

THE ACCOUNTABILITY OF SOFTWARE
DEVELOPERS FOR WAR CRIMES INVOLVING
AUTONOMOUS WEAPONS: THE ROLE OF THE
JOINT CRIMINAL ENTERPRISE DOCTRINE

Elliot Winter

ISSN 0041-9915 (print) 1942-8405 (online) • DOI 10.5195/lawreview.2021.822
<http://lawreview.law.pitt.edu>



This work is licensed under a Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 United States License.



This site is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.

THE ACCOUNTABILITY OF SOFTWARE DEVELOPERS FOR WAR CRIMES INVOLVING AUTONOMOUS WEAPONS: THE ROLE OF THE JOINT CRIMINAL ENTERPRISE DOCTRINE

Elliot Winter*

ABSTRACT

This Article considers the extent to which the joint criminal enterprise doctrine could be invoked to hold software developers criminally accountable for violations of international humanitarian law involving autonomous weapons. More specifically, it considers whether the third part of the concept—which concerns common criminal purposes—might be brought to bear to achieve this end. The doctrine is deconstructed into five components, and each component is analyzed both in abstract and in terms of practical application. The Article establishes that, in certain contexts, software developers can and should be held accountable through this mechanism. Thus, it demonstrates that it is possible to avoid the emergence of a “responsibility gap” if, or more likely when, autonomous weapons with offensive capabilities are finally deployed on the battlefield.

* The author is a Lecturer (Assistant Professor) in International Law at Newcastle University Law School in the United Kingdom.

INTRODUCTION

The International Committee of the Red Cross (ICRC) defines an autonomous weapon as any weapon system with autonomy in its critical functions that can select and attack targets without human intervention.¹ The extent to which the use of autonomous weapons might be compatible with substantive obligations in international humanitarian law (IHL) is a complex issue. The author has written previously on the intersection of these “killer robots” with key humanitarian law principles such as distinction,² proportionality,³ and precaution.⁴ The present Article represents something of a departure because instead of considering whether the use of autonomous weapons would comply with the law, it focuses on how international criminal law secures individual accountability for violations of IHL involving such weapons. In other words, it considers potential criminal accountability where, for example, a machine targets a civilian, acts in a disproportionate manner, or fails to issue the appropriate warning.

This issue is important because the value of any substantive legal rule is dependent, at least in part, on how amenable that rule is to enforcement. As the United Nations (UN) Special Rapporteur, Christof Heyns, noted: “Without the promise of accountability, deterrence and prevention are reduced, resulting in lower protection of civilians and potential victims of war crimes.”⁵ Thus, if there are no clear consequences for misusing autonomous weapons, individuals who wish to operate them may see this as a license to deploy machines that are not capable of complying with the law. The effect of this would be the deterioration of real-world protections for civilians. Of course, “robots have no moral agency” and cannot be

¹ INT’L COMM. RED CROSS, AUTONOMOUS WEAPON SYS.: IMPLICATIONS OF INCREASING AUTONOMY IN THE CRITICAL FUNCTIONS OF WEAPONS 8 (2016), https://icrcndresourcecentre.org/wp-content/uploads/2017/11/4283_002_Autonomus-Weapon-Systems_WEB.pdf.

² Elliot Winter, *The Compatibility of Autonomous Weapons with the Principle of Distinction in the Law of Armed Conflict*, 69 INT’L & COMPAR. L.Q. 845 (2020).

³ Elliot Winter, *Autonomous Weapons in Humanitarian Law: Understanding the Technology, Its Compliance with the Principle of Proportionality and the Role of Utilitarianism*, 6 GRONINGEN J. INT’L L. 183 (2018).

⁴ Elliot Winter, *The Compatibility of the Use of Autonomous Weapons with the Principle of Precaution in the Law of Armed Conflict*, 58 MIL. L. & L. WAR REV. 240 (2020).

⁵ Christof Heyns (Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions), *Rep. on the Extrajudicial, Summary, or Arbitrary Executions*, ¶ 75, U.N. Doc. A/HRC/23/47 (Apr. 9, 2013), https://www.ohchr.org/Documents/HRBodies/HRCouncil/RegularSession/Session23/A-HRC-23-47_en.pdf [hereinafter Heyns].

held liable for violations of IHL themselves.⁶ The Group of Governmental Experts (GGE)—formed by states under the auspices of the Convention on Conventional Weapons (CCW) to investigate the autonomous weapons phenomenon, confirmed this point.⁷ They concluded that “Human responsibility for decisions on the use of weapons systems must be retained since accountability cannot be transferred to machines.”⁸ This lack of agency on the part of machines means that we must look elsewhere to identify who should be held liable for war crimes committed through autonomous weapons.

