

The Interval, Not the Autonomy

Why Ten Years of Defining Autonomous Weapons Produced Nothing, and What to Regulate Instead: the Gap Between the Last Human Decision and the Effect

Christopher I. V. Farmer

Independent researcher, Geneva, Switzerland

17 August 2026 · Version 2, with a postscript dated 15 October 2026

Suggested citation. Christopher I. V. Farmer, ‘The Interval, Not the Autonomy’ (2026) preprint, version 2, DOI 10.5281/zenodo.21970777.

Status. Preprint; not yet peer reviewed. The body of this version is identical to the version of record of 17 August 2026; the Postscript records the report the Group of Governmental Experts adopted on 4 September 2026.

DOI. All versions [10.5281/zenodo.21970777](https://doi.org/10.5281/zenodo.21970777); version 1, 17 August 2026, [10.5281/zenodo.21970778](https://doi.org/10.5281/zenodo.21970778); this version, 15 October 2026, [10.5281/zenodo.23073917](https://doi.org/10.5281/zenodo.23073917).

Keywords. lethal autonomous weapons systems; Convention on Certain Conventional Weapons; Article 36 weapons review; meaningful human control; precautions in attack; international humanitarian law; military artificial intelligence

Licence. Creative Commons Attribution 4.0 International (CC BY 4.0).

Contents

The Interval, Not the Autonomy	2
Why Ten Years of Defining Autonomous Weapons Produced Nothing, and What to Regulate Instead: the Gap Between the Last Human Decision and the Effect	2
Abstract	2
Note on sources and method	4
Part I—Introduction	4
1. The problem stated	4
2. Why the definitional route was never going to arrive	4
3. Why the object is wrong	5
4. What is new here, and what is not	7
Part II—The Mandate as Cause	8
5. The mandate, read closely	8
6. Why this is a cause and not a description	9
7. The evidence, including a deadline that has already passed	10
7.1 What follows for a proposal	11
Part III—What the Rolling Text Actually Says	11
8. The text, box by box	11
9. The gap the text’s critics identify	13
10. What Parts II and III establish together	14
Part IV—The Review Obligation and the Wrong Axis	14
11. What Article 36 requires, and when	14
12. Where that fails, stated precisely	15
13. What the practice adds, and it is not encouraging	15
14. Why the answer is a second occasion rather than a wider scope	16
Part V—The Interval	16
15. The variable, stated	16
16. Why this is the governing variable and not a fourth parameter	18
17. The fact-decay rate	19
18. What must be recorded	20
19. Objections	21
Part VI—Three Cases	22
20. Why these three	22
21. Loitering munitions, where the interval is the design feature	22
22. Automated air and missile defence, where the interval is compressed and the practice is old	24
23. Emplaced offline ground systems, where the interval runs to weeks	25
24. What the three cases establish	25
Part VII—Responsibility	26
25. The gap as it is usually stated	26
26. What the interval framework does to the gap	27
27. How this maps onto the existing rules	27
28. The objection that this lets the commander off	28
Part VIII—The Element, Drafted	28
29. What the drafting has to satisfy	28
30. What it would cost the principal states to accept	28
31. The element	30

32. Notes on the drafting	31
33. Adoption without an instrument	31
Part IX—Conclusions	31
34. Conclusions on each element	31
35. What would change the analysis	32
36. Summative remarks	33
Appendix—Worked Records	33
A.1 Terminal air and drone defence	34
A.2 Loitering munition, urban employment within communications range	34
A.3 Emplaced ground system, thirty-day denial task	34
A.4 What the three records show	35
A.5 What this Appendix does not establish	35
Postscript—The September 2026 Outcome	35
P.1 What happened	35
P.2 The formula analysed in Part III has been deleted	36
P.3 The characterisation has changed, and the change bears on section 8.1	37
P.4 Section 16 is borne out	37
P.5 Part IV is borne out, by subtraction	38
P.6 The element’s route of entry has changed; its constraints have not	38
P.7 What has not changed	38
Table of Instruments and Official Documents	40
Bibliography	40

The Interval, Not the Autonomy

Why Ten Years of Defining Autonomous Weapons Produced Nothing, and What to Regulate Instead: the Gap Between the Last Human Decision and the Effect

Christopher I. V. Farmer

Abstract

The Group of Governmental Experts on lethal autonomous weapons systems reports to the Seventh Review Conference of the Convention on Certain Conventional Weapons in November 2026, after a decade of work. It has produced a rolling text and no instrument. This paper argues that the reason is not the disagreement usually blamed for it, and that the object the debate has organised itself around is the wrong one.

The structural finding comes first, because it changes what a proposal has to do. The Group’s mandate is to formulate, by consensus, a set of elements of an instrument *without prejudging its nature*. States therefore have no mandate to negotiate. A consensus rule combined with an express instruction not to prejudge the outcome is not a slow route to a treaty; it is a procedure with no terminus, and the decade of definitional argument is its symptom rather than its cause. Any proposal offered into the 2026 window must be capable of adoption by a body that cannot negotiate, which is a severe constraint and is treated here as the design brief rather than as a complaint.

The substantive argument is that the decisive variable is temporal, and that both halves of it are already in the debate without being joined. The rolling text’s central formula—“context-appropriate human judgement and control”—carries its operative tier. That formula cannot be verified: no state can be held to it and no inspector can check it, because nothing in it specifies what would count. Within

the same text sits a scale-of-operation measure—limiting “the types of targets, duration, geographical scope and the number of engagements”—in which *duration* appears as one component among four, unexplained. Outside the text, the principal gap identified by its most careful critics is the absence of any requirement that a user adequately understand a system **and its effects in the context of use**.

This paper’s claim is that those two are halves of one variable, and that the variable is not one component among four but the one determining whether the others can be satisfied at all:

the interval between the last human decision and the effect, together with what was knowable across that interval.

Distinction, proportionality and precautions are judgements made at a moment, on facts obtaining at that moment. The legal question is whether those facts still hold when the effect arrives. “Limit duration” says only that a system should not run too long. The reframing says that the lawful interval is a function of **the rate at which the relevant facts decay**, that this rate is a property of the situation rather than of the machine, and that it is estimable in advance and checkable afterwards. A system may lawfully loiter for a week over empty ocean and unlawfully for ninety seconds above a moving crowd. No single figure belongs in a treaty. A rule relating permitted interval to fact-decay rate does, and unlike “context-appropriate” it can be complied with, audited and breached.

The second finding concerns the review obligation and it is precise. Article 36 of Additional Protocol I requires a State to determine the legality of a new weapon, means or method of warfare that it studies, develops, acquires or adopts. The temporal application of that provision is routinely described as very broad, and within the acquisition pipeline it is. All four listed stages are nevertheless pre-employment. Article 36 has wide reach across procurement and **no reach at all across operational life**. Where a system’s behaviour is a function of data encountered since adoption, the object employed is not the object reviewed, and the provision designed to catch novel means has no temporal purchase on the novelty that matters. The defect is not that Article 36 is too narrow in subject matter. Its breadth runs along the wrong axis.

The framework is then applied to the three cases that decide whether it is useful: loitering munitions, where the interval is the design feature; automated air and missile defence, where the interval is compressed below human reaction time and the practice is long-standing and largely unobjectionable; and emplaced offline ground systems of the kind specified in the companion engineering paper, where the interval runs to weeks and a recall channel may not exist. The framework distinguishes them. “Autonomy” does not.

The paper closes with drafted protocol text in the form it would have to take to be tabled, together with an account of what it would cost each of the principal states to accept, and why China’s position is the one that determines an instrument’s reach.

Keywords: lethal autonomous weapons systems; Convention on Certain Conventional Weapons; Article 36 weapons review; meaningful human control; precautions in attack; international humanitarian law; military artificial intelligence

Postscript, 15 October 2026. The Group of Governmental Experts adopted its report on 4 September 2026 and the text this paper analyses has been superseded. The argument is unaffected and in three respects is borne out. The changes, and what they do to each claim, are set out in the Postscript following the Appendix. The body of the paper is reproduced unchanged from the version of record.

Note, 18 October 2026. The companion paper on closed-system ground robotics, cited in the body as in preparation, is *The Emplaced Machine: Energy, Detection, Recall and Control in Closed Ground-Denial Systems* (preprint, v2.12.5, 17 October 2026) <https://civfarmer.github.io/publications/manuscripts/the-emplaced-machine-v2.12.5.pdf>, <https://doi.org/10.5281/zenodo.23073940>.

Note on sources and method

Every authority in this paper has been checked against a primary or authoritative record before citation, and anything that could not be verified has been removed rather than approximated. Treaty text is cited to the instrument. Statements of position are cited to the document in which they were made, and where a position is reported by a third party rather than stated by the state itself, the note says so.

Three cautions apply to the material, and they govern how the paper's claims should be read.

The rolling text is a working document and not a draft treaty. It has no agreed status, its contents have changed between sessions, and citing it establishes what has been discussed rather than what has been decided. Where this paper relies on its content it identifies the version in the note.

Non-governmental analysis is used for what it observes, not for what it advocates. Several of the sharpest available readings of the negotiating record come from organisations campaigning for a particular outcome. Their observations about the text are checkable and are used; their conclusions about what should follow are not adopted, and where this paper reaches a similar conclusion it does so on its own reasoning.

The author is not neutral about the process and says so. The argument that the mandate is structurally incapable of producing an instrument is a criticism, and a reader should discount for it accordingly. It is offered because the evidence for it sits on the face of the mandate, and because a proposal ignoring that constraint would be worse than useless.

Part I—Introduction

1. The problem stated

A decade of multilateral work on autonomous weapons has produced a rolling text and no instrument.

The chronology is not in dispute. The Group of Governmental Experts on lethal autonomous weapons systems met for ten days in Geneva across 2025, on 3 to 7 March and 1 to 5 September, with an informal consultation in New York on 12 and 13 May.¹ It meets again on 2 to 6 March and 31 August to 4 September 2026, and it is to submit a final report to the Seventh Review Conference of the Convention on Certain Conventional Weapons in November 2026.² In the meantime the General Assembly's First Committee adopted a resolution on autonomous weapons for the third consecutive year on 6 November 2025, by 156 votes to five with eight abstentions.³ In September 2025, forty-two states delivered a joint statement calling for negotiations to begin on the basis of the rolling text.⁴

That is a process with overwhelming majority support, a working document, a deadline, and no instrument. The usual explanation is that states disagree about what an autonomous weapon is. This paper argues that the definitional disagreement is a symptom, that its cause sits on the face of the mandate, and that the object the debate has organised itself around is the wrong one.

2. Why the definitional route was never going to arrive

The argument of Part II is stated here in short because everything else depends on it.

The Group's mandate is to formulate, by consensus, a set of elements of an instrument **without prejudging its nature**.⁵ Two features of that formulation are together decisive. Consensus gives every participating state a veto over every element. And an express instruction not to prejudge the instrument's nature removes

¹The sessions are listed at UNODA, 'Convention on Certain Conventional Weapons—Group of Governmental Experts on Lethal Autonomous Weapons Systems (2026)' <https://meetings.unoda.org/ccw-/convention-on-certain-conventional-weapons-group-of-governmental-experts-on-lethal-autonomous-weapons-systems-2026> accessed 5 August 2026.

²*ibid.*

³Stop Killer Robots, '156 states support UNGA resolution on autonomous weapons' <https://www.stopkillerrobots.org/news/156-states-support-unga-resolution/> accessed 5 August 2026, reporting the First Committee vote of 6 November 2025.

⁴Joint statement to the September 2025 session of the CCW GGE LAWS' (UNODA documents library) <https://docs-library.unoda.org/> accessed 5 August 2026.

⁵The mandate formula is set out and analysed at Part II §5.

the one question whose settlement would force the others: a body that may not decide whether it is drafting a treaty, a political declaration or a set of best practices cannot resolve what its text is for, and a text whose purpose is undetermined cannot be tested against anything.

The consequence is that states have **no mandate to negotiate**. They have a mandate to accumulate agreed language, which is a different activity with a different terminus, and the terminus is wherever agreement runs out. A decade of argument about the definition of autonomy is what a body does when it has been instructed to produce agreement without being permitted to produce an outcome.

That matters for a proposal and not merely for a diagnosis. A proposal offered into the 2026 window must be adoptable by a body that cannot negotiate. It must therefore be capable of being stated as an element rather than bargained as a package; it must not require a state to abandon a position it has defended for a decade; and it must be assessable against something other than agreement. Those are the constraints Part VIII's drafted text is written to satisfy, and they are why it is short.

3. Why the object is wrong

In the revision under which most of the debate ran, the rolling text's central prohibition was on employing systems that operate without "context-appropriate human control and judgement". In the status of 18 December 2025, which was the text before the 2026 sessions, the same formula stands instead as a stated need—"context-appropriate human judgement and control is needed to ensure" compliance—supported by measures states "should" take.⁶

Consider what a state would have to do to breach the prohibition, or fail the requirement, and what an inspector would have to establish to show it. Neither question has an answer. "Context-appropriate" is not a standard whose content varies by context; it is a placeholder for content nobody has supplied. A rule that cannot be breached is not a weak rule but an absent one, and agreement accumulated around it is agreement about a form of words.

The migration between revisions is itself evidence. From the November 2024 revision to the December 2025 status the formula moved down the normative scale, from a prohibition to a stated need supported by recommendatory measures. Part III §8.2 takes up what that migration means.

The positive claim is that the content is available, that both halves of it are already in the debate, and that they have not been joined.

The first half is inside the text. The rolling text would have states limit "the scale of the operation, which includes the types of targets, duration, geographical scope and the number of engagements" a system may undertake autonomously.⁷ *Duration* appears there as one component among four, with no account of why it belongs or what sets it.

The second half is outside the text, as the principal gap its most careful critics identify: there is no clear requirement that a user adequately understand a system **and its potential effects in the context of use**.⁸

Joined, they are one variable:

the interval between the last human decision and the effect, together with what was knowable across that interval.

⁶GGE on LAWS, 'Rolling text, status date: 18 December 2025' (UNODA documents library, CCW GGE on LAWS (2026)) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf) accessed 14 August 2026, part III paras 6–7. The prohibition formulation is that of the revised rolling text of 8 November 2024, discussed through the 2025 sessions; there was an intermediate revision of 12 May 2025. The formula and the structure are examined at Part III §8.

⁷Rolling text, status 18 December 2025 (n 6), part III para 7(B); and see Part III §8.2.

⁸Article 36, 'Key comments on the current CCW GGE "rolling text" on autonomous weapons' (17 February 2025) <https://article36.org/statements-updates/key-comments-on-the-current-ccw-gge-rolling-text-on-autonomous-weapons> accessed 5 August 2026.

The reason this is the governing variable rather than a fourth parameter is that the substantive rules are all rules about a moment. Distinction requires a judgement that a person or object is a lawful target. Proportionality requires an assessment of expected incidental harm against anticipated military advantage. Precautions require that feasible steps be taken to verify. Each is made at a time, on facts obtaining at that time. **The legal question an autonomous system raises is not whether a machine made the judgement but whether the facts the judgement rested on were still true when the effect arrived.**

That reframing yields what the definitional route cannot. “Limit duration” says a system should not run too long without saying what too long is, and any figure written into a treaty would be wrong in most cases. The interval framework says the permitted interval is a function of **the rate at which the relevant facts decay**—a property of the situation rather than of the machine, which a commander already estimates informally before every operation, and which can be recorded before the fact and audited after it. A loitering system over open ocean and the same system over a market are governed by one rule producing different answers, because the facts decay at different rates.

