; an unlimited obligation to anticipate every conceivable misuse would be unadministrable; but it also means that the "in some or all circumstances" language, which reads as maximally demanding on its face, is in practice narrowed by the reviewing State's own account of what "normal" use will look like. This gives us a general understanding that the reviewing authority, in most cases the State authorities, shall determine the nature, extent and standards of this review.

The substantive legal framework the review is meant to apply, including the specific treaty and customary prohibitions, the general rules on indiscriminate effects and superfluous injury, and, residually, the Martens Clause⁷ where no specific rule exists, is comparatively well settled and is not where this paper locates its argument. The more revealing material lies in Part 2 of the ICRC's own Guide to the Legal Review of New Weapons,⁸ which turns from what the review must assess to how it must be conducted, and here the pattern changes sharply; where it is observed that article 36 does not specify by what instrument a review mechanism should be established, whether it should sit within the ministry of defence or

⁶ Yves Sandoz, Christophe Swinarski & Bruno Zimmermann eds., *Commentary on the Additional Protocols of 8 June 1977 to the Geneva Conventions of 12 August 1949* ¶ 1469 (ICRC 1987).

⁷ Protocol I, *supra* note [2], art. 1(2).

⁸ ICRC, *A Guide to the Legal Review of New Weapons, Means and Methods of Warfare: Measures to Implement Article 36 of Additional Protocol I of 1977* (2006).

elsewhere, whether it should be a committee or a single official, what expertise must be represented, at what precise point in procurement it must occur beyond "the earliest possible stage," how a reviewing body should reach its decision, whether that decision binds anyone, whether conditions may be attached to an approval, whether an adverse finding may be appealed, or whether any record of the review must ever be kept or disclosed. On each of these questions the Guide's method is identical as it surveys the six or seven States that have made their procedures public, notes that their practices diverge, and declines to resolve the divergence into a rule. Under this, the obligation to verify is treaty-based and thereby bound by the general rules of international law and law of armed conflict; however, the parameters to review are left at the discretion of the State itself.

This is not a criticism of the Guide, which accurately reports the state of the law; it is an observation about what the law, accurately reported, turns out to permit that is a State may lawfully discharge its Article 36 obligation through a single official, internal to the ministry procuring the weapon, applying a non-binding standard, subject to no appeal, with no requirement that a record be kept and no obligation that any conclusion; whether favourable or adverse; ever be disclosed to the public, to the ICRC, or to other States Parties beyond a general commitment under Article 84⁹ to share "procedures" rather than outcomes. Every element of what a legal system would ordinarily treat as the minimum architecture of accountability; including, but not limited to, an external or at least independent decision-maker, a binding determination, a reviewable record, a consequence for an adverse finding; is present in the Guide only as an option some States have chosen, never as a condition Article 36 requires. The obligation, in other words, mandates that a process occur; it does not mandate that the process be capable of being checked.

It is this specific feature; an internationally binding obligation to review, coupled with an entirely permissive architecture for how review is institutionally realised; that the remainder of this paper argues is not an oversight awaiting correction, but the precondition for Article

⁹ Protocol I, *supra* note [2], art. 84

36 to function as an accountability sink. Before that argument can be made, however, the concept itself needs to be set out on its own terms; that shall be done ahead.

IMPORTING THE ACCOUNTABILITY SINK

The term "accountability sink" is not a legal concept, and it should not be dressed up as one; it comes from Dan Davies' 2024 study of institutional failure, which asks a narrower and more useful question than the one lawyers usually ask about responsibility of, not "who is to blame," but "where does the demand for an explanation go when it enters a system, and does the system have any component capable of producing one." Davies' claim is that many modern institutions; corporate, regulatory, administrative; are built from long chains of delegated decision-making in which each actor's authority is bounded by rules set elsewhere in the chain, so that when an outcome goes wrong, every individual within the system can truthfully say they followed the process that applied to them. The explanation the outside world wants; why did this happen, and who will answer for it; has a specific basis of determination, and it is not refused or absorbed, but the system has sunk it.