This Article will focus on one class of individuals in particular: software developers. This category is of special interest because autonomous weapons are, more than anything else, a product of code designed by an array of military and civilian programmers. If those people are the creators of an entity that has no free agency of its own, arguably, they should be responsible if that entity violates international law. Of course, programmers generally work in an inherently legal manner—they are not in the business of breaking the rules. Rather, it will be those individuals who use the products of developers’ labors that will directly violate the law, i.e., the combatants and fighters who actually use the technology in war. Therefore, this work considers the extent to which programmers could be held accountable by virtue of sharing a goal with these individuals. In particular, it considers whether they could be held liable through the joint criminal enterprise doctrine (JCE), notably “JCE III.”

The ambit of the Article does not include accountability for violations of human rights obligations,⁹ accountability for violations of aspects of domestic law such as

⁶ *Id.* ¶ 76.

⁷ 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, Oct. 10, 1980, 1342 U.N.T.S. 137.

⁸ Grp. Governmental Experts, *Rep. of the 2019 Session of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems*, at 13, U.N. Doc. CCW/GGE.1/2019/3 (Sept. 25, 2019); see also U.N. Secretary-General’s Message to Meeting of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons System (Mar. 25, 2019), <https://www.un.org/sg/en/content/sg/statement/2019-03-25/secretary-generals-message-meeting-of-the-group-of-governmental-experts-emerging-technologies-the-area-of-lethal-autonomous-weapons-systems>.

⁹ Andrea Spagnolo, *What Do Human Rights Really Say About the Use of Autonomous Weapons Systems for Law Enforcement Purposes?*, in *USE AND MISUSE OF NEW TECHNOLOGIES: CONTEMPORARY CHALLENGES IN INTERNATIONAL AND EUROPEAN LAW* 55 (Elena Carpanelli & Nicole Lazzarini eds., 2019).

tortious obligations,¹⁰ or the attribution issues that arise in the context of the law on state responsibility.¹¹ These issues are all worthy of consideration, but they cannot be accommodated here.

I. THE NEED FOR SOFTWARE DEVELOPER ACCOUNTABILITY

Naturally, the default position for individual criminal accountability is that the person who directly perpetrates the crime carries the blame. This means that the operator of a means of warfare—whether it is a pistol or a grenade—is responsible if they misuse that weapon. According to the Rome Statute of the International Criminal Court (the Rome Statute), “a person shall be criminally responsible and liable for punishment for a crime . . . if that person . . . commits such a crime . . . as an individual.”¹² A number of extensions accompany this default position whereby other individuals can be drawn into the mix due to having involvement in the commission of a crime even if they did not pull the trigger themselves. For example, in armed conflict, the commander of a weapon operator is responsible for a crime if he or she orders the misuse of the weapon. More specifically, the Rome Statute provides that a person commits a crime if he or she “orders, solicits or induces the commission of . . . a crime which in fact occurs or is attempted.”¹³ Thus, if a soldier decided to use an autonomous weapon with the intention of committing an international crime—or if a commander ordered a soldier to use such a weapon to commit a crime—then criminal liability would arise in the normal way. This is because, as Amoroso put it, the autonomous weapon “would be nothing but a tool in the criminal hands of human agents” and this would make “responsibility ascription relatively unproblematic.”¹⁴ Consequently, for example, Schmitt notes that the operator of an autonomous weapon which “cannot distinguish civilians from

¹⁰ Daniele Amoroso & Benedetta Giordano, *Who Is to Blame for Autonomous Weapons Systems' Misdoings?*, in *USE AND MISUSE OF NEW TECHNOLOGIES: CONTEMPORARY CHALLENGES IN INTERNATIONAL AND EUROPEAN LAW* 211, 226–29 (Elena Carpanelli & Nicole Lazzarini eds., 2019).

¹¹ *Id.* at 224–25.

¹² Rome Statute of the International Criminal Court, art. XXV(3)(a), July 7, 1998, 2187 U.N.T.S. 90 [hereinafter Rome Statute].

¹³ *Id.* art. XXV(3)(b).

¹⁴ Amoroso & Giordano, *supra* note 10, at 217.

combatants [but who] employs the system in an area where the two are intermixed has committed the war crime of indiscriminate attack.”¹⁵

Those cases would indeed be relatively straightforward. However, these categories assume that military personnel such as soldiers and commanders retain control over weapons throughout the course of their deployment. The novelty of autonomous weapons is that they would be released onto the battlefield and would act independently thereafter. Their actions would be their own or, rather, would result from the amalgam of programming that came together to form their code. Thus, there is a question over the extent to which the individual accountability assigned to soldiers or commanders might need to be displaced by, or augmented with, the individual accountability of other actors. There are myriad candidates to whom accountability for the misdeeds of autonomous weapons might conceivably be attributed. Heyns noted, for example, that accountability for the crimes of killer machines could theoretically fall to “software programmers, those who build or sell hardware, [procurement officials], military commanders, subordinates who deploy these systems and political leaders.”¹⁶ Similarly, Boothby posited that, in addition to traditional military commanders, the range of those with potential criminal responsibility for violations of IHL could include procurement officials, engineers, scientists, computer programmers, technicians, operators, lawyers, and planners.¹⁷