3.1 The structure of the argument

Part II reads the mandate closely and argues that the failure to produce an instrument is a property of the instruction rather than of the disagreement, that the claim is testable, and that four constraints follow for anything offered into the 2026 window.

Part III sets out what the rolling text actually says, box by box, and establishes that its central formula cannot be breached—and in the December 2025 drafting is no longer even stated as a prohibition—before identifying the two components this paper joins.

Part IV takes the review obligation that ought already to have caught the problem and states its defect precisely: Article 36’s breadth runs along the procurement axis and has no reach across operational life. The remedy is a second occasion rather than a wider scope, and precautions in attack already supply the axis it should run along.

Part V is the paper’s core. It states the variable, credits what it owes to the literature on meaningful human control, distinguishes it from the unpredictability test, argues that it governs rather than sits alongside the other limitation parameters, specifies the fact-decay rate and what must be recorded, and answers five objections.

Part VI applies the framework to loitering munitions, automated air and missile defence, and emplaced offline ground systems, and shows that it separates them where the autonomy framing cannot.

Part VII takes the responsibility gap and argues that it narrows—not because the philosophical problem is solved but because the question becomes answerable—then maps the framework onto individual, command and state responsibility as they stand.

Part VIII tests the proposal against the positions of China, the United States, the European supporters of negotiation and the United Kingdom, drafts the element, annotates the drafting, and sets out four routes to adoption, three of which need nobody’s agreement.

Part IX concludes and states what would change the analysis. The Appendix works the required record for each of Part VI’s three cases, so that the burden the proposal imposes can be judged rather than taken on trust.

A reader pressed for time should take Part I, the table at Part VI §24, the element at Part VIII §31 and the conclusions at Part IX §34; the intervening Parts are where those claims are earned.

4. What is new here, and what is not

4.1 What is not new

That autonomy in weapons raises problems of distinction, proportionality and precautions is the subject of the field and requires no restatement.

That a temporal limit belongs among the constraints on an autonomous system is already in the rolling text, and this paper does not claim to have introduced it.⁹ The claim made here concerns what governs that limit, not that the limit exists.

That users must understand a system's effects in the context of use is identified as a gap by others and is credited to them.¹⁰

That Article 36 review is weak in practice is established: only a limited number of states are known to have review mechanisms at all, and there is little practice indicating what internal procedures should look like.¹¹

That existing international humanitarian law already governs these systems, and that they are not unlawful per se, is the position of the leading treatment of the question and is not disputed here.¹² This paper's argument is compatible with it and, if anything, depends on it: the contention is not that new law is needed but that one existing rule—the obligation to cancel or suspend—already looks along the right axis and lacks only the content that would make it operable.

That unpredictable systems should be ruled out is the position of the International Committee of the Red Cross and is not this paper's contribution.¹³ Part V §15.2 argues that the unpredictability test and the interval framework catch different failures and that a serious instrument would contain both.

That an implicit requirement of human judgement is already present in the law of armed conflict is Asaro's argument, and this paper agrees with it while differing on what follows.¹⁴

That responsibility for autonomous force is contested is a literature of long standing, which Part VII engages rather than announces.

4.2 What is offered as new

First, the mandate as cause rather than context (Part II). The definitional impasse is generally treated as a disagreement between states. The argument here is that a consensus rule combined with an instruction not to prejudge the instrument's nature is a procedure without a terminus, and that this is visible on the face of the mandate rather than inferable from the conduct of the parties. If that is right, proposals aimed at resolving the disagreement are aimed at the wrong thing.

Second, the joining of the two halves (Part V). Duration sits inside the text unexplained; understanding-in-context sits outside it as an acknowledged gap. That they are halves of one variable, and that joined they supply the content “context-appropriate” lacks, is the paper's central claim. It is modest in that both components are already present in the debate. It is not modest in consequence, because the joined variable is measurable where the formula it replaces is not.

Third, the fact-decay rate as the quantity that sets the lawful interval (Part V §17). This is the operative content of the second claim and the part most exposed to objection. It proposes that the permitted interval between the last human decision and the effect be governed by how quickly the facts underpinning that decision cease to hold, that the rate be estimated and recorded before employment, and that the recorded estimate be what conduct is later assessed against.

⁹Rolling text (n 6), and see Part III §8.2.

¹⁰Article 36 (n 8).

¹¹See Part IV §11.

¹²Michael N Schmitt and Jeffrey S Thurnher, “Out of the Loop”: Autonomous Weapon Systems and the Law of Armed Conflict’ (2013) 4 *Harvard National Security Journal* 231. See also William H Boothby, *Weapons and the Law of Armed Conflict* (2nd edn, Oxford University Press 2016), on the framework within which such systems fall to be assessed.

¹³See Part V §15.2 and the ICRC position there cited.

¹⁴See Part V §15.2 and Asaro there cited.

Fourth, the temporal-axis defect in Article 36 (Part IV). The provision’s temporal reach is described in the authoritative guidance as very broad, and across the procurement pipeline it is. Every stage it names is nevertheless pre-employment. Stating the defect as an axis error rather than a scope gap is offered as new, and it changes the remedy: the answer is not to widen what review covers but to give it a second occasion.

Fifth, the application to the three hard cases (Part VI). Loitering munitions, automated air defence and emplaced offline ground systems are usually treated as three arguments about autonomy. Under the interval framework they are three points on one axis, and the framework explains why the second is largely unobjectionable while the first and third are not—which the autonomy framing cannot do without special pleading.

4.3 What would falsify each claim

The **mandate argument** fails if the Group produces an instrument, or agrees a negotiating mandate, before the November 2026 Review Conference. That is a clean test with a date on it.

The **joining claim** fails if some state or commentator has already treated duration and contextual understanding as components of a single variable rather than as separate requirements. The author has not found such a treatment; that is not the same as there being none, and a reader who finds one should treat the claim as reduced to exposition.

The **fact-decay rate** fails if the rate cannot be estimated before employment to a standard capable of supporting later assessment. That is the objection this paper takes most seriously, and Part V §19.1 answers it at length rather than briefly.

The **Article 36 argument** fails if state practice already includes re-review on material change during operational life, which would mean the gap is closed in practice whatever the text says. Part IV asks for that evidence and reports what it found.

The **three-case application** fails if the framework cannot distinguish the cases without importing a judgement it has not derived—that is, if the interval analysis turns out to restate intuitions about which systems are acceptable rather than supply a means of testing them.

Part II—The Mandate as Cause

5. The mandate, read closely

This section aims to establish that the failure to produce an instrument is a property of the instruction the Group was given, and not of the disagreement between the states carrying it out. Our purpose is to show that the conclusion follows from the text of the mandate rather than from an inference about motives, because a claim about motives could not be tested and this one can.

The mandate was decided by the 2023 Meeting of the High Contracting Parties. It instructs the Group to

further consider and formulate, by consensus, a set of elements of an instrument, without prejudging its nature, and other possible measures to address emerging technologies in the area of lethal autonomous weapon systems, taking into account the example of existing Protocols within the Convention, proposals presented by High Contracting Parties and other options related to the normative and operational framework on emerging technologies in the area of lethal autonomous weapon systems, building upon the recommendations and conclusions of the Group, and bringing in expertise on legal, military, and technological aspects.¹⁵

¹⁵Mandate decided by the 2023 Meeting of the High Contracting Parties to the Convention on Certain Conventional Weapons. The text is reproduced in the working papers and briefings of the Group; see UNODA, ‘Briefing by the Chair of the CCW GGE on LAWS on the margins of the First Committee’ <https://disarmament.unoda.org/en/updates/briefing-chair-ccw-gge-laws-margins-first-committee> accessed 5 August 2026.

The Group is to report to the Seventh Review Conference in 2026, completing its work as soon as possible and **preferably before the end of 2025**. The report must be adopted by consensus.¹⁶

Five features of that instruction bear on whether it can terminate. Taken singly, each is ordinary multilateral drafting. Taken together they describe a procedure with no exit.

First, the verbs are “consider and formulate”, not “negotiate”. The distinction is not pedantry: negotiation is an activity directed at an agreed text between parties who have accepted that a text is the object. Consideration and formulation are directed at an output whose status is undetermined, and they are complete whenever the participants say they are.

Second, consensus applies twice. The elements must be formulated by consensus, and the report containing them must be adopted by consensus. Every state therefore holds a veto at both the level of content and the level of the vehicle, and a state content with the present position can exercise the second veto without ever having to defend a position on the first.

Third, the object is “a set of elements of an instrument”, not an instrument. Elements are components whose assembly is somebody else’s task, and the mandate does not say whose. A set of elements can be complete in the sense the mandate requires while remaining incapable of binding anyone.

Fourth, “without prejudging its nature” removes the question whose settlement would force the others. Whether the output is to be a protocol, a political declaration, a set of good practices or a chair’s summary determines what its provisions must be capable of doing. A prohibition drafted for a binding protocol and a prohibition drafted for a declaration are different texts, because only the first has to survive contact with a state that does not want to be bound. Instructed not to settle the nature of the instrument, the Group cannot test its language against the thing the language would have to do—and a text that cannot be tested cannot be finished, only accumulated.

Fifth, and least remarked, the mandate is disjunctive. The Group is to formulate elements **“and other possible measures”**. That phrase supplies a second terminus which is not an instrument at all. A Group unable to agree elements may discharge its mandate by reporting other possible measures, and a state preferring no instrument can pursue that branch in good faith throughout, without obstruction and without ever refusing anything.

6. Why this is a cause and not a description

The obvious objection is that this reverses the order of explanation. Mandates are drafted by the same states that later disagree, so a weak mandate is the *product* of disagreement rather than its cause, and the analysis above merely restates the difficulty in procedural language.

The objection is partly right and the concession should be made plainly: the 2023 mandate was itself adopted by consensus among states holding the positions they hold, and a mandate to negotiate would not have been agreed. To that extent the mandate records a disagreement.

What the objection cannot account for is the **asymmetry the mandate introduces once it exists**. A disagreement between states can be resolved by persuasion, by coalition, by the accumulating cost of holding out, or by a majority moving without the holdouts—all of which are ordinary and all of which have produced instruments elsewhere. None is available here. Persuasion cannot produce a text that has not been permitted to have a nature. A coalition cannot proceed, because the report requires consensus. And the cost of holding out is nil, because a state need never hold out: it can support elements while preferring “other possible measures”, and the mandate treats that as compliance.

The mandate therefore does more than record the disagreement. **It converts a distribution of positions that would elsewhere be resolvable into one that is not**, by removing every mechanism by which such distributions are ordinarily resolved. That is a causal claim about a procedure, and it is testable in the way §7 sets out.

¹⁶ibid. The report is due to the Seventh Review Conference in 2026; the mandate expresses a preference for completion before the end of 2025, and adoption of the report is by consensus in accordance with the Convention’s practice.

7. The evidence, including a deadline that has already passed

Three things are observable and none requires an assumption about what any state wants.

The mandate's own preferred deadline has been missed. Work was to be completed as soon as possible and preferably before the end of 2025.¹⁷ It was not. The Group met twice more in 2026 and the report is now due to a Review Conference in November of that year. A procedure that overruns its own preferred completion date by more than a year, while retaining majority support and a working text, is not a procedure being obstructed by a minority. It is a procedure with no mechanism for stopping.

Overwhelming majority support has not converted into an outcome. The General Assembly's First Committee adopted a resolution on autonomous weapons for the third consecutive year in November 2025, by 156 votes to five with eight abstentions.¹⁸ In September 2025, forty-two states delivered a joint statement calling for negotiations to begin on the basis of the rolling text.¹⁹

The first 2026 session widened the majority and moved nothing. The Group met from 2 to 6 March 2026 and conducted a first reading of the chair's text in its December 2025 status; on the Wednesday night the chair circulated revised formulations for three of its five boxes; and by the close of the week the states endorsing negotiation on the basis of the text had grown from the forty-two of September to more than seventy, after a group of African states joined the call.²⁰ The support widened; the procedure it is offered into is unchanged; and the session of 31 August–4 September 2026 is the Group's last before it must report.²¹

And the resolution as adopted preserves the very disjunction described at §5, in its own text. Resolution 80/57, adopted on 1 December 2025, notes in its preamble the Secretary-General's calls "to commence negotiations of a legally binding instrument on autonomous weapons systems", and then asks the High Contracting Parties, in its single operative call, "to work towards completing the set of elements for an instrument being developed within the mandate of the Group of Governmental Experts ..., **with a view to future negotiations**".²²

Read those two passages against each other. The organ recording the Secretary-General's call to begin negotiating asks the parties to finish assembling elements, and locates the negotiation in a future it does not schedule, mandate or assign to anyone. That is not a criticism of the drafters, who were writing for adoption in a body where the alternative was no resolution. It is the strongest available evidence for this Part's claim: after nine years, **the most that could be agreed by the General Assembly is a request to complete the components of an instrument nobody is yet instructed to negotiate.** The mandate's disjunction has been reproduced at the level above it. Elsewhere in disarmament practice, a majority of that size is decisive. Here it is not, and the reason it is not is the consensus rule at §5.

The one state whose position would be most consequential for a European instrument has not taken one. When France, Germany and more than forty other states agreed in September 2025 that the draft rules were ready to negotiate, the United Kingdom did not join them and did not oppose them.²³ Silence is a lawful and unremarkable position in a consensus body, and that is the point: a procedure in

¹⁷ibid.

¹⁸Stop Killer Robots (n 3).

¹⁹Joint statement (n 4).

²⁰Laura Varella, 'The Final Stretch Before the Finishing Line' (2026) 14(2) CCW Report (Reaching Critical Will, 11 March 2026) <https://reachingcriticalwill.org/images/documents/Disarmament-fora/ccw/2026/gge/reports/CCWR14.2.pdf> accessed 14 August 2026. The revised formulations of 4 March 2026 covered boxes I–III.

²¹Indico.UN, '2026 GGE on LAWS, Second session (31 August–4 September 2026)' <https://indico.un.org/event/1019358/> accessed 14 August 2026.

²²UNGA Res 80/57 (1 December 2025) UN Doc A/RES/80/57, preambular paragraph noting the calls by the Secretary-General to commence negotiations of a legally binding instrument, and operative paragraph 3. The resolution was adopted at the 52nd plenary meeting on the report of the First Committee (A/80/534, para 7).

²³Article 36, 'UK stays mute as France, Germany and 40 more states agree that draft rules on "killer robots" are ready to negotiate' (12 September 2025) <https://article36.org/statements-updates/uk-stays-mute-as-france-germany-and-40-more-states-agree-that-draft-rules-on-killer-robots-are-ready-to-negotiate> accessed 5 August 2026. The characterisation of the United Kingdom's position is that of the reporting organisation; what is verifiable from the record is that the United Kingdom was not among the signatories to the joint statement.

which silence costs nothing will generate silence, and a state that neither supports nor blocks cannot be brought to a decision by anything the procedure contains.

7.1 What follows for a proposal

The analysis is only worth making if it changes what should be offered, and it does. Four constraints follow, and Part VIII's drafted text is written to satisfy them.

It must be statable as an element. A proposal that requires a package, a set of trades, or a sequence of concessions is asking the Group to negotiate, which it may not do. A proposal capable of standing as a single element can enter the rolling text on the terms the mandate actually permits.