Three features of this structure matter for present purposes, and each has a direct legal analogue. Firstly, an accountability sink depends on delegation to process rather than to a person; that is, authority is exercised through following a procedure correctly, not through a judgment for which the actor remains answerable. Secondly, the sink requires a broken feedback loop; the actor who designs or approves the process is insulated from the consequences experienced by whoever the process fails, so the process is never tested against its own failures in a way that forces revision. Thirdly, and most relevant to Article 36 specifically, certification functions within the sink as what Davies calls an information-attenuating technology, that is, a stamp of approval that is cheaper to produce and easier to point to than the underlying substance it supposedly certifies, and which; once issued; ends to be treated by everyone downstream as a satisfactory answer to the question "has this been checked," foreclosing further inquiry rather than inviting it.

The Second section has already shown why Article 36 maps onto this structure with unusual precision, and it is worth being explicit about why the fit is not merely rhetorical. Delegation to process; the obligation is to conduct a review, not to reach a correct or independently verifiable answer, and the reviewing State defines for itself what "normal or expected use" means for the purpose of the test it must pass. Broken feedback loop; because review outcomes are, on the ICRC's own account, not required to be published, kept as an accessible record, or subjected to appeal in the great majority of State practice, there is no mechanism by which a weapon's real-world performance; thereby its failures, its unintended effects, the civilian harm it in fact causes; is fed back into the review that approved it, still less into the reviews of other States. Information-attenuation through certification; a completed Article 36 review, wherever one is known to exist, functions internationally exactly as Davies predicts a certification will; as the fact that forecloses the question. A State that can say a weapon "underwent legal review" has, in the practice of international bodies and public debate alike, generally been treated as having discharged the inquiry, regardless of who conducted the review, what standard they applied, or whether the conclusion was ever disclosed.

It is worth pondering what this reframing changes and what it does not; we understand that it does not suggest that States conducting Article 36 reviews are acting in bad faith, or that no State review is ever rigorous; Davies' own examples are drawn from institutions with no shortage of individually diligent, well-intentioned participants. The sink is a property of the structure, not an accusation against the people operating within it; a conscientious official working inside a review mechanism with no binding authority, no public record, and no external check is still, whatever their diligence, working inside a sink. Nor does it suggest that the "accountability gap" literature is wrong about its symptoms; the diffusion of criminal responsibility across programmers, commanders and operators that the further part takes up is a real and correctly identified problem. What the reframing changes is the diagnosis of cause, and therefore what would count as a cure. A gap is filled by adding what is missing. A sink is not fixed by adding more process to it.

AUTONOMOUS WEAPON SYSTEMS AS THE AGGRAVATING CASE

Autonomous weapon systems did not create the accountability sink described in Part III; what they do is remove the specific feature of conventional weapons review that had, until now, made the sink's operation easy to overlook, wherein the presence, at the point of attack, of a human being who could always be pointed to as the person who last exercised judgment, however constrained. Davison's account of the "human control"¹⁰ standard; developed by the ICRC and now the dominant vocabulary in CCW discussions;¹¹ is best read not as a proposed solution to this problem, but as an attempt to relocate the point at which a human's judgment can be said to have occurred, spreading it backward across the development, activation and operation stages precisely because it can no longer be located reliably at the moment of the attack itself. This relocation is doctrinally sensible and is also, on the sink analysis, exactly the mechanism by which responsibility is most effectively absorbed, because a standard that can be satisfied by any of several stages is a standard that no single actor at any one stage can be shown, on their own, to have failed to meet.

The clearest evidence of this is Davison's own treatment of individual criminal liability, which the ICRC's chapter states with unusual candour, wherein, A programmer who "intentionally programmes an autonomous weapon to operate in violation of IHL," or a commander who "activates a weapon that is incapable of functioning lawfully in that environment," would plainly be criminally liable; all from the same account. But the chapter immediately concedes that this is the easy case, and that the hard case; the one AWS actually generates; is different, under which programmers "might not have knowledge of the concrete situations" in which the weapon will later be deployed, and commanders, at the point of activation, "may not know the exact time and location where an attack would take place." The knowledge and intent that international criminal law requires is not absent from the

¹⁰ Neil Davison, *Autonomous Weapon Systems Under International Humanitarian Law*, in UNODA Occasional Papers No. 30