From the range of options outlined above, the group with the most potential in terms of accountability for an autonomous weapon’s actions is software developers. This is because the decision-making influence individual soldiers and military commanders are losing is likely to be replaced principally by the decision-making influence of those who are tasked with coding the machines. Several commentators have expressed this sentiment. As Schmitt put it, “a human must decide how to program the system [and], self-evidently, that individual would be accountable for programming it to engage in actions that amounted to war crimes.”¹⁸ For McFarland, “the higher the level of autonomous operation exhibited by a weapons system, the less control will be exercised by a human operator in the field and the more control

¹⁵ Michael N. Schmitt, *Autonomous Weapon Systems and International Humanitarian Law: A Reply to the Critics*, HARV. NAT’L SEC. J. (Feb. 5, 2013), <https://harvardnsj.org/2013/02/autonomous-weapon-systems-and-international-humanitarian-law-a-reply-to-the-critics/>.

¹⁶ Heyns, *supra* note 5, ¶ 77.

¹⁷ William H. Boothby, *Highly Automated and Autonomous Technologies*, in *NEW TECHNOLOGIES AND THE LAW IN WAR AND PEACE* 137, 152–53 (William H. Boothby ed., 2019).

¹⁸ Schmitt, *supra* note 15, at 33.

will be exercised by developers of the control software.”¹⁹ More dramatically, Amoroso cautioned that the rise of autonomous weapons has the potential to reduce soldiers and commanders “to puppets in the hands of war geek-criminals”²⁰ in the form of software developers who “might well have greater weight than final users on the way [autonomous weapons] take targeting decisions.”²¹

In short, as the influence of soldiers and commanders on battlefields wanes, it is the influence of software developers that will wax the most. It is only right to apportion criminal responsibility accordingly.

II. SCOPING THE OPTIONS FOR SOFTWARE DEVELOPER ACCOUNTABILITY

Numerous mechanisms exist that might enable the prosecution of software developers for misdeeds involving their creations. They cannot all be considered in detail here, but it seems prudent to briefly survey the field before turning to the option that is the main focus of this Article.

First, and most obviously, a software developer could be prosecuted on the basis of individual accountability in the event that they programmed an autonomous weapon, intentionally or recklessly, in such a way that it would violate IHL. After all, as we saw above in the context of soldiers and commanders, the Rome Statute provides that people are criminally responsible for crimes they commit “as an individual.”²² For example, a rogue developer working within the military might decide to input or alter code with the effect of causing an autonomous weapon to misidentify civilians as combatants. In that case, the individual could be held criminally liable simply on the basis of their own conduct. However, as Schmitt observed, “it is hopefully improbable that an autonomous weapons system would be programmed to commit war crimes [and] much more likely would be a case in which a system that has not been so programmed is nevertheless used in a manner that constitutes such crimes.”²³ This point is surely correct. Those who are selected to work on this sort of programming are unlikely to have anything to gain from deliberately writing code that might one day precipitate a violation of international law—indeed, it could lose them their employment or even their liberty. Further, on

¹⁹ Tim McFarland & Tim McCormack, *Mind the Gap: Can Developers of Autonomous Weapons Systems Be Liable for War Crimes?*, 90 INT’L L. STUD. 361, 369–70 (2014).

²⁰ Amoroso & Giordano, *supra* note 10, at 219.

²¹ *Id.* at 218.

²² Rome Statute, *supra* note 12.

²³ Schmitt, *supra* note 15, at 34.

a more practical level, such programmers are unlikely to work in isolation and will almost always be working in a team where their peers would detect any attempt to corrupt the program. In any event, this sort of accountability does not present anything novel from a legal point of view—individual accountability is the best-understood form of criminal responsibility. For these reasons, this option will not be considered further here.

Second, a software developer could be liable through the doctrine of indirect perpetration. According to the Rome Statute, “a person shall be criminally responsible . . . if that person . . . commits . . . a crime . . . through another person, regardless of whether that other person is criminally responsible.”²⁴ This would bridge the gap in cases where software developers, acting like puppet masters, perpetrate violations of IHL through soldiers or other operators of autonomous weapons. Moreover, it would ensure that the chain of causation remains intact regardless of whether the operators had any knowledge of the developers’ malfeasance. This latter point may be important given that operators are unlikely to be fully versed in all the detailed programming that will comprise the artificial intelligence aboard an autonomous weapon—indeed, no human could be familiar with the millions of lines of code that would be involved. While indirect perpetration remains an important option to keep open, as with individual accountability, it is unlikely to be terribly useful in practice. This is because, again, software developers are unlikely to be interested in deliberately programming war machines to violate IHL. There is little or no motive to do so. Indeed, as the North Atlantic Treaty Organization (NATO) has noted, the violation of IHL has more potential to damage the strategic objectives of one’s own side than to harm the enemy in the long run.²⁵

Third, a software developer could be held criminally liable if he or she “aids, abets or otherwise assists” in the commission of a crime—including providing the means for its commission.²⁶ This option may appear to be of more utility as it does not come with the baggage of being predicated on the assumption that the software developer deliberately wants to corrupt an autonomous weapon. Rather, it might seem to apply to them as wholly disinterested suppliers of technology. McFarland, for example, assumes that this head of liability would apply in a similar way to software developers as to those who provided Zyklon B to the Nazis in the Second

²⁴ Rome Statute, *supra* note 12.