It must not require any state to abandon a position. A proposal that resolves the definitional dispute by picking a side loses the states on the other side, and consensus makes that fatal. The interval framework at Part V is offered partly because it is **orthogonal** to the definitional dispute: a state may hold any view about what constitutes an autonomous weapon and still accept a rule about the interval between decision and effect, because the rule does not depend on the classification.

It must be assessable against something other than agreement. The defect at §5's fourth feature is that the Group cannot test its language against a purpose. A rule stated so that compliance and breach are observable supplies its own test, which is a property the Group can use without being permitted to decide what the instrument is.

And it should be capable of operating whether or not an instrument follows. A rule that requires a treaty to have effect is worth nothing if no treaty arrives. A rule expressed as a matter of what a state must record before employing a system, and against which its conduct can later be assessed, can be adopted unilaterally, in national doctrine, in an Article 36 review procedure, or in a coalition's rules of engagement, and it improves matters in each case. That is not an argument against a treaty. It is an argument for drafting the rule so that the treaty is not the only route by which it does work.

Part III—What the Rolling Text Actually Says

8. The text, box by box

This section aims to establish what has in fact been agreed at the level of language, because the argument of Part V is that the rolling text contains one half of the answer without recognising it, and that claim cannot be made without setting out the text. We take the status of 18 December 2025, which was the basis of the first reading at the March 2026 session and remains the text before the second; where the argument depends on how the text has changed, the revisions of 8 November 2024 and 12 May 2025 are identified. At each point we identify whether this paper is relying on the text itself or on analysis of it.²⁴

The document is organised in five numbered parts which the room still calls boxes, and they are boxes rather than articles, which is itself informative: boxes hold material for assembly and articles bind. A reader who has followed the rolling text through its revisions may pass to §10, which states what Parts II and III together establish.

8.1 The characterisation, and a defect in it

The text characterises the systems in issue as “a functionally integrated combination of one or more weapons and technological components, that can identify, select, and engage a target, without intervention by a

²⁴GGE on LAWS, ‘Rolling text, status date: 18 December 2025’ (UNODA documents library, CCW GGE on LAWS (2026)) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf) accessed 14 August 2026. Earlier revisions are of 8 November 2024 and 12 May 2025; revised formulations for boxes I–III were circulated at the March 2026 session (Laura Varela, ‘The Final Stretch Before the Finishing Line’ (2026) 14(2) CCW Report (Reaching Critical Will, 11 March 2026)). The text has no agreed status and its content has changed between sessions; references here establish what was discussed, not what has been decided.

human operator in the execution of these tasks”.²⁵

That formulation does real work and should be credited with it. It identifies the functions that matter—identification, selection, engagement—rather than attempting to define autonomy in the abstract, and it locates the question at the absence of human intervention *in the execution* of those functions rather than in their design. A decade of definitional argument has at least produced a characterisation aimed at conduct.

One objection to it is well made and matters more than it appears. Describing such a system as a functionally integrated combination risks implying that the functions must sit within a single physical unit, when in practice they are dispersed across sensors, processing, communications and effectors that may belong to different platforms and different formations.²⁶ A characterisation that attaches to a unit will therefore miss a system assembled from parts none of which is individually the thing described.

This paper’s framework does not have that difficulty, and the reason should be stated early because it is one of the framework’s practical advantages. **An interval is a property of a chain of events, not of an object.** The question of how long elapsed between the last human decision and the effect, and what was knowable across that period, can be asked of a dispersed system exactly as it is asked of an integrated one, because it does not require the analyst to decide which box the system is in.

8.2 The operative rules, and the phrase that carries them

The December 2025 status arranges the operative material as follows. Part II restates, in five propositions, that international humanitarian law applies in full to these systems and that its obligations fall on states and individuals, not machines. Part III opens with prohibitions—systems of a nature to cause superfluous injury or unnecessary suffering, inherently indiscriminate systems, systems otherwise incapable of use in compliance with IHL, and attacks on civilians—which restate what existing law already prohibits.²⁷ It adds a duty to refrain from manufacturing, acquiring, stockpiling or transferring such systems, with a destruction requirement for what is retained.²⁸ The central formula then arrives twice: “context-appropriate human judgement and control is needed to ensure” that use and effects comply with IHL, and states “should ensure” it “through, inter alia” five families of measures—a responsible chain of command; anticipation and control of effects, including limits on “the scale of the operation, which includes the types of targets, duration, geographical scope and the number of engagements”; a bar on the system significantly modifying its own mission parameters, including through real-time machine learning; timely deactivation, or self-destruct, self-deactivation or self-neutralisation mechanisms; and confinement to a defined perimeter and to military objects.²⁹ A design-and-development duty closes the part.³⁰

In the November 2024 revision the same formula stood as the central prohibition: it was prohibited in all circumstances to employ systems operating **without context-appropriate human control and judgement**. The December status states it as a need, supported by measures states “should” take. Nothing in this paper turns on which formulation survives to November. The analysis at §8.3 applies to both, and the migration from prohibition to recommendation between revisions is Part II at work on the text’s own core.

The two-tier structure has wide support and should be retained; nothing in this paper argues against it. What this paper argues is that the tier that matters is empty—and that the December drafting has made it, in terms, less than a tier.

²⁵ibid, part I para 1. The November 2024 revision had “identify and/or select” and “human user”; the refinement is drafting rather than a change of object.

²⁶Article 36 (n 8). The observation is the analysing organisation’s; the inference drawn from it in the following paragraph is the author’s.

²⁷Rolling text (n 24), part III paras 1–3 and part II.

²⁸ibid, part III para 4.

²⁹ibid, part III paras 6–7.

³⁰ibid, part III para 8. The lifecycle measures—legal review at study, development, acquisition or adoption; review of modifications with significant impact; testing, evaluation and training; bias and automation-bias measures—are at part IV, and the responsibility provisions, including internal investigation of incidents, at part V.

8.3 Why “context-appropriate” cannot do the work

The objection is not that the phrase is vague. Legal standards are routinely vague and function well, because vagueness is resolved by application: reasonableness, proportionality and due diligence acquire content through repeated determination by bodies with the authority to determine them.

“Context-appropriate human judgement and control” is differently situated in three respects, and the combination is what defeats it.

There is no determining body. No court, tribunal or treaty organ has jurisdiction to hold that a given deployment fell short. The phrase will not accumulate content by application because there is no application, and the states that would apply it are the states whose conduct it governs.

The variable is unnamed. “Reasonable” is vague about a threshold on a known scale—reasonable *care*, reasonable *time*. “Context-appropriate” does not say appropriate along which dimension. Control could mean supervision, or ability to intervene, or authorship of the decision, or knowledge of the circumstances, and the phrase does not choose. A state may therefore comply on any reading it prefers and no reading is wrong.

Compliance is not observable. Even granted a determining body and a chosen dimension, a state’s conduct in respect of the phrase leaves no trace. Nothing must be recorded, nothing must be produced, and there is no artefact after the fact from which the presence or absence of context-appropriate control could be established.

The first 2026 session put the second respect on display. During discussion of the revised part III, the United States proposed replacing “human control” with “good faith human judgement and care”, and many delegations rejected the substitution as insufficient.³¹ A negotiation conducted at that level is a dispute about which unmeasured words to prefer. Neither formula names the dimension, the estimate, or the record from which an inspector could distinguish compliance from breach; Part V offers that content, for whichever formula survives.

The consequence is the finding of this Part. **A prohibition that cannot be breached is not a weak prohibition. It is an absent one**—and the December 2025 status no longer drafts the formula as a prohibition at all. The agreement accumulated around this form of words is agreement about a form of words. That is not a criticism of the states that agreed it: under a mandate that forbids prejudging the instrument’s nature, language which does not commit anyone is precisely the language most likely to attract consensus, and Part II explains why such language is what the procedure selects for.

9. The gap the text’s critics identify

The most careful published analysis of the rolling text identifies its principal shortcoming as the absence of a clear requirement that **users adequately understand a system and its potential effects in the context of use**.³²

That was written against the text as it stood through 2025, and the December status answers part of it. Part IV now requires that a system’s capabilities, limitations and anticipated effects “are understood”, and specifies the means: testing and evaluation, realistic simulation, guidance, training and instructions.³³ Each of those is discharged before employment, and what they produce is understanding of the system in its expected use. The state of knowledge of the person who authorises *this* employment—what they knew of the situation the effect would land in, and when—is still nowhere required. The gap survives in its exact form.

The observation is more important than its brief statement suggests. Every requirement in §8.2 attaches either to the system—predictability, reliability, explainability—or to limits imposed on its employment. None

³¹Laura Varella, ‘The Final Stretch Before the Finishing Line’ (2026) 14(2) CCW Report (Reaching Critical Will, 11 March 2026) <https://reachingcriticalwill.org/images/documents/Disarmament-fora/ccw/2026/gge/reports/CCWR14.2.pdf> accessed 14 August 2026. The reported formulation is the reporting organisation’s record of the United States’ proposal; the inference drawn from the episode is the author’s.

³²Article 36 (n 8).

³³Rolling text (n 24), part IV para 4; the formula recurs across the lifecycle at part IV para 3.

attaches to the **state of knowledge of the person who authorises it**. A system may be predictable in the laboratory and unpredictable in the valley it is flown into; explainability is a property of the model and not of the officer’s grasp of what it will do here, today, against this population, in this weather.

The same analysis separates automation bias from harmful bias and treats the first as a distinct risk to meaningful control, on the ground that it produces nominal rather than actual human decision-making.³⁴ The distinction is right and it sharpens the point. A requirement of human control that can be satisfied by a human who defers is satisfied by nothing, and the difference between a decision and a ratification is a difference in what the human knew and had time to weigh.

Both observations converge on a quantity the text does not contain: **what the authorising human actually knew about the situation the effect would land in, and when they knew it**.

10. What Parts II and III establish together

Three things, and they set up the remainder.

The procedure cannot terminate, so a proposal must be capable of doing work without waiting for it. That is Part II.

The operative prohibition is unverifiable, so the useful contribution is content rather than agreement. That is §8.3.

And the missing content has two components which are already present in the debate but separated: **dura-****tion**, sitting inside the text as one component of its scale-of-operation measure with no account of what sets it; and **understanding in the context of use**, sitting outside the text as its acknowledged principal gap.

Part V argues that these are halves of one variable, that joined they supply what “context-appropriate” lacks, and that the joined variable is measurable, recordable and capable of being breached. Part IV takes first the review obligation that ought already to have caught the problem, and explains why it does not.

Part IV—The Review Obligation and the Wrong Axis

11. What Article 36 requires, and when

This section aims to establish that the provision designed to catch novel means of warfare has no purchase on the novelty that matters here, and to state the defect precisely, because the imprecise version of the criticism is common and is answerable.

Article 36 of Additional Protocol I provides that in the study, development, acquisition or adoption of a new weapon, means or method of warfare, a High Contracting Party is under an obligation to determine whether its employment would, in some or all circumstances, be prohibited by the Protocol or by any other rule of international law applicable to that Party.³⁵

The authoritative guidance describes the temporal application of that obligation as **very broad**, requiring assessment at the stages of study, development, acquisition and adoption.³⁶ That description is accurate. A State cannot wait until a weapon is in service; the obligation bites from the earliest stage at which the weapon is being studied, and it captures acquisition of an existing weapon as well as development of a new one.

Now observe what the four stages have in common. **Study, development, acquisition and adoption are all pre-employment**. The obligation is broad across the procurement pipeline and terminates where the pipeline delivers. Nothing in Article 36 attaches to the weapon in service.

³⁴ibid.

³⁵Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I) (adopted 8 June 1977, entered into force 7 December 1978) 1125 UNTS 3, art 36.

³⁶International Committee of the Red Cross, *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* (ICRC 2006), published at (2006) 88 *International Review of the Red Cross* 864. The Guide was prepared following an expert meeting hosted by the ICRC in January 2001.

That is not an oversight. For the weapons the provision was drafted against it is exactly right, because their properties are fixed at manufacture. A munition's fragmentation pattern, a projectile's deformation on impact, a laser's output, a mine's fusing—each is a property of the thing as built, and a determination made before adoption remains true across the service life. The temporal reach is broad along the axis on which those weapons vary, and for them there is no second axis to reach along.

12. Where that fails, stated precisely

The imprecise criticism is that Article 36 is too narrow, or that autonomous systems are too complex to review. Neither is the defect and both are answerable: complexity is a resourcing problem, scope is a drafting problem, and States with serious review mechanisms already assess software-defined systems.³⁷

The defect is structural. **A system whose behaviour is a function of data encountered after adoption is not, at the point of employment, the object that was reviewed.**

Three cases, in ascending order of difficulty.

A system updated in service. The reviewed object has been replaced by a different one. Practice varies, and a serious review mechanism will re-review on a material software change, but nothing in Article 36 compels it, because the update is not a new weapon the State is studying, developing, acquiring or adopting. It is the same weapon, altered.

A system that learns in service. Here the alteration has no author and no release. There is no version to re-review and no moment at which a decision to change was taken; the object drifts. A review conducted at adoption certifies a distribution of behaviours the system has since left, and a State re-reviewing on every material change would be re-reviewing continuously, which is to say doing something Article 36 does not describe.

A system whose behaviour depends on a context it was not reviewed against. This is the case that matters most and the one the first two obscure. Even a wholly static, fully deterministic system that has never been updated may behave lawfully in one environment and unlawfully in another, because discrimination performance is a joint property of the model and the scene. The object has not changed at all. What has changed is the world it has been asked to operate in, and Article 36 has nothing to say about that, because it reviews the weapon rather than the employment.

The defect is therefore not that the provision's reach is too narrow. Its breadth runs along the wrong axis. Wide across the procurement pipeline, zero across operational life—and the properties determining whether these systems comply with the law of targeting vary along the second axis and not the first.

13. What the practice adds, and it is not encouraging

Two features of State practice compound the structural point, and both are established rather than argued.

Few States review at all. Only a limited number are known to have put mechanisms or procedures in place to conduct legal reviews of weapons, and there is little practice indicating what internal procedures should be established.³⁸ The obligation is universal among Parties to the Protocol; the compliance is not.

Article 36 does not specify manner or authority. It does not state in what manner or under what authority reviews are to be constituted.³⁹ That is a deliberate margin of appreciation with the ordinary consequence: practice varies widely, comparison between States is difficult, and no external standard exists against which a review's adequacy could be measured.

The combination matters for §12. Where a provision is weakly implemented and unspecified as to procedure, the natural repair is to read it purposively—to say its object requires re-review when the reviewed object

³⁷On the scope and conduct of review generally see Boothby (n 12), and Justin McClelland, 'The Review of Weapons in Accordance with Article 36 of Additional Protocol I' (2003) 85 *International Review of the Red Cross* 397.

³⁸*ibid.*

³⁹*ibid.*

materially changes. That reading is available and some States take it. But a purposive reading cannot supply what the structure lacks. Even a State that re-reviews diligently on every update is reviewing the *system*; nothing in the obligation directs attention to the **employment decision**, or to what the deciding human knew about the environment the effect would land in.

Whether any State’s practice already includes a second determination of that kind is a question this paper puts and cannot answer from the public record. Review procedures are not generally published, and their contents are among the least transparent parts of an already opaque practice. If such practice exists, the argument of Part V is a proposal to generalise it rather than to invent it, and the author would prefer that to be the case.

14. Why the answer is a second occasion rather than a wider scope

The remedy follows from the diagnosis and differs from what is usually proposed.