¹¹ Convention on Certain Conventional Weapons, Fifth Review Conference, Recommendations to the 2016 Review Conference Submitted by the Chairperson of the Informal Meeting of Experts, ¶ 2(b) (Apr. 2016), *reprinted in* Final Document of the Fifth Review Conference, U.N. Doc. CCW/CONF.V/10 (Dec. 23, 2016)

system; it is present, but distributed across actors, none of whom individually possesses enough of it to meet the mens rea threshold on their own. This is structurally close to the problem the ICC's Appeals Chamber confronted in the Bemba case¹² in the different context of command responsibility, where the tribunal split precisely over how much concrete knowledge of subordinates' conduct a commander must be shown to have possessed before "should have known" liability attaches. Bemba was decided within a chain of command a single individual formally headed, but AWS removes even that formal apex; the "chain" runs through a design process, a procurement decision, and a review mechanism with no binding decision-maker at all, so that the closest analogue to Bemba's position; someone with authority over the whole operation; does not clearly exist.

This is where Article 36 stops being an ordinary preventive safeguard operating upstream of a genuine problem, and becomes, on the sink analysis, part of the problem's architecture. Because the review that is supposed to catch an unlawful weapon before deployment is itself non-binding, internally conducted, and undisclosed in most State practice, it cannot supply the missing feedback loop that might otherwise compensate for the diffusion of criminal responsibility Davison describes. A binding, published, appealable review would not resolve the mens rea difficulty at the moment of attack, but it would at least create a single, identifiable decision; the approval of the weapon for use, which a court or a public could interrogate, and a named decision-maker who could be asked to justify it; however, Article 36 as it presently operates supplies neither. The predictability and reliability standards Davison sets out as the substantive core of a legally adequate review are, on paper, demanding; but a demanding standard applied by an unaccountable process only raises the quality of the certification's language, not the checkability of what it certifies. The result is a structure in which the more autonomous the weapon, the more work "human control" is asked to do, and the more that phrase is deployed at exactly the stage; development, procurement, review;

¹² *Prosecutor v. Bemba Gombo*, ICC-01/05-01/08-3636-Red, Judgment on the Appeal of Mr Jean-Pierre Bemba Gombo against Trial Chamber III's "Judgment Pursuant to Article 74 of the Statute" (June 8, 2018)

where the sink identified, already operates on conventional weapons, and where it operates with the least public visibility of all.

EVIDENCE: OPACITY AS FEATURE, NOT FAILURE

If Article 36 functions as an accountability sink, the prediction that follows is a specific one; the absence of public information about how States actually conduct their reviews should not be a correctable oversight, but a structural condition the mechanism depends on. The available evidence is consistent with this prediction, though it needs to be stated carefully, because the same fact, near-universal non-disclosure, is equally consistent with the more conventional "accountability gap" account, which reads it simply as a shortfall to be corrected through better transparency norms. The two accounts diverge not on the fact but on what should be expected to happen if the fact were corrected.

Of the roughly 175 States Parties to Additional Protocol I,¹³ fewer than twenty¹⁴ have ever made public any account of how they conduct Article 36 reviews, and of those, a smaller number still have published the substantive outcome of any individual review rather than a general description of their procedure. The 2006 ICRC Guide, canvassing the small set of States willing to share their instruments at all, records only descriptions of *mechanisms*; who sits on a committee, at what stage a request is triggered; and explicitly notes that whether a reviewing authority's decision is binding, and whether its records are accessible, varies by State and is frequently unresolved even where a mechanism's existence is acknowledged. Two decades of subsequent advocacy by the ICRC itself, by SIPRI, and by civil-society initiatives, have not materially changed this picture; the same handful of States are still cited

¹³ Int'l Comm. of the Red Cross, IHL Treaties Database: Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts (Protocol I), <https://ihl-databases.icrc.org/en/ihl-treaties/api-1977>

¹⁴ Maaïke Verbruggen & Vincent Boulanin, SIPRI Compendium on Article 36 Reviews, Stockholm Int'l Peace Rsch. Inst. (Dec. 2017), <https://doi.org/10.13140/RG.2.2.24199.70563>.

as the available examples, and no comprehensive account exists of what any of them concluded about any specific weapon.