²⁵ N. Am. Treaty Org., Allied Joint Doctrine for Joint Targeting, ¶ 0205, AJP-3.9 (2016).

²⁶ Rome Statute, *supra* note 12, art. XXV(3)(c).

World War.²⁷ However, this is unlikely to be the case. The International Criminal Tribunal for the Former Yugoslavia (ICTY) noted in *Tadić* that an “aider and abettor carries out acts specifically directed to assist, encourage or lend moral support to the perpetration of a certain specific crime (murder, extermination, rape, torture, wanton destruction of civilian property, etc.).”²⁸ Evidently, the degree of interest and knowledge required here must be high. The individual must be *specifically* contributing to a *specific* crime. This will rule out the application of the aiding and abetting category to software developers in most, if not all, cases, as they are very unlikely to be specifically directing their efforts towards the commission of a particular war crime.

In short, none of the above options are likely to secure the accountability of software developers for violations of IHL involving autonomous weapons.

III. SOFTWARE DEVELOPERS AND THE JOINT CRIMINAL ENTERPRISE DOCTRINE

The potential lacuna in accountability for software developers left by the individual accountability, indirect perpetration, and aider and abettor models could be filled by the JCE doctrine. JCE was developed principally by the ICTY building, as Haan explained, on a number of World War II cases.²⁹ The impetus for the revival of JCE emanated from the United Nations Secretary-General (UNSG), and in particular from the Report of the Secretary-General on the establishment of the ICTY, which stated that “all persons who participate in the planning, preparation or execution of serious violations of [IHL] . . . are individually responsible for . . . violations” of IHL.³⁰ The ICTY Statute, created by the United Nations Security Council (UNSC), went on to state that anyone who contributes to the “planning, preparation or execution of a crime . . . shall be individually responsible for the

²⁷ McFarland & McCormack, *supra* note 19, at 370.

²⁸ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 229(iii) (Int’l Crim. Trib. for the Former Yugoslavia July 15, 1999).

²⁹ Verena Haan, *The Development of the Concept of Joint Criminal Enterprise at the International Criminal Tribunal for the Former Yugoslavia*, 5 INT’L CRIM. L. REV. 167 (2005).

³⁰ U.N. Secretary-General, *Rep. of the Secretary-General Pursuant to Paragraph 2 of Security Council Resolution 808*, ¶ 54, U.N. Doc. S/25704 (May 3, 1993).

crime.”³¹ In other words, not only would direct perpetrators be held accountable but so too would those who made broader contributions to a crime’s commission.

The jurisprudence of the ICTY began to build on this starting point in *Tadić* in which it was held that criminal accountability arises where “several persons having a common purpose embark on criminal activity that is then carried out either jointly or by some members of this plurality of persons” and in which it was also reaffirmed that “whoever contributes to the commission of crimes by the group of persons or some members of the group, in execution of a common criminal purpose, may be held to be criminally liable.”³² The Appeals Chamber explained that this approach is “dictated by . . . the very nature of many international crimes which . . . do not result from the criminal propensity of single individuals but constitute manifestations of collective criminality [where] the participation and contribution of [others] is often vital in facilitating the commission of the offense in question.”³³ Consequently, “international criminal responsibility embraces actions perpetrated by a collectivity of persons in furtherance of a common criminal design.”³⁴ In *Tadić*, the Appeals Chamber found an implicit basis for JCE in Article 7(1) of the ICTY Statute since “the commission of crimes . . . might also occur through participation in the realization of a common design or purpose” and because the article “does not exclude those modes of participating.”³⁵ The logic of JCE applies, *a fortiori*, in the context of autonomous weapons, where the participation of many people is required for such machines to be fielded and for an offense to be committed. When software developers, company executives, militaries, and others work together to deploy autonomous weapons which go on to cause violations of international law, there is clear potential for collective criminality of this sort.

There is nuance in the sense that JCE comes in three distinct categories. The first category, “JCE I,” which concerns co-perpetratorship was relied on in some early Trial Chamber reasoning but was rejected by the Appeals Chamber as a basis for responsibility in *Stakić* when it held that “‘co-perpetratorship’ . . . does not have support in customary international law or in the settled jurisprudence of this Tribunal,

³¹ S.C. Res. 827, Statute of the International Criminal Tribunal for the Former Yugoslavia, art. VII(1) (May 25, 1993) [hereinafter ICTY Statute].

³² *Tadić*, IT-94-1-A ¶ 190.

³³ *Id.* ¶ 191.

³⁴ *Id.* ¶ 193.