If the defect were scope, the answer would be to enlarge what a review must consider—assessment of learning behaviour, of training data, of failure modes under distributional shift. Proposals of that kind exist and are sensible as far as they go.⁴⁰ They do not reach the third case at §12, because no amount of pre-employment assessment can establish how a system will perform in an environment nobody has yet specified.

If the defect is an axis error, the answer is a **second occasion** on which a determination is made: not attaching to the weapon before adoption, but to the **employment** before it occurs, and asking a different question. Not “is this weapon capable of lawful use in some circumstances”, which is Article 36’s question and which nearly every such system passes, but:

in the circumstances now obtaining, and across the interval this employment will run, are the facts on which the targeting judgement rests likely to hold?

That question is the subject of Part V. Three things about it belong here, because they are what make it a repair rather than an aspiration.

It is already required in substance. Precautions in attack oblige those who plan or decide upon an attack to do everything feasible to verify the objective, to take all feasible precautions in the choice of means and methods, and to cancel or suspend an attack if it becomes apparent that the objective is not a military one or that the incidental harm would be excessive.⁴¹ The obligation to cancel or suspend is continuous and attaches to the attack rather than to the weapon. **It is the one existing rule that already looks along the right axis.** What it lacks is an account of how it applies where the system executing the attack cannot be reached, and Part V supplies one.

It requires no amendment of Article 36. A State may adopt the second occasion within its own review and targeting procedures without any change to the Protocol—which matters given Part II’s finding that the multilateral route cannot terminate.

And it produces an artefact. A determination made before employment, recorded, and expressed in terms of an interval and an information state is a document that exists afterwards. That is the difference between a rule capable of assessment and the formula at Part III §8.3 which is not, and it is why this paper’s proposal is framed around a record rather than around a standard of conduct.

Part V—The Interval

15. The variable, stated

This section states the paper’s central claim, positions it against the existing literature on meaningful human control, and then does the work the claim requires: specifying what sets the permitted interval, what must

⁴⁰SIPRI, *Implementing Article 36 Weapon Reviews in the Light of Increasing Autonomy in Weapon Systems* (SIPRI Insights on Peace and Security No 2015/1) <https://www.sipri.org/sites/default/files/files/insight/SIPRIInsight1501.pdf> accessed 5 August 2026.

⁴¹Additional Protocol I (n 35), art 57(2).

be recorded, and how a breach would be shown.

The variable is:

the interval between the last human decision bearing on an engagement and the effect of that engagement, together with what was knowable to the deciding human across that interval.

Two components, and neither is new to the debate. Duration sits inside the rolling text as a component of its scale-of-operation measure.⁴² Understanding of a system's effects in the context of use sits outside the text as the principal gap its critics identify.⁴³ **What is offered here is the joining**, and the claim that joined they are not two requirements but one variable with two dimensions, which together determine whether the substantive rules of targeting can be satisfied at all.

15.1 What this owes to the existing literature, which is a good deal

The concept of meaningful human control has been examined for a decade, and the most useful work on it anticipates part of this argument. That should be stated before the contribution rather than after it.

Horowitz and Scharre's primer set out the concept's components and its ambiguities early, and identified the three elements—information, action and accountability—that later formulations have largely retained.⁴⁴ Roff and Moyes developed the requirement into a set of key elements bearing on predictability and the human's role in the decision.⁴⁵ Most directly for present purposes, Ekelhof argued that abstract concepts are of little use if they ignore the operational context in which they must be applied, and that any workable account must recognise the **distributed** character of control in military decision-making, attending to the whole chain including humans exercising control at points less proximal to the system's operation.⁴⁶

Those propositions are correct and this paper adopts them. The criticism of “context-appropriate human judgement and control” at Part III §8.3 is in substance an application of Ekelhof's general point to a particular formula, and it is credited accordingly rather than presented as a discovery.

Where this paper goes further is in supplying the quantity the distributed account leaves open. If control is distributed along a chain, the question becomes **where the chain ends and what happens after it does**. In an ordinary engagement the chain terminates in a decision made on information current at the moment of effect, or close enough that the difference does not matter. In the engagements this paper concerns, the chain terminates and the system continues. **The interval is the length of the chain's unsupported tail**, and the information state is what the last human on that chain knew about the ground the tail would run across. That is a measurement of Ekelhof's structure rather than a rival to it.

15.2 What the framework catches that the unpredictability test does not

One further position must be distinguished, because it is the most developed alternative and the two are complementary rather than competing.

The International Committee of the Red Cross has recommended that unpredictable autonomous weapon systems be expressly ruled out, on the ground of their indiscriminate effects, and that this is best achieved by prohibiting systems designed or used such that their effects cannot be sufficiently understood, predicted and explained. It considers that some such systems would already be inherently indiscriminate under existing law, including those whose effects in their normal or expected circumstances of use could not be sufficiently

⁴²Rolling text (n 24), part III para 7(B).

⁴³Article 36 (n 8).

⁴⁴Michael C Horowitz and Paul Scharre, *Meaningful Human Control in Weapon Systems: A Primer* (Center for a New American Security Working Paper, March 2015).

⁴⁵Heather M Roff and Richard Moyes, *Meaningful Human Control, Artificial Intelligence and Autonomous Weapons* (Article 36 Briefing Paper, April 2016).

⁴⁶Merel A C Ekelhof, 'Moving Beyond Semantics on Autonomous Weapons: Meaningful Human Control in Operation' (2019) 10 *Global Policy* 343.

understood, predicted and explained.⁴⁷

That test is directed at **a property of the system**: whether its behaviour can be anticipated. The interval framework is directed at **a property of the situation**: whether the facts the authorisation rested on survive to the effect. The two catch different failures, and neither subsumes the other.

A perfectly predictable system—one whose behaviour is fully understood, and which does precisely what it was designed to do—may still deliver an effect on a proportionality judgement that expired days earlier. Nothing in the unpredictability test reaches that, because nothing about the system was unpredictable. Conversely an unpredictable system is objectionable however short its interval, and the interval framework does not reach that, because the failure is in the machine rather than in the passage of time.

They are therefore two necessary conditions rather than rival formulations of one, and a serious instrument would contain both. This paper argues only that the second is missing and can be supplied.

The Committee's fullest statement of position confirms both the proximity and the difference. Its October 2025 position paper requires users to take into account "all reasonably foreseeable changes in circumstances over the entire area and duration of the AWS operation between activation and any potential strike"; identifies "a significant risk that changes in the operating environment would invalidate legal determinations made and precautions taken at the time of the AWS activation"; and observes that the likelihood and scale of such changes "increase with the size of the area, the duration of the operation and the dynamic nature of the situation".⁴⁸ That is the decay of a targeting judgement, stated qualitatively, with its governing variables named. What the position does not do is make the observation operable: nothing requires the rate of change to be estimated before employment, recorded, or made the measure against which the employment is later assessed. The framework of this Part is that step. And the design features the Committee would require of systems that do not permit timely human intervention—"an effective self-destruction or self-neutralization mechanism and have a backup self-deactivation feature"—are the effect-end architecture the interval rule needs, now mirrored in the rolling text's own deactivation and perimeter measures.⁴⁹

There is a further point of agreement worth recording. Asaro argued that an implicit requirement for human judgement can be found within the existing law of armed conflict, and rested a case for prohibition partly on that foundation.⁵⁰ This paper agrees that the requirement is implicit in the existing law and disagrees about what follows: the argument here is that making it explicit and temporal is more useful than building a prohibition on it, because an explicit temporal requirement can be complied with, recorded and assessed, whereas a prohibition requires a definition of the prohibited class and Part III explains why that route has not arrived.

16. Why this is the governing variable and not a fourth parameter

The rolling text folds duration into its scale-of-operation measure, alongside the types of targets, geographical scope and the number of engagements.⁵¹ Listing implies parity. The argument here is that the four are not of equal standing, because three are constraints on an engagement while the fourth is a constraint on the **continuing validity of the judgement that authorised it**.

Each substantive rule of targeting is a rule about a moment.

Distinction requires a determination that a person or object is a lawful target.⁵² The obligation is conven-

⁴⁷International Committee of the Red Cross, *ICRC Position on Autonomous Weapon Systems* (12 May 2021), reproduced with its background paper at (2021) 102 *International Review of the Red Cross* 915 <https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systems> accessed 14 August 2026.

⁴⁸ICRC, *Autonomous Weapon Systems and International Humanitarian Law* (position paper, October 2025; version with executive summary, March 2026) 18–19 https://www.icrc.org/sites/default/files/2026-03/4896_002_Autonomous_Weapons_Systems_-_IHL-ICRC.pdf accessed 14 August 2026.

⁴⁹*ibid* 18; Rolling text (n 24), part III para 7(D)–(E).

⁵⁰Peter Asaro, 'On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making' (2012) 94 *International Review of the Red Cross* 687.

⁵¹Rolling text (n 24), part III para 7(B).

⁵²Additional Protocol I (n 35), arts 48, 51 and 52.

tional and customary alike.⁵³ That determination is made at a time, on the appearance and conduct of the person or object at that time, and on intelligence of a particular vintage.

Proportionality requires that expected incidental civilian harm not be excessive in relation to the concrete and direct military advantage anticipated.⁵⁴ It is Rule 14 of the customary study in materially identical terms.⁵⁵ Both sides of that comparison are estimates about a future state of affairs, formed at a moment on the facts then obtaining. Who is present, what the weather is doing, what else is happening nearby, and what the objective is worth are all quantities that change.

Precautions require those who plan or decide upon an attack to do everything feasible to verify the objective, to choose means and methods with a view to minimising incidental harm, and—the obligation that matters most here—to **cancel or suspend** an attack if it becomes apparent that the objective is not a military one or that the incidental harm would be excessive.⁵⁶ The customary formulation of precautions is instructive for this paper’s argument, because the measures it enumerates include target verification, the choice of appropriate means and methods, assessment of the effects of attacks, advance warning, and **ensuring appropriate control**—so that control already appears within the precautions rule rather than as a free-standing requirement invented for autonomous systems.⁵⁷

Three consequences follow, and the third is the one the debate has missed.

The rules are satisfied or breached by reference to facts, not intentions. A commander who forms an honest and reasonable proportionality judgement on facts that have since changed has not thereby made the attack lawful. The obligation to cancel or suspend exists precisely because the law contemplates that facts change between decision and effect.

Every weapon has an interval, and the law already accommodates short ones. A rifle’s interval is milliseconds; an artillery mission’s is minutes; a cruise missile’s may be hours. The law has never required the interval to be zero. What it requires is that the interval be short enough, relative to the rate at which the relevant facts change, that the judgement remains good—and where it is not, that the attack be capable of cancellation.

And that is why duration is not a parameter alongside the others. Restricting target types, geography or scale narrows what an engagement may do. Restricting the interval determines whether the authorisation was still valid at the moment the effect occurred. A system operating within perfectly drawn geographic and target-type limits, at modest scale, may still deliver an effect on a proportionality judgement that expired a fortnight earlier. The other three parameters cannot detect that. **The interval is the parameter on which the others depend for their meaning.**

17. The fact-decay rate

The operative content of the claim is that the permitted interval is not a fixed number but a function of how quickly the facts underlying the authorising judgement cease to hold.

Call that **the fact-decay rate**: the rate at which the propositions on which the distinction and proportionality determinations rest lose their truth.

17.1 What it is a property of

It is a property of the **situation**, not of the machine. That point makes the framework usable and it inverts the usual assumption that the relevant properties belong to the system.

⁵³Jean-Marie Henckaerts and Louise Doswald-Beck (eds), *Customary International Humanitarian Law, Volume I: Rules* (Cambridge University Press 2005) rules 1 and 7.

⁵⁴*ibid.*, arts 51(5)(b) and 57(2)(a)(iii).

⁵⁵*ibid.*, rule 14.

⁵⁶*ibid.*, art 57(2), and in particular art 57(2)(b).

⁵⁷*ibid.*, rule 15. The enumeration of precautionary measures there includes ensuring appropriate control, which is the point relied on in the text.

The propositions that decay are of a familiar kind. That this building is being used for military purposes. That no civilians are within the effects radius. That the vehicles in this column are military. That the pattern of life in this area is as reconnaissance reported. Each has a characteristic lifetime, and the lifetime is set by the environment rather than by the weapon: a hardened static installation in a depopulated area supports propositions holding for weeks, and a road junction in a city at dusk supports propositions holding for minutes.

Two systems with identical autonomy, identical software and identical review histories are therefore governed differently by the same rule, because they are pointed at different worlds. Conversely, one system may be lawfully employed for a week in one place and unlawfully for two minutes in another. **That variability is not a defect of the framework. It is why no single figure could ever have been written into a treaty, and why “limit duration” without more was never going to be operative.**

17.2 That commanders already estimate it

The framework asks for something the profession does informally and continuously.

A targeting cell treating intelligence as perishable is estimating a fact-decay rate. A rule of engagement requiring re-observation before a strike on a mobile target is a fact-decay judgement expressed as a procedure. A decision to hold a dynamic target for positive identification rather than strike on a two-hour-old report is the same judgement made in the moment. The concept is not exotic and does not need to be taught. What it lacks is a name, a place in the legal analysis, and any requirement that it be written down.

17.3 What the rule would say

Stated as an element, and this is the form Part VIII drafts:

An autonomous engagement is permitted only where the interval between the last human decision bearing on it and its effect is shorter than the period over which the facts underpinning that decision may reasonably be expected to hold; and where it is not, only if the system can be cancelled or suspended within that period.

The second limb matters as much as the first and is often the operative one. It is not a new obligation: it is Article 57(2)(b) applied to a system that will still be running when the obligation bites.⁵⁸ Its practical content is that **a system employed across an interval longer than the fact-decay rate must be recallable within that rate**—a design requirement with a legal source, and one the companion engineering paper shows can be met at negligible cost even in a system deliberately built to be unreachable.⁵⁹

18. What must be recorded

A rule is only as good as the artefact it produces, and Part III §8.3 identified the absence of an artefact as the reason “context-appropriate” cannot be assessed. The answer is a short record made before employment.

Four items, and the brevity is deliberate, because a burdensome requirement will not be adopted and will not be completed honestly under operational pressure:

One—the interval. The maximum period, from the authorising decision, across which the system may deliver an effect without further human input.

Two—the decay estimate. The period over which the facts underpinning the authorisation are expected to hold, with the principal propositions named. Not a probability distribution; a stated expectation and the reasons for it.

Three—the cancellation capability. Whether the system can be cancelled or suspended within the estimated period, and by what means.

⁵⁸Additional Protocol I (n 35), art 57(2)(b).

⁵⁹See the companion paper on closed-system ground robotics, in which the recall channel is priced at approximately one part in a thousand of the platform’s energy budget and the emission it entails is borne by the control station rather than by the emplaced system.

Four—the information state. What the authorising human knew about the environment at the time, and the vintage of that knowledge.

Three features of that record do the work.

It is **short enough to be produced**. A commander asked to record four items has been asked for something achievable; a commander asked to demonstrate context-appropriate control has been asked for nothing in particular.

It **discloses the reasoning rather than the conclusion**, which is what makes later assessment possible. A record stating that the facts were expected to hold for six hours, that the system was employed across seventy-two, and that it could not be cancelled, establishes a breach on its face without anyone having to determine what “appropriate” means.

And it is **assessable without hindsight bias**, because what is assessed is the estimate rather than the outcome. A decay estimate reasonable when made is not falsified by an unlucky result, and one plainly unreasonable is not saved by a fortunate one. That is the ordinary structure of legal assessment of targeting decisions, and the framework fits it rather than replacing it.