On the gap account, this stasis is puzzling; two decades of consistent advocacy for a modest, low-cost transparency norm; publishing procedures, not necessarily outcomes; ought to have produced at least incremental convergence, particularly among States that otherwise present themselves as committed to IHL compliance and to the broader "responsible military AI" agenda, which does not seem to be the case. We observe that publication of a review's substance would do more than inform outside observers; it would supply exactly the feedback loop identified earlier as structurally absent; a record against which a later failure of the weapon in the field could be measured, and a named decision or decision-maker that a court, a treaty body, or public opinion could hold to that record. It is not obvious that any State, having built a review mechanism whose principal legal function is to satisfy Article 36 while incurring no binding consequence, would find it advantageous to introduce the one feature capable of generating consequence. Consistent with this, the closest customary-international-law scholarship¹⁵ comes to a general practice is a practice of *verbal* commitment under which States affirm, in CCW statements and elsewhere, that they conduct reviews, without those affirmations being accompanied by the conduct-level evidence; actual disclosed reviews; that would ordinarily be required to establish a customary rule from State practice rather than from *opinio juris* alone. A pattern of declared commitment substituting for demonstrated conduct is precisely what an information-attenuating certification, in Davies' sense, would be expected to produce; the statement that review occurs, then functions as the discharge of the inquiry, standing in for evidence that it occurred well.

None of this shows that no State review is ever competent, nor that opacity is the product of a coordinated intention to evade scrutiny. Herein, the claim is narrower and, for that reason, more durable as to, whatever the intentions of individual States are, the mechanism as

¹⁵ Natalia Jevglevskaja, *Weapons Review Obligation Under Customary International Law*, 94 Int'l L. Stud. 186 (2018).

designed does not require disclosure to function as a legally sufficient discharge of Article 36, and disclosure would introduce a cost; reviewability; that the mechanism's design gives no State a legal reason to bear.

IMPLICATIONS

Building upon this diagnosis, it is understood that it carries an uncomfortable implication for the reform proposals that dominate the existing literature; harmonised national procedures, model legislation, a published register of reviews, an agreed international standard¹⁶ for what "meaningful human control" requires. Each of these proposals treats Article 36 as a gap, that is something incompletely built, which more processes will complete. But a sink does not fail for want of process; on Davies' account, sinks are frequently constructed from process, and tend to absorb additional procedural requirements as readily as they absorb the original obligation, converting each new safeguard into a further stage at which delegation to process, rather than to an accountable judgment, can occur. A harmonised international standard for what human control must consist of, unaccompanied by a binding decision-maker, a disclosed record, and a consequence for an adverse finding, risks becoming not a check on the sink but an additional wall of it, that is a more sophisticated certification, attenuating information more effectively rather than less.

This paper does not attempt to specify what an adequate alternative institutional design would look like; that is a substantial question in its own right, and one this paper's diagnostic aims do not require it to resolve. A fix confined to the review process itself; more expertise on the committee, a clearer substantive test, more detailed guidance from the ICRC; treats the symptom the "accountability gap" literature correctly identifies without addressing the structural feature identified here: that nothing currently connects the outcome of a review to any consequence external to the process that produced it. Whatever institutional form

¹⁶ Boulanin & Verbruggen, *supra* note [14].

eventually closes that loop, it will need to do so from outside the reviewing State's own discretion, since the discretion is precisely what the sink is built from.

CONCLUSION

Article 36 has survived nearly five decades of IHL scholarship largely uncontested as a safeguard, its adequacy questioned only at the margins; whether it is customary law, whether it needs updating for autonomous weapons, whether States should publish more. This paper has argued that the more fundamental question has gone unasked; not whether Article 36 does its job well, but whether "doing its job," on its own terms, requires anything a legal system would ordinarily recognise as accountability at all. The concept of the accountability sink makes us understand that Article 36 is not a preventive rule awaiting completion; instead, it is a functioning mechanism that converts the demand for accountability into the appearance of its discharge, and autonomous weapon systems have not broken that mechanism; they have simply removed the last convenient place to pretend it was still working.