³⁵ *Id.* ¶ 190.

which is binding on the Trial Chambers.”³⁶ In any event, according to Badar, that model would have required “(i) an explicit agreement or silent consent between two or more individuals to reach a common goal; (ii) coordinated co-operation; and (iii) joint control over the criminal conduct.”³⁷ Software developers would be unlikely to satisfy these criteria. In particular, many programmers are unlikely to share the goal of the militaries they supply (indeed, they may be unaware that military forces will use the technology), and almost none will have “joint control” over the conduct of an autonomous weapon in the field. The second category, “JCE II,” deals with “concentration camp” scenarios.³⁸ Here, “the requisite mens rea comprises knowledge of the nature of the system of ill-treatment and intent to further the common design of ill-treatment.”³⁹ This is unlikely to be relevant as software developers will not be contributing to a concentration-camp-style system of any sort—let alone one that is intended to further the ill-treatment of civilians. In short, neither JCE I nor JCE II seems especially apposite for dealing with the problem at hand. This leaves the final category, “JCE III.”

According to the Appeals Chamber in *Tadić*, JCE III is applicable when there is “a common design to pursue one course of conduct where one of the perpetrators commits an act which, while outside the common design, was nevertheless a natural and foreseeable consequence of the effecting of that common purpose.”⁴⁰ In *Stakić*, which was decided a few years later, the position was clarified, and it was required that, in addition to the existence of a common purpose, the following tests must be satisfied:

- (a) crimes outside the Common Purpose have occurred; (b) these crimes were a natural and foreseeable consequence of effecting the Common Purpose; and (c) the participant in the joint criminal enterprise was aware that the crimes were a possible consequence of the execution of the Common Purpose, and in that awareness, he nevertheless acted in furtherance of the Common Purpose.⁴¹

³⁶ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 62 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

³⁷ Mohamed Badar, “Just Convict Everyone!”—*Joint Perpetration: From Tadić to Stakić and Back Again*, 6 INT’L CRIM. L. REV. 293, 296 (2006).

³⁸ *Tadić*, IT-94-1-A ¶ 202–03.

³⁹ *Id.* ¶ 220.

⁴⁰ *Id.* ¶ 204.

⁴¹ *Stakić*, IT-97-24-A ¶ 87.

For present purposes then, JCE III can be distilled into five criteria. First, there must be a common purpose. Second, a crime outside that common purpose must have occurred. Third, the crime must have been a natural and foreseeable consequence of the common purpose. Fourth, the accused must have been aware that the crime was a possible consequence of the common purpose. Fifth, the accused must have acted in furtherance of the common purpose. Criteria one, two, and five are the *actus reus* components of JCE III. Criteria three and four are the *mens rea* components. The remainder of the Article will proceed by exploring how each component of this five-part system would apply in the context of a violation of IHL that involves the actions of an autonomous weapon.

Before that, it should be noted that the status of JCE III in customary international law is a matter of some dispute. The ICTY asserted that JCE is a mode of liability that is “firmly established in customary international law and is routinely applied in the Tribunal’s jurisprudence.”⁴² The doctrine also exists in the context of the International Criminal Tribunal for Rwanda (ICTR).⁴³ That Tribunal has articulated its understanding of the doctrine in cases such as *Ntakirutimana* in which JCE III is described as the “extended form” of joint criminal enterprise.⁴⁴ It also appears in the International Convention for the Suppression of Terrorist Bombing, which the UN General Assembly adopted by consensus.⁴⁵ However, the International Criminal Court (ICC) does not endorse JCE III and instead favors a standard based on intent to commit a crime or knowledge of an impending crime.⁴⁶ Also, Guilfoyle notes that the ICTY’s revival of JCE was done “controversially and on a limited survey of trials.”⁴⁷ This Article, by demonstrating that JCE III can partially fill an accountability gap left by other heads of liability, will continue to make the case for the full acceptance of JCE III into customary international law.

⁴² *Id.* ¶ 62.

⁴³ S.C. Res. 955, Statute of the International Criminal Tribunal for Rwanda, art. VI(1) (Nov. 8, 1994) [hereinafter ICTR Statute].

⁴⁴ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17, Judgment, ¶ 465 (Int’l Crim. Trib. for Rwanda Dec. 13, 2004).

⁴⁵ International Convention for the Suppression of Terrorist Bombing, art. II(3), Jan. 12, 1998, 2149 U.N.T.S. 256.

⁴⁶ Rome Statute, *supra* note 12, art. XXV(3)(d).

⁴⁷ DOUGLAS GUILFOYLE, INTERNATIONAL CRIMINAL LAW ¶ 12.6.1 (2016).