19. Objections

19.1 That the decay rate cannot be estimated

This is the objection the paper takes most seriously, because if the rate cannot be estimated to a standard capable of supporting later assessment, the framework supplies a form to be filled in and nothing else.

Four replies, in ascending strength.

It is already estimated, as §17.2 sets out. The objection must therefore be that estimates made routinely for operational purposes are too poor to bear legal weight, which is a different and much larger claim about targeting practice generally.

The standard is reasonableness, not accuracy. The framework does not ask whether the facts in fact held. It asks whether the estimate of how long they would hold was reasonable on the information available. That is the standard applied to every other targeting judgement and it is no more demanding here.

Uncertainty is handled by the second limb. A commander unable to estimate the decay rate with confidence is not licensed to proceed; they are directed to the cancellation requirement at §17.3, which is the correct answer to uncertainty and what a careful commander would do in any event.

And an unreasonable estimate is easier to identify than an unreasonable degree of control. The framework does not require an assessor to determine the correct decay rate. It requires them to determine whether the recorded estimate was one a reasonable commander could have made—a comparative exercise on a stated quantity, which is tractable, against an exercise on an unstated quality, which is not.

19.2 That it privileges time over the substance of the decision

The objection is that a framework organised around an interval will license bad decisions made quickly and prohibit good ones made slowly.

It would, if it replaced the substantive rules. It does not, and on that point this paper is with those who hold that the existing law of armed conflict already governs these systems adequately in principle.⁶⁰ The framework adds an interpretive content to one existing obligation rather than a new obligation. Distinction, proportionality and precautions apply in full and the framework adds a further condition. An engagement authorised on a defective proportionality assessment is unlawful whatever the interval; the framework asks the additional question whether a sound assessment had ceased to be sound by the time the effect arrived. A rule adding a necessary condition cannot license anything.

⁶⁰Schmitt and Thurnher (n 12).

19.3 That states will set the interval to whatever suits them

They will set it to whatever they can defend, which is a different thing and is how every reasonableness standard operates.

The constraint is that the estimate is recorded before the outcome is known. A state recording a seventy-two-hour decay estimate for an urban engagement has created a document that will not survive examination, and it has created it in advance. That is materially different from the present position, in which nothing is recorded and no examination is possible.

19.4 That it does nothing about anti-personnel systems

Correct, and the concession should be plain. The particular difficulties of systems targeting people—the danger of discrimination arising from bias in system design, and the objection that reducing a person to a target signature is wrong independently of accuracy—are not addressed by this framework.⁶¹ They require their own treatment, the rolling text gestures at them through the restriction to objects that are military objectives by nature, and this paper does not claim to have resolved them.

19.5 That it is a design constraint dressed as a legal rule

Partly true, and it is a feature. The second limb at §17.3 is a legal obligation with a direct engineering consequence: systems intended for employment across intervals longer than the local decay rate must be recallable within it. That is what a workable rule in this area should look like, because a rule with no design consequence changes nothing about what gets built.

Part VI—Three Cases

20. Why these three

This section applies the framework to the cases that decide whether it is useful. A framework which cannot distinguish systems that plainly differ is a restatement; one which distinguishes them on a stated ground, and explains a difference the prevailing analysis cannot, is doing work.

The three are chosen because they occupy widely separated positions on the interval axis while sitting close together on the autonomy axis. All three exist; none is a projection. Crootof made the point a decade ago that autonomous weapon systems are not weapons of the future but are already integrated into states' armed forces, and the systems examined below bear that out.⁶² Each identifies, selects and engages without human intervention in the execution of those tasks, so each falls within the rolling text's characterisation.⁶³ If autonomy were the governing variable they would be governed alike. They are not, and this Part explains why.

21. Loitering munitions, where the interval is the design feature

21.1 The Harpy and the Harop, which are a natural experiment

Israel Aerospace Industries' Harpy entered service in 1989 for the suppression of enemy air defences, homing on radar emissions.⁶⁴ It is described as a fire-and-forget autonomous weapon, and the later Harpy NG retains

⁶¹Article 36 (n 8), on the specific challenges posed by anti-personnel systems and on harmful bias.

⁶²Rebecca Crootof, 'The Killer Robots Are Here: Legal and Policy Implications' (2015) 36 *Cardozo Law Review* 1837. On the control question specifically see also Rebecca Crootof, 'A Meaningful Floor for "Meaningful Human Control"' (2016) 30 *Temple International and Comparative Law Journal* 53.

⁶³Rolling text (n 24), part I para 1.

⁶⁴The system entered service in 1989 for suppression of enemy air defences, homing on radar emissions; published figures give 2.7 m length, 2.1 m wingspan, 135 kg gross weight, maximum speed 185 km/h and operational range 200 km. See 'IAI Harpy' https://en.wikipedia.org/wiki/IAI_Harpy accessed 14 August 2026, and the descriptions collected at Automated Decision Research, 'Israel Aerospace Industries HAROP loitering munition' <https://automatedresearch.org/weapon/israel-aerospace-industries-harop-loitering-munition/> accessed 14 August 2026. These are manufacturer and open-source figures, not independently verified, and the argument in the text does not depend on their precision.

that character with an airborne endurance of nine hours.⁶⁵

The Harop is the same manufacturer’s development of the same airframe lineage, with comparable endurance of nine hours and a 23 kg warhead.⁶⁶ It differs in one respect that matters here. It operates in a **man-in-the-loop** mode with a remote mission operator, and **it can be aborted and re-attacked**—a capability the published descriptions characterise as a safety feature for complex urban or coalition operations.⁶⁷

Set the two side by side. Same manufacturer, same class, same endurance, same broad function, both answering to the rolling text’s characterisation. The difference is not the degree of autonomy in the abstract. It is that one can be recalled within its loiter and one cannot.

That difference is the second limb of the rule at §17.3, and industry arrived at it commercially before the law arrived at it doctrinally. The stated justification for the abort capability is the environment: urban and coalition operations. Those are precisely the environments in which the facts underpinning a targeting judgement decay fastest—population moves, friendly forces are present, and the identity of what is below changes within minutes. A manufacturer offering abort for urban use, and not troubling to offer it for open-country anti-radar use, has priced the fact-decay rate. It has not called it that.

21.2 The finding the range figures produce

The Harop’s published specifications contain a second and sharper illustration, and it is quantitative.

Its communications range is given as 200 km. Its operational range is given as 1,000 km.⁶⁸

The consequence is that **the same airframe, on the same sortie, passes through a threshold at which its legal character changes.** Within 200 km it is a system whose engagement can be cancelled or suspended by the operator, and Article 57(2)(b) can be complied with in the ordinary way.⁶⁹ Beyond 200 km it is a system that cannot be reached, whose authorising judgement must hold unaided for the remainder of the sortie, and in respect of which the obligation to cancel or suspend has become impossible to discharge by any means the operator possesses.

Nothing about the machine’s autonomy changes at that boundary. The software is the same, the seeker is the same, the release authority granted before launch is the same. What changes is the interval, and with it whether a continuing legal obligation can be performed.

Three points follow, and the third is the one to carry forward.

The autonomy framing cannot see this. A characterisation asking whether the system identifies, selects and engages without human intervention returns the same answer at 150 km and at 900 km. The rolling text’s four limitation parameters do not help either: target type, geographical scope and scale are unchanged, and duration as a bare dial does not distinguish nine hours within communications range from nine hours outside it.

The interval framing sees it immediately, and produces an operational rule a targeting cell could apply tomorrow: employment beyond communications range is employment at an unbounded interval, and is permissible only where the fact-decay estimate for the target area exceeds the remaining time of flight.

And it identifies the design remedy. The gap between 200 km and 1,000 km is not a law of physics but a consequence of the communications architecture chosen. A receive-only recall channel—which requires the munition to listen but never to transmit, and therefore costs it nothing in detectability, since direction-finding

⁶⁵IAI Harpy NG’ https://en.wikipedia.org/wiki/IAI_Harpy_NG accessed 14 August 2026, giving nine hours’ airborne time and describing the system as fire-and-forget.

⁶⁶IAI Harop’ https://en.wikipedia.org/wiki/IAI_Harop accessed 14 August 2026: nine hours’ endurance, 23 kg warhead, 135 kg gross weight, 400 km/h, ceiling 4,600 m.

⁶⁷ibid; and Automated Decision Research (n 64). The abort-and-re-attack capability, and its characterisation as a safety feature for complex urban or coalition operations, are drawn from the published descriptions. The inference in the text that this amounts to commercial pricing of the fact-decay rate is the author’s.

⁶⁸ibid: communications range 200 km, operational range 1,000 km.

⁶⁹Additional Protocol I (n 35), art 57(2)(b).

locates transmitters—is among the cheapest capabilities available in a system of this kind.⁷⁰ A rule expressed in terms of the interval tells a manufacturer what to build. A rule expressed in terms of context-appropriate control does not.

21.3 What this does not establish

The paper does not assert that the Harpy is unlawful or that the Harop is lawful. Both are capable of lawful and unlawful employment, and which it is depends on the circumstances of the particular use, the target set, and the judgement made before launch. What the comparison establishes is that **the legally salient difference between them is the interval and the recall capability, and that this is visible on the published descriptions without any need to resolve what autonomy means.**

22. Automated air and missile defence, where the interval is compressed and the practice is old

The second case is the one usually raised against restrictive proposals, and it is a fair challenge the framework must answer rather than evade.

Automated terminal defences against incoming rockets, artillery, mortars and missiles operate on timelines shorter than human reaction. A system engaging a supersonic sea-skimming missile, or a rocket in terminal descent, cannot pause for authorisation, and states have operated such systems for decades without serious suggestion that doing so is unlawful. Ukraine’s Sky Sentinel is a recent instance of the same logic applied to drone defence: radar-cued, engaging targets closing at speeds up to 500 mph, operating without a soldier working the turret.⁷¹

The interval framework explains why these are unobjectionable, and does so without needing an exception.

The interval is short—often seconds. The facts on which the authorising judgement rests are, in the paradigm case, still true when the effect occurs, because almost no time has passed.

The decay rate is low relative to the interval, not merely absolutely. This is the point that matters, and it is why the analysis does not reduce to “short is fine”. The propositions underpinning the authorisation are of an unusually durable kind: that this is a defended installation, that the airspace above it is closed, that objects approaching at that speed and profile on that bearing are hostile. Those hold across the hours or days for which the system is switched on, not merely across the seconds of an engagement.

And the human decision is genuinely present, at the right level. The authorising judgement is the decision to activate the system in a defined mode, in a defined place, against a defined class of incoming object, for a defined period. That is a decision by a human on facts they knew, and it remains valid across the period it covers. Ekelhof’s distributed account describes this exactly: control resides at a point less proximal than the engagement, and it is real control because the facts have not moved.⁷²

The framework therefore yields the conventional answer for the conventional case, which is a necessary condition for taking it seriously. More usefully, it identifies what would change the analysis. **If the defended area ceases to be one in which the durable propositions hold—if civilians enter the engagement envelope, if friendly aircraft operate in it, if the system is moved somewhere the airspace is not closed—the decay rate rises and the same system at the same settings becomes unlawful without anything about it having changed.** That is a real and operationally familiar risk, it is the mechanism by which such systems have caused fratricide and civilian harm, and no analysis conducted in terms of autonomy captures it.

⁷⁰See the companion paper on closed-system ground robotics, on receive-only emission postures and the asymmetry that direction-finding locates transmitters rather than receivers.

⁷¹ibid, on Ukraine’s Sky Sentinel turret and the reporting there collected.

⁷²Ekelhof (n 46).

23. Emplaced offline ground systems, where the interval runs to weeks

The third case is the hardest, it is not hypothetical, and the companion engineering paper specifies it in detail.⁷³

A ground platform emplaced to deny an approach for thirty days, operating on local sensing and on-platform inference, at an emission posture chosen so that it cannot be found, presents an interval three to four orders of magnitude longer than a loitering munition's and six orders longer than a terminal defence engagement. A machine gun on a tracked chassis is reported to have held a fighting position on the Ukrainian front for more than a month.⁷⁴

The framework's treatment is severe, and it should be.

The interval is the emplacement. Thirty days is the period across which the authorising judgement must hold, unaided, in contested ground where the pattern of life is exactly what an adversary is attempting to change.

The decay rate in such an environment is high, and the propositions concerned are of the fragile kind rather than the durable kind at §22. That no non-combatant will enter this sector is a proposition about a contested area over a month, not a proposition about closed airspace over a defended installation.

So the first limb fails and everything turns on the second. An emplacement at a thirty-day interval, in an environment whose facts decay in days or hours, is permissible under §17.3 only if the system can be cancelled or suspended within the decay period. The legal question is therefore not whether such a system may exist but whether it can be reached—and that is a question about communications architecture, which is a design choice and not a fact of nature.

Here the engineering and the law converge in a way worth stating plainly, because it is the strongest practical argument in either paper. A receive-only recall channel costs such a system approximately one part in a thousand of its energy budget, requires it to emit nothing, and places the associated emission on the control station to the rear rather than on the emplaced platform.⁷⁵ **The capability that makes a thirty-day emplacement lawful is among the cheapest things in it.** A state fielding such a system without a recall channel is not economising on anything material; it is declining a capability that costs almost nothing and that is the difference between an employment able to comply with Article 57(2)(b) and one that cannot.

24. What the three cases establish

Case	Interval	Decay rate in the paradigm environment	Recall available	Result under §17.3
Terminal air defence	Seconds	Low across the activation period	Not needed at that interval	Permissible; the conventional answer, reached without an exception
Loitering munition within communications range	Hours	Variable; high in urban use	Yes	Permissible where the estimate is honest; abort is the operative safeguard

⁷³See the companion paper on closed-system ground robotics, specifying a thirty-day armed emplacement.

⁷⁴ibid, citing Ukrainian reporting of a machine-gun platform holding a fighting position for over a month. The figure is reported by a party to the conflict about its own performance.

⁷⁵ibid, on the tiered receiver and the placement of the recall emission at the control station.

Case	Interval	Decay rate in the paradigm environment	Recall available	Result under §17.3
Loitering munition beyond communications range	Hours, unbounded	Same	No	Permissible only where the decay estimate exceeds remaining flight
Emplaced ground system	Weeks	High in contested ground	Only if designed in	Impermissible without a recall channel

Three observations.

The framework reaches the conventional answer where there is one, which is the test any proposed rule must pass before its harder applications are worth considering.

It distinguishes cases the autonomy framing cannot, most sharply within a single airframe on a single sortie at §21.2, where the legal character changes at a range threshold while the autonomy does not change at all.

And in every case the operative question is the same and is answerable: how long must the authorising judgement hold, how quickly do the facts underpinning it decay, and can the system be reached if they decay faster than expected. Three questions with answers, in place of one question—what is context-appropriate—with none.

Part VII—Responsibility

25. The gap as it is usually stated

This section aims to establish that the responsibility gap is narrower than the literature asserts, and that what narrows it is not a philosophical answer but a procedural one: giving the responsible human something specific to be responsible for.

Matthias framed the difficulty first, in terms of learning automata whose behaviour is not fully determined by their programmers, and argued that a gap opens in which nobody can properly be held responsible for what such a machine does.⁷⁶ Sparrow applied the argument to weapons and pressed it to a conclusion. Asking who should be held responsible when an autonomous system is involved in an atrocity, he considered the designers, the commanding officer who ordered its use, and the machine itself, and found each inadequate: the system is independent enough that its makers may disclaim its choices, yet too unlike us to be a fit subject of blame. Because he treats it as a necessary condition of fighting justly that someone can justly be held responsible for the deaths caused, he concludes that deploying such systems would be unethical.⁷⁷

The argument has been contested. Responses have questioned whether the dilemma is genuine, whether responsibility must attach in the way Sparrow requires, and whether the gap is a feature of automated warfare or an artefact of how the question is posed.⁷⁸ This paper does not resolve that debate and does not need to. Its contribution is narrower and, for that reason, more usable.

⁷⁶Andreas Matthias, ‘The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata’ (2004) 6 *Ethics and Information Technology* 175.

⁷⁷Robert Sparrow, ‘Killer Robots’ (2007) 24 *Journal of Applied Philosophy* 62.

⁷⁸Among the replies see Marc Champagne and Ryan Tonkens, ‘Bridging the Responsibility Gap in Automated Warfare’ (2015) 28 *Philosophy & Technology* 125, and Daniel W Tigard, ‘There Is No Techno-Responsibility Gap’ (2021) 34 *Philosophy & Technology* 589. The debate is not resolved here, and the argument in the text does not depend on which side of it is correct.

26. What the interval framework does to the gap

Sparrow's dilemma is generated by a particular question: **who is responsible for what the machine did?** Posed that way it is hard, because the machine's conduct is the product of design decisions, training, deployment context and circumstance, and no single person authored it.

The interval framework does not ask that question. It asks a different one:

Was the estimate of how long the authorising judgement would remain good a reasonable one to make on the information available?

That question has a determinate subject—the human who made the authorising decision—and a determinate object—an estimate they made, recorded, and which exists afterwards as a document. Four consequences follow, and the third is the one the literature has not drawn.

The responsible party is identified without any inquiry into the machine. Whoever authorised the employment made the decay estimate. That person is not being held responsible for the system's behaviour, which they did not author, but for a judgement they did make: how long the facts would hold. Commanders are already responsible for judgements of exactly that character.

The object of responsibility is an ordinary military judgement, not a technical one. Estimating how quickly a situation will change requires no understanding of a model's internals. It requires knowledge of the ground, the population, the tempo and the enemy—the commander's own expertise rather than the engineer's. This matters, because much of the responsibility-gap literature assumes that holding a commander responsible would require them to understand a system they cannot understand. On this framing it does not.

And the gap shrinks not because the philosophical problem has been solved but because the question has been made answerable. Sparrow asks whether anyone can be responsible for an outcome nobody authored. The framework asks whether a particular person made a particular estimate reasonably. The first may well be unanswerable. The second is the kind of question courts martial, boards of inquiry and tribunals answer routinely, on the same reasonableness standard applied to every other targeting judgement.

What remains outside is stated rather than absorbed. Where the system does something no reasonable decay estimate would have anticipated—a failure of the model rather than a change in the world—the framework does not assign responsibility to the commander, and it should not. That residue is real, it is Sparrow's case at its strongest, and it belongs with the design, testing and review obligations rather than with the employment decision. **The framework does not close the gap. It confines it to the cases that actually generate it,** which is a smaller set than the literature's framing implies.

27. How this maps onto the existing rules

The framework needs no new doctrine of responsibility. It attaches to three that exist.

Individual responsibility for the attack decision. The obligations in Article 57 fall on those who plan or decide upon an attack.⁷⁹ A decay estimate is part of deciding upon an attack, and a commander who records an estimate no reasonable commander could have made, and whose system then causes unlawful harm across an interval that estimate was used to justify, has failed the precautionary obligation in an identifiable way.

Command responsibility. Where a subordinate commits a breach, a superior is not absolved if they knew, or had information which should have enabled them to conclude, that the breach was being or was about to be committed, and did not take feasible measures within their power to prevent or repress it.⁸⁰ Commanders are further required to prevent and, where necessary, to suppress and report breaches.⁸¹ The recording requirement at §18 bears directly here, because it produces the very information by reference to which such knowledge is later assessed. **A superior who saw a seventy-two-hour interval recorded against a**

⁷⁹Additional Protocol I (n 35), art 57(2).

⁸⁰ibid, art 86(2).

⁸¹ibid, art 87.

six-hour decay estimate had information enabling them to conclude. That is a materially different evidential position from the present one, in which no such document exists.

State responsibility. The conduct of organs of a state is attributable to it, including where an organ exceeds its authority or contravenes instructions, and a state is responsible for its internationally wrongful acts.⁸² Nothing in autonomy disturbs that. A state employing a system across an interval exceeding the decay period, without a recall capability, has through its organs made a choice for which it answers; the absence of a human at the moment of effect creates no attribution difficulty, because **the attributable act is the employment decision, and a human took it.**

28. The objection that this lets the commander off

The objection deserves a direct answer: that a commander who records a plausible-looking estimate and is then unlucky escapes, so the framework weakens responsibility rather than strengthening it.

Three replies.

The comparison is with the present position, not with an ideal one. At present nothing is recorded, no estimate is made explicit, and there is nothing against which conduct can be assessed. A commander who records nothing cannot be shown to have been unreasonable about anything. The framework creates the document that makes assessment possible, and creates it before the outcome is known.

Reasonableness is not a low standard when the estimate is contemporaneous and specific. A recorded expectation that the facts would hold for a stated period, with the principal propositions named, is a hostage to fortune in a way that a general assertion of context-appropriate control is not. It can be compared against what the intelligence actually said, against what the commander was told, and against what other commanders estimated in comparable circumstances.

And the second limb does not depend on the estimate at all. Where the interval exceeds the decay period the system must be recallable. That is a hard requirement about a capability rather than a soft one about a judgement, and no quality of estimate satisfies it. A state fielding an unrecallable system across a long interval in a fast-changing environment breaches the rule whatever it wrote down.

Part VIII—The Element, Drafted

29. What the drafting has to satisfy

Part II established four constraints on anything offered into the 2026 window: it must be statable as a single element rather than bargained as a package; it must not require a state to abandon a position held for a decade; it must be assessable against something other than agreement; and it must do work whether or not an instrument follows.⁸³ This section tests the proposal against the positions of the states that matter, then drafts it.

30. What it would cost the principal states to accept

30.1 China, whose position determines an instrument's reach

China has called for a treaty prohibiting the **use** of lethal autonomous weapons systems, while not seeking to prohibit their development or production.⁸⁴ Its framework is two-tiered: outright prohibition of systems

⁸²International Law Commission, *Articles on Responsibility of States for Internationally Wrongful Acts* (2001) arts 4 and 7, annexed to UNGA Res 56/83 (12 December 2001) UN Doc A/RES/56/83; and Additional Protocol I (n 35), art 91, on liability to pay compensation.

⁸³Part II §7.1.

⁸⁴Human Rights Watch, *Stopping Killer Robots: Country Positions on Banning Fully Autonomous Weapons and Retaining Human Control* (10 August 2020) <https://www.hrw.org/report/2020/08/10/stopping-killer-robots/country-positions-banning-fully-autonomous-weapons-and> accessed 14 August 2026. The characterisation of China's position as directed at use rather than development is that of the reporting organisation, and is consistent with Working Paper of the People's Republic of China on Lethal Autonomous Weapons Systems (2022) <https://documents.unoda.org/wp-content/uploads/2022/07/Working-Paper->

operating without meaningful human involvement, and regulation of the remainder to secure appropriate human involvement.⁸⁵ Its 2022 working paper defines the prohibited category by cumulative traits, among them lethality and full autonomy without human intervention.⁸⁶ It regards the Convention as the appropriate forum and supports negotiation of a legally binding instrument **when the conditions are mature**.⁸⁷

Three observations, and the first is the most consequential for this paper.

A rule about the interval is a rule about use. It does not restrict development, production, testing or possession. It attaches at the point of employment, to a decision a commander makes, and it is indifferent to what a state builds. That is the axis on which China has staked its position, and the proposal sits on the permitted side of it. **A state whose stated objective is to prohibit use while preserving development can accept an interval rule without moving.** No definitional prohibition has that property, because a definitional prohibition necessarily reaches the systems themselves.

The cumulative-traits approach illustrates the definitional route’s weakness rather than solving it. Defining a prohibited class by a conjunction of characteristics means a system escapes by failing any one of them. A prohibition constructed that way is narrow by construction, and its narrowness is a consequence of the method rather than an accident of drafting. The interval framework defines no class, which is why nothing escapes it by falling outside one.

And “when the conditions are mature” is Part II’s argument arriving from another direction. It is a formula that never expires and cannot be falsified. A state may support a binding instrument in principle, indefinitely, at no cost, and the mandate contains nothing capable of testing whether conditions have matured. That is not a criticism of China in particular. It describes what any state may rationally do in a body constituted as this one is.

30.2 The United States

United States policy is set by Department of Defense Directive 3000.09, issued in November 2012 and updated in January 2023.⁸⁸ The 2023 revision maintains the structure and substance of its predecessor rather than departing from it.⁸⁹

The Directive’s architecture is a senior review requirement for certain systems, together with a design standard directed at allowing commanders and operators to exercise appropriate levels of human judgement over the use of force. The proposal here is compatible with that architecture and adds one thing: a **contemporaneous employment-stage record**, where the Directive’s review occurs at development and fielding. In the terms of Part IV, the Directive is a developed instance of the first occasion, and the framework supplies the second.⁹⁰

The proposal therefore asks the United States for a procedural addition to a system it already operates, rather than for a concession of position. That is a low price and it is deliberate.

30.3 The European position and the United Kingdom’s absence

France, Germany and more than forty other states have said the draft rules are ready to negotiate.⁹¹ For those states an interval rule adds content to a text they already support, and the constraint at §29 that it

of-the-Peoples-Republic-of-China-on-Lethal-Autonomous-Weapons-Systems%EF%BC%88English%EF%BC%89.pdf accessed 14 August 2026.

⁸⁵Lieber Institute West Point, ‘Human Oversight with Chinese Characteristics: Lethal Autonomous Weapons in the CCW GGE’ <https://lieber.westpoint.edu/human-oversight-chinese-characteristics-lethal-autonomous-weapons-ccw-gge/> accessed 14 August 2026.

⁸⁶Working Paper of the People’s Republic of China (n 84).

⁸⁷ibid.

⁸⁸US Department of Defense, Directive 3000.09, *Autonomy in Weapon Systems* (21 November 2012, updated 25 January 2023) <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf> accessed 14 August 2026.

⁸⁹On the 2023 revision maintaining rather than departing from the 2012 structure, see Human Rights Watch, ‘Review of the 2023 US Policy on Autonomy in Weapons Systems’ (14 February 2023) <https://www.hrw.org/news/2023/02/14/review-2023-us-policy-autonomy-weapons-systems> accessed 14 August 2026.

⁹⁰Part IV §14.

⁹¹Joint statement (n 4); and Article 36 (n 23).

be statable as a single element is what allows it to be added without reopening anything.

The United Kingdom neither joined that statement nor opposed it.⁹² The framework offers a state in that position something specific: an element adoptable domestically, within existing review and targeting procedures, without waiting for the multilateral process and without taking a position on the definitional question it has so far declined to answer. Part II §7.1's fourth constraint exists for this case.

30.4 What the framework does not buy

It does not resolve the disagreement between states seeking prohibition and states seeking regulation, and is not intended to. Its claim is to be **orthogonal** to that disagreement: a state on either side can accept it without conceding anything. That is what makes it adoptable in a consensus body, and it is also the limit on how much it achieves.

31. The element

Drafted as it would have to be to enter the rolling text. It is short because §29 requires it to be.

Element: temporal validity of the authorising judgement

1. *Definitions.* For the purposes of this element—

- (a) “authorising decision” means the last decision by a human being to employ a weapon system in circumstances in which it may select and engage a target, whether taken at the point of engagement or in advance of it;
- (b) “interval” means the period between the authorising decision and the effect of an engagement resulting from it;
- (c) “decay period” means the period over which the factual circumstances on which the authorising decision was based may reasonably be expected to continue to obtain.

2. *The rule.* A weapon system may be employed in circumstances in which it may select and engage a target only where—

- (a) the interval does not exceed the decay period; or
- (b) where the interval exceeds the decay period, the system can be cancelled or suspended within the decay period by those who plan or decide upon the attack.

3. *Record.* Before employment, the authorising decision shall be recorded together with—

- (a) the maximum interval for which the system may operate without further human decision;
- (b) the decay period, and the principal factual propositions to which it relates;
- (c) whether the system can be cancelled or suspended within that period, and by what means; and
- (d) the information available to the deciding person concerning the environment of employment, and its currency.

4. *Standard.* The decay period is to be assessed by reference to what was reasonable on the information available at the time of the authorising decision, and not by reference to events subsequent to it.

5. *Preservation.* Nothing in this element derogates from any obligation under Articles 48, 51, 52, 57 or 58 of Additional Protocol I or under customary international humanitarian law, and the obligations are cumulative.

⁹²Article 36 (n 23).

32. Notes on the drafting

Paragraph 1(a) locates the decision rather than the system. It says “the last decision by a human being”, which makes the element indifferent to whether the functions are integrated in one unit or dispersed across several—the difficulty Part III §8.1 identified in the rolling text’s characterisation.⁹³ It also accommodates the distributed account of control without adopting it as a matter of definition: wherever along the chain the last human decision falls, that is where the interval starts.⁹⁴

Paragraph 2 is the operative rule and it is disjunctive. A state may comply either by keeping the interval short or by retaining the ability to intervene. That is deliberate. It does not prohibit long intervals, which would be unacceptable to several states and wrong in cases such as terminal air defence, and it does not require continuous supervision, which is unattainable.

Paragraph 3 is the innovation and the burden. It is what makes the rule assessable, and it is held to four items so that it can be completed honestly under operational conditions. A longer form would be a worse rule.

Paragraph 4 forecloses hindsight. Without it the element would be assessed on outcomes, which would be unfair to commanders and would create an incentive to record nothing.

Paragraph 5 makes the obligation cumulative rather than substitutive. Its function is to prevent the element being read as a safe harbour: a state complying with the interval rule has not thereby satisfied distinction or proportionality. The drafting says so expressly, because the objection at Part V §19.2 would otherwise have force.

33. Adoption without an instrument

Part II found that the multilateral procedure cannot terminate. The element is therefore drafted so as not to depend on it.

It can be adopted **as a national standard**, by inserting paragraph 3’s record into targeting doctrine and paragraph 2’s rule into rules of engagement. It can be adopted **within an Article 36 review procedure**, as the second occasion Part IV §14 argues for, with the review certifying not merely that a system is capable of lawful use but the intervals and environments in which it is. It can be adopted **as a coalition standard**, which is the form in which most operational law actually converges. And it can be adopted **as an element of the rolling text**, which is the form drafted above.

Each does real work, and the first three require nobody’s agreement. That is the practical answer to a procedure with no exit: draft the rule so that the exit is not needed.

Part IX—Conclusions

34. Conclusions on each element

The mandate is the cause and the finding is testable. The Group was instructed to consider and formulate, by consensus, a set of elements of an instrument without prejudging its nature, and other possible measures. Five features of that instruction together remove every exit: the verbs are not “negotiate”; consensus operates twice, on the elements and on the report; the object is elements rather than an instrument; the bar on prejudging the nature prevents the text being tested against any purpose; and the disjunctive limb supplies a terminus that is not an instrument at all. The mandate’s own preferred completion date of end-2025 has passed. The claim fails if the Group produces an instrument, or agrees a negotiating mandate, before November 2026.