A. *Criterion 1: Common Purpose*

The first criterion that must be satisfied for prosecution under JCE III is that the accused and the person(s) who directly committed a crime must share a common purpose. The Appeals Chamber in *Tadić* explained that what is needed is a “common design to pursue one course of conduct.”⁴⁸ More particularly, there must be “a common plan, design or purpose which amounts to or involves the commission of a crime.”⁴⁹ Similarly, in the more recent *Stakić* decision, it was held that “the existence of a common purpose which amounts to or involves the commission of a crime” is required.⁵⁰ In other words, it is not enough for the plan to be a legal one which later goes awry in its implementation. Rather, the design itself must be inherently illegal. Many judgments such as *Kvočka* refer to “common purpose” and “common criminal purpose” interchangeably, so the terms are synonymous here.⁵¹ Given the nature of the international tribunals most heavily involved in developing JCE III, it is unsurprising that the plan usually at issue is ethnic cleansing. In *Stakić*, for example, the Appeals Chamber noted that the “common purpose consisted of a discriminatory campaign to ethnically cleanse the Municipality of Prijedor by deporting and persecuting Bosnian Muslims and Bosnian Croats in order to establish Serbian control.”⁵² In *Ntakirutimana*, concerning Rwanda rather than the former Yugoslavia, a typical issue was a group plan “to forcibly remove at gun-point members of one ethnicity from their town, village or region (to effect ‘ethnic cleansing’).”⁵³

The question for present purposes is what the criminal design might be in situations where a shared endeavor led to a violation of international law through the use of an autonomous weapon. Certainly, ethnic cleansing, genocide, and similar offenses would still be captured in the context of autonomous weapons. It is possible to imagine a future in which one ethnic group tries to remove another from a geographical area by means of autonomous weapons technology. Israel, for example,

⁴⁸ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 204 (Int’l Crim. Trib. for the Former Yugoslavia July 15, 1999).

⁴⁹ *Id.* ¶ 227.

⁵⁰ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 64 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

⁵¹ Prosecutor v. Kvočka, Case No. ICTY-98-30/1-A, Appeal Judgment (Int’l Crim. Trib. for the Former Yugoslavia Feb. 28, 2005).

⁵² *Stakić*, IT-97-24-A ¶ 73.

⁵³ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17, Judgment, ¶ 465 (Int’l Crim. Trib. for Rwanda Dec. 13, 2004).

has invested heavily in autonomous weapons and might one day be tempted to use this technology to gain control of disputed Palestinian territories.⁵⁴ However, when it was operating, the ICTY could have considered any common purpose which involved the commission of a crime provided for in its Statute. After all, the ICTY Statute also covered “grave breaches” of the Geneva Conventions including “(a) wilful killing; (b) torture or inhuman treatment . . . ; (c) wilfully causing great suffering or serious injury to body or health [and] (d) extensive destruction and appropriation of property, not justified by military necessity and carried out unlawfully and wantonly.”⁵⁵ Similarly, its remit included “violations of the laws or customs of war” such as:

(a) employment of poisonous weapons or other weapons calculated to cause unnecessary suffering; (b) wanton destruction of cities, towns or villages, or devastation not justified by military necessity; (c) attack, or bombardment, by whatever means, of undefended towns, villages, dwellings, or buildings; (d) seizure of, destruction or wilful damage done to institutions dedicated to religion, charity and education, the arts and sciences, historic monuments and works of art and science [and] (e) plunder of public or private property.⁵⁶

Other international tribunals, such as the ICC (if it ever changes its position and adopts JCE III), could rely on a similarly broad range of criminal purposes to satisfy this first test. The particular crime relevant in any given case will, of course, depend on the facts.

By way of a caveat, certain actions involving autonomous weapons would not amount to criminal purposes. First, the development of autonomous weapons would not, in itself, amount to a violation of IHL. This is because international law merely provides that “in the study, development, acquisition or adoption of a new weapon . . . [states are] under an obligation to determine whether its employment would . . . be prohibited.”⁵⁷ In other words, while the law requires states to keep the compatibility of autonomous weapons under development with IHL in mind, it does

⁵⁴ See Ingvild Bode & Hendrik Huelss, *Autonomous Weapons Systems and Changing Norms in International Relations*, 44 REV. INT’L STUD. 393, 398 (2018).

⁵⁵ ICTY Statute, *supra* note 31, art. II.

⁵⁶ *Id.* art. III.

⁵⁷ Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts, art. XXXVI, June 8, 1977, 1125 U.N.T.S. 3 [hereinafter Protocol Additional].

not actually prevent the states from developing weapons whose deployment would violate IHL. Second “under international criminal law it is not a crime to conspire to commit war crimes or crimes against humanity.”⁵⁸ The only exception to this conspiracy rule is genocide—conspiring to commit genocide via autonomous weapons would be a crime even if the weapon was never actually developed or deployed.⁵⁹

Of course, establishing that a group holds a common criminal design is not enough in itself. It is also necessary to show that the accused was sufficiently connected with, or integrated into, that group. This is because, as Ambos summarized, individual liability in this context is “essentially based on . . . membership in the group pursuing the JCE.”⁶⁰ The criteria here are not especially demanding. According to *Stakić*, any “plurality of persons”⁶¹ could suffice, and the members “need not be organised in a military, political or administrative structure.”⁶² This relaxed approach to delineating what does and does not, count as integration into a group has proven controversial. Sluiter has scoffed that the lack of a rigorous standard here warrants a reinterpretation of JCE’s acronym to “just convict everybody.”⁶³ Guilfoyle recognized the concern with prosecuting individuals who had only loose links with the direct perpetrators⁶⁴ and suggested that courts “could restrict the application of the doctrine to small, closed, or identifiable groups . . . thus requiring . . . a close connection between all participants.”⁶⁵ This would shrink the JCE III net and result in fewer individuals being considered as sufficiently integrated into the group. However, the jurisprudence of both the ICTY and ICTR has shown an eagerness to keep a wide net. For example, it was held in *Karemera* that common

⁵⁸ GUILFOYLE, *supra* note 47, ¶ 12.3.3.