The operative formula is unverifiable. “Context-appropriate human judgement and control” has no determining body, names no dimension along which appropriateness is measured, and leaves no artefact

⁹³Part III §8.1, and the observation on dispersal at Article 36 (n 8).

⁹⁴Ekelhof (n 46).

from which compliance could be established. A prohibition that cannot be breached is absent rather than weak—and the December 2025 status no longer drafts the formula as a prohibition at all. That is not the fault of the states that agreed it: under a mandate forbidding any decision about the instrument’s nature, language committing nobody is what the procedure selects for.

Article 36’s defect is an axis error and the remedy follows from that. The provision’s temporal reach is broad across study, development, acquisition and adoption—all four pre-employment—and zero across operational life. For weapons whose properties are fixed at manufacture that is exactly right. For systems whose behaviour depends on data encountered after adoption, or simply on an environment they were not reviewed against, the object employed is not the object reviewed. The answer is a second occasion rather than a wider scope, and it requires no amendment: precautions in attack already attach to the attack rather than the weapon, and the obligation to cancel or suspend is already continuous. **It is the one existing rule that looks along the right axis.**

The two halves were already in the debate and had not been joined. Duration sits inside the rolling text as one component of its scale-of-operation measure with no account of what sets it. Understanding of a system’s effects in the context of use sits outside as the principal gap its critics identify. Joined they are one variable—the interval between the last human decision and the effect, together with what was knowable across it—and it is not a fourth parameter but the one on which the others depend for their meaning, because every substantive rule of targeting is a rule about a moment and the legal question is whether the facts of that moment survived to the effect.

The fact-decay rate is the operative content and it is a property of the situation. Two systems with identical autonomy are governed differently because they are pointed at different worlds; one system is lawful for a week in one place and unlawful for two minutes in another. No figure could have been written into a treaty, which is why “limit duration” without more was never going to be operative. Commanders already make this estimate; what it has lacked is a name, a place in the legal analysis, and any requirement that it be written down.

The three cases separate on the interval and not on autonomy. Terminal air defence is unobjectionable without needing an exception, because the decay rate is low relative to the interval and the authorising decision remains valid across the period it covers. The Harpy and the Harop are a natural experiment: same manufacturer, same nine-hour endurance, same characterisation, and the difference is that one can be aborted. The abort capability is justified in the published descriptions by reference to urban and coalition operations—which is the fact-decay rate priced commercially before it was named legally. And the Harop’s 200 km communications range against 1,000 km operational range means the same airframe on the same sortie crosses a threshold at which its legal character changes while its autonomy does not.

The responsibility gap narrows because the question becomes answerable. Sparrow asks who is responsible for what a machine did, which may be unanswerable. The framework asks whether a named person made a recorded estimate reasonably, which is the kind of question tribunals answer routinely. The residue—where a system does what no reasonable estimate could have anticipated—is real, is Sparrow’s case at its strongest, and belongs with design and review rather than with the employment decision. The gap is confined rather than closed.

And the element is drafted to be adoptable by a body that cannot negotiate. It is a rule about *use*, which is the axis on which China has staked its position, so a state seeking to prohibit use while preserving development can accept it without moving. It is a procedural addition to the architecture the United States already operates. It adds content to a text France, Germany and forty others already support. And it gives a state that has taken no position something it can adopt unilaterally.

35. What would change the analysis

An instrument, or a negotiating mandate, before November 2026. This would defeat Part II directly and the rest of the paper would need rereading in that light.

Evidence of existing employment-stage determination in state practice. Review procedures are not published. If states already record something equivalent to §18’s four items, the proposal is a generalisation

of practice rather than an addition to it, and the author would prefer that.

A demonstration that decay estimates cannot be made to a standard supporting later assessment. This is the objection the paper takes most seriously. It is answered at Part V §19.1 by reference to the reasonableness standard and to the second limb, and the answer is arguable rather than conclusive.

A treatment of anti-personnel systems. The framework does not reach the objection that reducing a person to a target signature is wrong independently of accuracy, and does not pretend to. That requires separate work.

36. Summative remarks

Ten years of argument about what an autonomous weapon *is* have produced a rolling text whose central formula cannot be breached, under a mandate that cannot terminate, reporting to a Review Conference in November 2026.

The argument of this paper is that the question was mis-set. What the law of targeting requires is not that a human pull a trigger but that the facts a targeting judgement rests on should still be true when the effect arrives. That requirement is temporal, it is already latent in the obligation to cancel or suspend an attack, and it can be stated in a form a commander can comply with, a superior can supervise and a tribunal can assess—which is more than can be said for the formula it would replace.

Both halves of the necessary variable were already in the room. One sat inside the rolling text as an unexplained parameter and the other sat outside it as an acknowledged gap, and nobody had put them together. Joined, they price the thing the debate has been unable to name: how long an authorisation is good for, and what has to be true for it to remain so.

The element drafted at Part VIII is short, it costs the principal states very little, and three of its four routes to adoption need nobody's permission. That last property is not modesty. It is the only sensible response to a procedure which, on the evidence of its own instruction, was never going to finish.

What to do with it is equally short. A delegation that wants the second session to produce something assessable can table the element as a working paper: it stands as a single element, and it asks no state to abandon a position. A ministry that prefers not to wait can put paragraph 3's four-line record into targeting doctrine, where it begins doing work the first time an employment is authorised against it. And a reader who takes nothing else should take the question the framework makes askable of any system, any formula and any future text: **how long did the authorisation have to hold, how fast did the facts decay, and could the system be reached when they did.**

A companion engineering paper specifies the emplaced ground system treated at Part VI §23, and prices the recall capability on which its lawfulness depends at approximately one part in a thousand of the platform's energy budget.

Appendix—Worked Records

The rule proposed at Part VIII §31 stands or falls on whether paragraph 3's record can be completed honestly, briefly, and under operational conditions. Assertion is not enough. This Appendix works the record for each of the three cases at Part VI, so that a reader can judge the burden rather than take it on trust.

The entries are the author's construction. They are not drawn from any state's practice, no operational document of this kind is public, and they should be read as illustrations of form rather than as evidence about content.

A.1 Terminal air and drone defence

Interval. Seconds from cue to engagement. The authorising decision is the activation of the system in the stated mode; the authorisation period is the watch, stated as 12 hours.

Decay period. 12 hours, coextensive with the authorisation. Principal propositions: this installation is a defended military objective; the airspace within the engagement envelope is closed to civil and friendly traffic by NOTAM and by coordination with the air component; objects approaching within the envelope at or above the stated speed and on the stated profile are hostile.

Cancellation. Yes—the system is under continuous local supervision and can be placed in a hold state within seconds by the duty operator.

Information state. Airspace closure confirmed at watch handover; friendly air movement schedule held; no civil aviation activity notified within the envelope.

Comment. The record is short because the case is easy, and that is the point: a rule whose paperwork is proportionate to the difficulty of the case will be completed. Note also what the record would expose if the facts changed. Were friendly aircraft to be operating within the envelope, the third proposition would be false, and the record made at watch handover would show that the authorisation was premised on its truth.

A.2 Loitering munition, urban employment within communications range

Interval. Up to 4 hours of loiter, plus terminal flight.

Decay period. 40 minutes. Principal propositions: the emitter at the recorded location is a military radar in continuous use; the two structures within the effects radius are unoccupied; the road on the eastern approach carries no civilian traffic during the hours of the authorisation.

Cancellation. Yes—within communications range throughout; abort available to the mission operator at any point up to terminal dive.

Information state. Imagery of 90 minutes' vintage; signals intelligence of 3 hours' vintage confirming emitter activity; no pattern-of-life product on the two structures more recent than the preceding day.

Comment. Here the record does the work the framework is for. The interval of 4 hours exceeds the decay period of 40 minutes by a factor of six, so paragraph 2(a) is not satisfied and the employment is lawful only under 2(b)—which it is, because abort is available. **The record also discloses the weakness in its own reasoning:** a proposition about the occupancy of two structures is being carried on pattern-of-life material a day old, and that is visible on the face of the document to anyone reviewing it, before anything has happened.

A.3 Emplaced ground system, thirty-day denial task

Interval. 30 days.

Decay period. 72 hours. Principal propositions: no non-combatant movement across the sector; no friendly patrolling within the arcs; the approach remains the only trafficable route into the position.

Cancellation. *This is the entry that decides the case.* If the answer is no, paragraph 2 is not satisfied on either limb—the interval exceeds the decay period tenfold and the system cannot be reached—and the employment is unlawful under the element. If the answer is yes, by a receive-only recall channel with a stated command-delivery expectation, the employment falls under 2(b) and the remaining question is whether the recall is reliable enough to be relied on.

Information state. Sector survey at emplacement; local knowledge of civilian movement patterns from the preceding month; no continuing observation after emplacement.

Comment. The record makes the legal question unmissable and reduces it to a single line. It also shows why the companion engineering paper's finding matters so much to the law: the entry that decides lawfulness is the cancellation entry, and the capability it asks about costs approximately one part in a thousand of the

platform's energy budget. A state that answers "no" to A.3's third entry has declined a nearly free capability and thereby made an otherwise defensible employment unlawful.

A.4 What the three records show

The burden is proportionate to the difficulty. The easy case takes four short lines. The hard case takes four short lines. What varies is not the length of the record but how much it exposes.

The record discloses reasoning, not conclusions, which is what makes it assessable. In none of the three does the commander assert that control was appropriate. Each states a period, the propositions it rests on, the vintage of the information, and whether the system can be reached.

And in the hardest case it converts a contested legal question into a factual one. Whether a thirty-day armed emplacement is lawful is, on the element, the question whether it can be cancelled within 72 hours. That is a question about a communications architecture, with a yes or no answer, recorded before the fact.

A.5 What this Appendix does not establish

It does not show that commanders would complete such records accurately under pressure, which is an empirical question about institutions and cannot be settled by drafting.

It does not show that the decay periods stated above are correct. They are illustrative figures chosen to make the form legible, and a real estimate would rest on intelligence this author does not have.

And it does not address the case where a system does something outside the range of behaviours any decay estimate contemplates. Part VII §26 concedes that residue and locates it with design and review rather than with the employment decision, and nothing in this Appendix reaches it.

Postscript—The September 2026 Outcome

Dated 15 October 2026. The body of this paper is reproduced unchanged from the version of record of 17 August 2026. This postscript records what the Group of Governmental Experts adopted at the close of its mandate, and what it does to the argument. Nothing in it requires a claim to be withdrawn. Three claims are borne out, an objection the paper makes is met in part, and the variable the paper is built on has appeared in the text, in the sentence that defines what the characterisation excludes.

P.1 What happened

The Group met from 31 August to 4 September 2026 and on 4 September adopted the report of its 2024-2025-2026 sessions, "as contained in CCW/GGE.1/2026/CRP.1/Rev.2, as orally amended", issued as CCW/GGE.1/2026/3 in an advance copy dated 14 September 2026.⁹⁵ The Group recommends that High Contracting Parties, at the Review Conference of 16 to 20 November 2026, "consider and endorse the set of elements of an instrument, without prejudging its nature, and other possible measures to address emerging technologies in the area of lethal autonomous weapon systems, as contained in paragraphs 25 to 44" of the report, and that the Review Conference "decides by consensus on the possible modalities of the future work of the Group".⁹⁶ That is endorsement and not a mandate to negotiate.

⁹⁵GGE on LAWS, *Report of the 2024-2025-2026 sessions of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems*, UN Doc CCW/GGE.1/2026/3 (advance copy, 14 September 2026) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW-GGE.1-2026-3-Advance_copy.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW-GGE.1-2026-3-Advance_copy.pdf) accessed 15 October 2026, para 46: "On 4 September 2026, the Group adopted its final report as contained in CCW/GGE.1/2026/CRP.1/Rev.2, as orally amended, to be issued as document CCW/GGE.1/2026/3." The advance copy states that it is issued without formal editing.

⁹⁶*ibid*, para 45(a)–(b).

The advance copy is issued without formal editing, and a reader should check the paragraph numbering against the edited document when it appears. The adopted text is not the draft of 3 September 2026: provisions were altered and removed between that draft and adoption the next day, and where the rolling text of 5 June and the drafts of 7 July and 3 September are mentioned below it is for what changed on the way to adoption.⁹⁷

The paper's opening proposition—that a decade of work has produced a rolling text and no instrument—now reads: an agreed set of elements, and no instrument. The second half is unchanged, and it is the half the argument rests on.

P.2 The formula analysed in Part III has been deleted

Part III is built against the rolling text's central formula, "context-appropriate human judgement and control". The words "context-appropriate" appear nowhere in the adopted text. Paragraph 34 provides instead that, to uphold compliance with the applicable rules, "control and human judgement with regard to LAWS are needed", that human beings exercise such control "directly or indirectly, including through measures taken before and/or during the use of LAWS", and that they "need to exercise judgement to determine those measures necessary to ensure compliance with applicable IHL" and "to implement such measures in practice".⁹⁸

The quotation in Part III is therefore historical. The diagnosis is not, and the reader should judge whether it was right. Section 8.3 argued that the formula could not do the work for three reasons in combination. All three stand against the adopted text, and the third stands only because of something that happened between the last draft and adoption.

There is still no determining body. Nothing in the adopted text creates one, and the recommendation to the Review Conference is to consider and endorse a set of elements without prejudging their nature.

The variable is still unnamed. The paper's objection was that "context-appropriate" does not say appropriate along which dimension. The qualifier has gone and the replacement names no dimension either: a requirement to exercise judgement to determine the measures necessary is a requirement to decide what is required. Between December and September the formula did not become more verifiable. It became shorter.

Compliance is still unobservable, and here the record is worth setting out. The 2026 drafts had moved toward the paper on this point. The draft of 3 September 2026, at paragraph 39(c), would have had States ensure that these systems "operate in an adequately reliable manner and that their functioning and anticipated effects are predictable and traceable to those responsible for their development and use"—not the record Part V's rule would require, but an obligation of traceability where the paper said there was none.⁹⁹ The adopted text does not carry it. Paragraph 38 asks for testing and evaluation and for guidance, training and instructions; paragraph 42 lists understandable human-machine interfaces among the means of promoting responsibility; nothing in paragraphs 25 to 44 requires anything to be recorded or produced from which the presence or absence of control could afterwards be established. The third limb of section 8.3 therefore describes the adopted text as it described the rolling text, and the sentence this paper would otherwise have had to qualify was removed before it could be relied on.

⁹⁷GGE on LAWS, *Report of the 2024-2025-2026 sessions*, UN Doc CCW/GGE.1/2026/CRP.1 (7 July 2026) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW-GGE.1-2026-CRP.1.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW-GGE.1-2026-CRP.1.pdf) accessed 15 October 2026; UN Doc CCW/GGE.1/2026/CRP.1/Rev.1 (3 September 2026) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW-GGE.1-2026-CRP.1-Rev.1_-_03_September_.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW-GGE.1-2026-CRP.1-Rev.1_-_03_September_.pdf) accessed 15 October 2026; and, before them, GGE on LAWS, 'Rolling text, status date: 05 June 2026' [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/GGE_LAWS_-_Rolling_Text_-_5_June_2026.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/GGE_LAWS_-_Rolling_Text_-_5_June_2026.pdf) accessed 15 October 2026. The intermediate draft CCW/GGE.1/2026/CRP.1/Rev.2, named in the adoption paragraph, has not been located; the comparisons here are with the texts of 5 June, 7 July and 3 September.