⁵⁹ Convention on the Prevention and Punishment of the Crime of Genocide, art. III(b), Dec. 9, 1948, 78 U.N.T.S. 277; ICTY Statute, *supra* note 31, art. IV(3)(b); ICTR Statute, *supra* note 43, art. II(3)(b).

⁶⁰ Kai Ambos, *Joint Criminal Enterprise and Command Responsibility*, 5 J. INT’L CRIM. JUST. 159, 168 (2007).

⁶¹ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 64 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

⁶² *Id.*

⁶³ Goran Sluiter, *Foreword*, 5 J. INT’L CRIM. JUST. 67, 67 (2007).

⁶⁴ GUILFOYLE, *supra* note 47, ¶ 12.6.4.

⁶⁵ *Id.* ¶ 12.6.1.

purposes can emerge among participants who have never met and where the crimes are “structurally or geographically remote from the accused.”⁶⁶

This relaxed standard for integration makes it possible to argue that software developers would be members of most “groups” when it comes to the deployment of autonomous weapons. That said, there are different categories of developers—those who work within militaries themselves, those who work in arm’s length defense contractors, and those who work in the generic commercial sector—whose products are just as likely to be sold off to civilian enterprises as military ones. Certainly, software developers who work within a military are sufficiently integrated as they wear a uniform, swear allegiance to a particular state, and operate within a chain of command. In such cases, the common purpose will likely be to prepare to engage enemy forces in a conflict of some sort. Defense contractor developers may also find themselves judged as part of a group because, again, there is no need for them to be formally absorbed into a “military, political or administrative structure.”⁶⁷ It may be that a future court asked to deal with this issue would consider the frequency and level of detail in communications between, say, a military and a software company. A defense contractor from BAE Systems seconded to the U.K. Ministry of Defence might be the prime example of a non-military programmer caught by the net as there would be hand-in-glove cooperation in such a case. Still, even freelance defense contractors working in their own offices might have sufficient contact with the military to be considered part of the “group.” When it comes to generic commercial software developers, integration into a particular military is unlikely to exist, and thus establishing group affiliation will be more difficult. Indeed, in many cases, it is likely that the country which produces the software will export it to third countries with which the developers have no connection. As wide as the net is, it seems unlikely that it would catch programmers such as these. Ultimately, the whole point of JCE was to widen the ambit of criminal law and, given the low level of integration required by the Appeals Chamber for an individual to become affiliated with a group, it seems to have stretched far enough to cover most software developers.

In summary, to establish a common purpose, there must be a criminal design and some integration of the accused into the group carrying out that design. In terms of the former, almost any war crime will suffice—it does not need to relate to ethnic cleansing or genocide even though most cases concerning JCE III to date have related to these phenomena. In terms of the latter, the integration test is set at a very low level, and there is no need for a formal structure to be in place to bind the accused

⁶⁶ *Karemera v. Prosecutor*, Case No. ICTR-98-44-AR72.5, Decision on Jurisdictional Appeals: Joint Criminal Enterprise, ¶ 14 (Int’l Crim. Trib. for Rwanda Apr. 12, 2006).

⁶⁷ *Ntakirutimana*, ICTR-96-17 ¶ 466.

and other members of the group. Rather, the standard is loose, and a court or tribunal could find evidence of affiliation between an individual and a group based on an array of unspecified contextual factors. In most cases, this criterion is likely to mean that the net will catch military programmers and defense contractor programmers but not generic commercial developers. This nuanced result should strike the balance for which JCE III constantly strives. It will ensure that military and defense contractor programmers can be held liable for any misdeeds that the technology they have helped produce goes on to perpetrate while ensuring that generalist programmers working for companies that have made more tangential contributions to the code are not unfairly brought within the ambit of international criminal law.

B. Criterion 2: Crime Outside the Common Purpose

Assuming the establishment of a common purpose, the second criterion that must be satisfied for prosecution under JCE III is that a crime *outside* the common purpose was committed by another party to that purpose. These could be expressed more neatly as “spin-off” crimes. Recall that, had the crime fallen *inside* the ambit of the common purpose, there would be a more proper charge against the accused under one of the individual liability headings discussed briefly above. As Guilfoyle summarized, the defendant “will remain liable for any crimes falling within the scope of the original plan”⁶⁸ and the *Blagojevic* case confirmed this.⁶⁹

The commission of spin-off crimes and the commission of crimes within the ambit of the common purpose might be equally wide in range. For example, while the crime inside the common purpose may have been restricted to causing “extensive destruction and appropriation of property,” the spin-off crime may have involved the “wilful killing” of civilians by another member of the group (or, though it seems less likely, vice versa).⁷⁰ Alternatively, the planned crime may have been the plunder of public or private property but, perhaps as a result of resistance, this may have escalated to the “attack, or bombardment . . . of undefended towns, villages, dwellings, or buildings.”⁷¹ In short, it is not the nature of the second violation that matters, provided it is a violation of IHL and that it falls outside the scope of the common purpose. In terms of real-world examples, the common criminal purpose

⁶⁸ GUILFOYLE, *supra* note 47, ¶ 12.6.2.