⁹⁸CCW/GGE.1/2026/3 (n 95), para 34.

⁹⁹CCW/GGE.1/2026/CRP.1/Rev.1 (n 97), para 39(c). The rolling text of 5 June 2026 (n 97), para 15(C), and the draft of 7 July 2026 (CCW/GGE.1/2026/CRP.1 (n 97), para 39(c)) had "traceable and explainable"; "explainable" did not survive to 3 September, and "traceable" did not survive to adoption.

P.3 The characterisation has changed, and the change bears on section 8.1

Section 8.1 quotes the rolling text’s characterisation of these systems as “a functionally integrated combination of one or more weapons and technological components, that can identify, select, and engage a target, without intervention by a human operator in the execution of these tasks”, and objects to “functionally integrated” because it risks attaching the characterisation to a single physical unit. The adopted text, at paragraph 25, no longer carries the phrase—the draft of 3 September had already dropped it—and characterises a lethal autonomous weapon system as “a combination of one or more weapons and technological components operating together, that is capable of identifying, selecting, and engaging a target in a way that results in loss of life, without intervention by a human being operating the system”.¹⁰⁰ The objection is met in part: “operating together” says less about physical integration than “functionally integrated” did, though it still describes the components as a set. The qualifier on human intervention is not new—section 8.1 quotes its December form—but it has been recast from “a human operator in the execution of these tasks” to “a human being operating the system”, and the question it raises is the one this paper is built to answer: whose intervention, and when. The effect is now stated as loss of life rather than, as the draft of 3 September had it, force; the provision that a capacity for non-lethal use does not take a system outside the characterisation has been carried since December. Nothing in this paper turns on either.

What matters most is the second sentence of paragraph 25(a). A system is not within the characterisation “when the personnel planning, deciding upon, or carrying out the attack determine the specific target or specific group of targets and the timing of the attack”. The draft of 7 July 2026 had the specific target; the words “and the timing of the attack” were added by 3 September and survived to adoption.¹⁰¹ That is the first appearance in the text of the variable this paper is built on. Whether a system is within the characterisation now turns, in part, on whether a human being fixed the time of the attack, and the text says so in the sentence that defines what the characterisation excludes. The element at Part VIII is drafted around “the last decision by a human being” and the interval between that decision and the effect precisely so that it is indifferent to where along the chain that person sits; the new sentence leaves inside the characterisation exactly the case the element is addressed to, a person who fixed the category and the area but not the time.

P.4 Section 16 is borne out

Section 16 argues that duration is not a fourth parameter of equal standing with the types of targets, geographical scope and the number of engagements, because three of the four constrain an engagement while the fourth constrains the continuing validity of the judgement that authorised it.

In the adopted text the list has changed. Paragraph 35(b) lists, among the measures of control and human judgement, “Limiting the scale of the operation, types of targets, duration, and geographical scope, as appropriate”.¹⁰² The number of engagements has gone; the draft of 7 July 2026 still listed it, the draft of 3 September did not, and the adopted text does not restore it. Duration has survived. The paper does not claim to have caused that and offers no view on why it happened. It records only that the list which implied parity has lost a member, that the member the paper argued was of a different kind is among those still standing, and that nothing in the adopted text says what sets it: “as appropriate” is the whole of the guidance.

¹⁰⁰CCW/GGE.1/2026/3 (n 95), para 25 and para 25(b). The draft of 7 July 2026 had kept “functionally integrated”, attached to the technological components rather than to the combination; the draft of 3 September 2026 had “a combination of one or more weapons and technological components that, operating together, can identify, select, and engage a target with force”.

¹⁰¹ibid, para 25(a), compared with CCW/GGE.1/2026/CRP.1 (n 97), para 25(a), “determines the specific target or group of targets to be engaged by the system”, and CCW/GGE.1/2026/CRP.1/Rev.1 (n 97), para 25(a), in which “and the timing of the attack” first appears. A suggestion circulated under the Chair’s authority on 31 August 2026 had proposed “and the moment of engagement”: GGE on LAWS, ‘CCW/GGE.1/2026/CRP.1 – Subsection A: Additional suggestions’ (31 August 2026) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/2026_08_31_CCW_GGE_LAWS_-_Additional_suggestions_-_Para_25.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/2026_08_31_CCW_GGE_LAWS_-_Additional_suggestions_-_Para_25.pdf) accessed 15 October 2026. The December 2025 text had no exclusion sentence at all; its part I para 1(A) said only that a human operator setting the parameters of the targets does not exclude the system.

¹⁰²CCW/GGE.1/2026/3 (n 95), para 35(b); CCW/GGE.1/2026/CRP.1/Rev.1 (n 97), para 37(c), “including, inter alia, by restricting the scale of the operation, types of targets, duration, and geographical scope”; CCW/GGE.1/2026/CRP.1 (n 97), para 36(b), “the scale of the operation, the types of targets, duration, geographical scope, and the number of engagements that LAWS can undertake autonomously”.

P.5 Part IV is borne out, by subtraction

Part IV argues that the review obligation in Article 36 of Additional Protocol I is broad along the wrong axis: wide across the procurement pipeline—study, development, acquisition, adoption—and absent across operational life, which is where the properties that determine whether these systems comply with the law of targeting vary. Section 12 sets out three cases the obligation does not reach, the first of which is a system updated in service.

The adopted text restates Article 36 and stops. Paragraph 36 provides that “in the study, development, acquisition, or adoption of a new LAWS, determination must be made whether its employment would, in some or all circumstances, be prohibited by international law”—the four pre-employment stages, in the Protocol’s words.¹⁰³ The drafts had reached further. The rolling text of 5 June 2026 and the draft of 7 July would have had States conduct legal reviews “as appropriate, across the life cycle of LAWS”, and ensure that “any further development of, or modifications to, existing LAWS that have a substantial impact on the operational functioning of their target identification, selection, and engagement functions, or on their anticipated effects, are reviewed”; the draft of 3 September tied the review to the four Article 36 stages but kept a re-review of modifications with a “material impact” on the same functions.¹⁰⁴ The adopted text carries neither the life-cycle framing of the review nor the re-review; the one life-cycle sentence that survives, at paragraph 40, attaches to responsibility and not to review. What remains of the drafts’ reach into service is paragraph 37, which asks that risks be considered “during the design, development, testing, and deployment” of these systems—a consideration of risk, not a determination of lawfulness—and paragraph 34’s judgement exercised “before and/or during the use”, which is the formula discussed at P.2.¹⁰⁵

The consequence for Part IV is direct. The re-review the drafts carried would have reached the first of the three cases at section 12 and neither of the other two; it was the one provision in the process that looked along the second axis, and it was removed. Nothing in the adopted text attaches a determination to the weapon in service or to the employment decision. The diagnosis at section 12—breadth across the pipeline, nothing across operational life—describes the adopted text as it describes Article 36, and the second occasion proposed at section 14 is as absent from the set of elements as it is from the Protocol. The falsification condition at section 4.3 for the Article 36 argument is a question about state practice, and a text that declines to require re-review does not answer it either way.

P.6 The element’s route of entry has changed; its constraints have not

Part VIII drafts the element as it would have to be drafted to enter the rolling text. There is no longer a rolling text open for insertion. There is a closed set of elements going to the Review Conference, which is asked to consider and endorse it, and what follows the Review Conference will be decided there.

The four constraints at section 29 are unaffected, because they are constraints on what a consensus body can adopt rather than on any particular document. A proposal must be statable as a single element rather than bargained as a package, must require no state to abandon a position, must be assessable against something other than agreement, and must be capable of operating whether or not an instrument follows. The fourth of them is now the one that matters. The element satisfies it by whichever of section 33’s four routes it is taken—a national standard, an Article 36 review procedure, a coalition standard, or an element of whatever instrument follows—and section 33 already says that the first three of those require nobody’s agreement. That remains the practical answer.

P.7 What has not changed

No claim in Parts I, II, V, VI or VII is affected. The three cases stand on the public record of the systems and not on the treaty text. The responsibility analysis at Part VII runs through Article 57 of Additional Protocol I, command responsibility and the rules of state responsibility; the adopted text restates, at paragraphs 40 to

¹⁰³CCW/GGE.1/2026/3 (n 95), para 36; compare Article 36 of Additional Protocol I, set out at section 11.

¹⁰⁴Rolling text, status 05 June 2026 (n 97), para 14, chapeau and sub-paragraph B; CCW/GGE.1/2026/CRP.1 (n 97), para 38, chapeau and sub-paragraph (b), in the same words; CCW/GGE.1/2026/CRP.1/Rev.1 (n 97), para 38(a)–(b).

¹⁰⁵CCW/GGE.1/2026/3 (n 95), paras 37 and 34.

44, that States, parties to armed conflict and individuals “remain at all times responsible” for their decisions on the use of these systems and that accountability “cannot be transferred to machines”, which is the premise Part VII starts from, and it adds nothing that changes the analysis.¹⁰⁶ The fact-decay rate, the two-limb rule at section 17.3 and the recording requirements at section 18 are unchanged.

The falsification conditions at section 4.3 and the limits at section 35 stand as written. The first of them—an instrument, or a negotiating mandate, before the November 2026 Review Conference—has not been met by the Group, which recommends endorsement; whether the Review Conference meets it is the question November will answer. To the listed conditions one may be added that follows from section 8.3: the analysis of Part III would change if a determining body were created, if a text named the dimension along which control is to be assessed, or if a text required something to be recorded or produced from which control could afterwards be established. None has happened, and the one provision in the 2026 drafts that moved toward the third of them—the traceability sub-paragraph carried, from the rolling text of 5 June to the draft of 3 September, in the paragraph on understanding a system’s capabilities and limitations—did not survive to adoption.

¹⁰⁶ibid, paras 40–44; the quoted words are at paras 40 and 41.

Table of Instruments and Official Documents

Treaties

Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May Be Deemed to Be Excessively Injurious or to Have Indiscriminate Effects (adopted 10 October 1980, entered into force 2 December 1983) 1342 UNTS 137

Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I) (adopted 8 June 1977, entered into force 7 December 1978) 1125 UNTS 3

Official documents

International Law Commission, *Articles on Responsibility of States for Internationally Wrongful Acts* (2001), annexed to UNGA Res 56/83 (12 December 2001) UN Doc A/RES/56/83

US Department of Defense, Directive 3000.09, *Autonomy in Weapon Systems* (21 November 2012, updated 25 January 2023)

Working Paper of the People's Republic of China on Lethal Autonomous Weapons Systems (2022)

Mandate decided by the 2023 Meeting of the High Contracting Parties to the Convention on Certain Conventional Weapons

'Joint statement to the September 2025 session of the CCW GGE LAWS' (UNODA documents library)

UNGA Res 80/57 (1 December 2025) UN Doc A/RES/80/57

GGE on LAWS, 'Rolling text, status date: 18 December 2025' (UNODA documents library, CCW GGE on LAWS (2026)); earlier revisions of 8 November 2024 and 12 May 2025

GGE on LAWS, *Report of the 2024-2025-2026 sessions of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems*, UN Doc CCW/GGE.1/2026/3 (advance copy, 14 September 2026); drafts CCW/GGE.1/2026/CRP.1 (7 July 2026) and CCW/GGE.1/2026/CRP.1/Rev.1 (3 September 2026)

GGE on LAWS, 'Rolling text, status date: 05 June 2026'

GGE on LAWS, 'CCW/GGE.1/2026/CRP.1 – Subsection A: Additional suggestions' (31 August 2026)

Bibliography

Article 36, 'Key comments on the current CCW GGE "rolling text" on autonomous weapons' (17 February 2025)

Article 36, 'UK stays mute as France, Germany and 40 more states agree that draft rules on "killer robots" are ready to negotiate' (12 September 2025)

Asaro P, 'On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making' (2012) 94 *International Review of the Red Cross* 687

Automated Decision Research, 'Israel Aerospace Industries HAROP loitering munition'

Boothby WH, *Weapons and the Law of Armed Conflict* (2nd edn, Oxford University Press 2016)

Champagne M and Tonkens R, 'Bridging the Responsibility Gap in Automated Warfare' (2015) 28 *Philosophy & Technology* 125

Crootof R, 'The Killer Robots Are Here: Legal and Policy Implications' (2015) 36 *Cardozo Law Review* 1837

Crootof R, 'A Meaningful Floor for "Meaningful Human Control"' (2016) 30 *Temple International and Comparative Law Journal* 53

Ekelhof MAC, ‘Moving Beyond Semantics on Autonomous Weapons: Meaningful Human Control in Operation’ (2019) 10 *Global Policy* 343

Farmer CIV, companion paper on closed-system ground robotics (in preparation): the emplaced armed platform whose specification, power budget and recall pricing Parts V and VI rely on

Henckaerts J-M and Doswald-Beck L (eds), *Customary International Humanitarian Law, Volume I: Rules* (Cambridge University Press 2005)

Horowitz MC and Scharre P, *Meaningful Human Control in Weapon Systems: A Primer* (Center for a New American Security Working Paper, March 2015)

Human Rights Watch, *Stopping Killer Robots: Country Positions on Banning Fully Autonomous Weapons and Retaining Human Control* (10 August 2020)

Human Rights Watch, ‘Review of the 2023 US Policy on Autonomy in Weapons Systems’ (14 February 2023)

‘IAI Harop’ (Wikipedia)

‘IAI Harpy’ (Wikipedia)

‘IAI Harpy NG’ (Wikipedia)

International Committee of the Red Cross, *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* (ICRC 2006)

International Committee of the Red Cross, *ICRC Position on Autonomous Weapon Systems* (12 May 2021), reproduced with its background paper at (2021) 102 *International Review of the Red Cross* 915

International Committee of the Red Cross, *Autonomous Weapon Systems and International Humanitarian Law* (position paper, October 2025; version with executive summary, March 2026)

Lieber Institute West Point, ‘Human Oversight with Chinese Characteristics: Lethal Autonomous Weapons in the CCW GGE’

Matthias A, ‘The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata’ (2004) 6 *Ethics and Information Technology* 175

McClelland J, ‘The Review of Weapons in Accordance with Article 36 of Additional Protocol I’ (2003) 85 *International Review of the Red Cross* 397

Roff HM and Moyes R, *Meaningful Human Control, Artificial Intelligence and Autonomous Weapons* (Article 36 Briefing Paper, April 2016)

Schmitt MN and Thurnher JS, ‘“Out of the Loop”: Autonomous Weapon Systems and the Law of Armed Conflict’ (2013) 4 *Harvard National Security Journal* 231

SIPRI, *Implementing Article 36 Weapon Reviews in the Light of Increasing Autonomy in Weapon Systems* (SIPRI Insights on Peace and Security No 2015/1)

Sparrow R, ‘Killer Robots’ (2007) 24 *Journal of Applied Philosophy* 62

Stop Killer Robots, ‘156 states support UNGA resolution on autonomous weapons’ (2025)

Tigard DW, ‘There Is No Techno-Responsibility Gap’ (2021) 34 *Philosophy & Technology* 589

Varela L, ‘The Final Stretch Before the Finishing Line’ (2026) 14(2) *CCW Report* (Reaching Critical Will, 11 March 2026)

URLs and access dates for online sources, and pinpoint references throughout, are given in the footnotes. Only works cited in this paper are listed.