⁶⁹ Prosecutor v. Blagojevic, Case No. ICTY-02-60-T, Trial Judgment, ¶ 700 (Int’l Crim. Trib. for the Former Yugoslavia Jan. 17, 2005).

⁷⁰ ICTY Statute, *supra* note 31, art. II.

⁷¹ *Id.* art. III.

has typically been ethnic cleansing—either in Yugoslavia or in Rwanda. The spin-off crimes have taken different forms, including murder, torture, rape, destruction of property, and assault. For example, in *Krstić*, where the common purpose was to rid Srebrenica of its Bosnian-Muslim inhabitants, the accused was convicted for “the incidental murders, rapes, beatings and abuses committed in the execution of this [primary] criminal enterprise.”⁷² In *Stakić*, the Appeals Chamber held that the relevant crimes were “murder (as both a war crime and a crime against humanity) and extermination.”⁷³ In that case, killing was not within the scope of the common purpose (which was nominally to move people out of Serbian territory), but individual perpetrators had exceeded this remit by killing members of the Muslim minority in detention facilities, convoys, and through police action in municipalities. In *Ntakirutimana*, before the ICTR, it was noted that a spin-off crime might arise where “one or more of the victims is shot and killed” as part of an area’s planned ethnic cleansing.⁷⁴ In short, there are no special requirements for, or limitations on, the nature of the spin-off crime.

When it comes to establishing a spin-off crime where autonomous weapons and their developers are involved, the possibilities are endless. Another party could perpetrate any one of the crimes mentioned above to the common purpose through the medium of an autonomous weapon. In particular though, as the author has demonstrated previously, technology as it stands is not sufficiently advanced to allow the production of autonomous weapons that are able to comply with distinction⁷⁵ or proportionality⁷⁶—core principles of IHL. Therefore, the crime outside the common purpose is perhaps more likely to be a violation of one of these principles than, say, murder or torture. If a military officer ordered the use of an autonomous weapon resulting in a violation of one of these rules, whether by design or by accident, that would be sufficient for the purposes of this criterion. Of course, the spin-off offense would need to be established in court. This would require establishing the *actus reus* and *mens rea* of the military commander (or other “direct” perpetrator), determining who deployed the weapon, leading evidence, hearing defenses, and so on.

⁷² Prosecutor v. Krstić, Case No. ICTY-98-33-T, Trial Judgment, ¶ 617 (Int’l Crim. Trib. for the Former Yugoslavia Aug. 2, 2001).

⁷³ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 89 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

⁷⁴ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17, Judgment, ¶ 465 (Int’l Crim. Trib. for Rwanda Dec. 13, 2004).

⁷⁵ Winter, *supra* note 2.

⁷⁶ Winter, *supra* note 3.

At any point along the track, the principal actor's criminal liability might derail. If this occurs, then the software developer could not be charged under JCE III as it is a contingent form of liability in the sense that it relies on someone else's crime. The court must establish the other actor's crime for JCE III to be available.

In summary, a wide range of direct violations of IHL involving autonomous weapons could qualify as spin-off crimes. They may be deliberate, or they may be accidental where the weapon's inadequate technology fails to comply with requirements such as distinction or proportionality. The difficulty is that the guilt of the principal actor—in our case, the person who deployed or operated the weapon—would need to be established before considering the contingent guilt of the software developer. This introduces a potential break in the chain of responsibility. Assuming, however, that the chain remains intact, the prosecution would still need to satisfy the remaining criteria for JCE III. It is to those that we now turn.

C. Criterion 3: Natural and Foreseeable Consequence

It was noted above in the distillation of the criteria of JCE III that foreseeability and, separately, awareness of the possibility of the crime are the third and fourth requirements respectively.⁷⁷ More particularly, they comprise the twin mens rea components of JCE III. These will be subjected to individual analysis shortly, but first, it is important to address why there are two mens rea tests in the first place.

The initial legal basis for the two-pronged approach to mens rea is in *Tadić*.⁷⁸ The case provides that “responsibility for [the spin-off crime] arises only if (i) it was foreseeable that such a crime might be perpetrated by . . . members of the group and (ii) the accused willingly took that risk.”⁷⁹ Guilfoyle summarizes that “the further crime must have been objectively foreseeable and also actually foreseen and accepted by the accused.”⁸⁰ Thus, the mens rea for JCE III effectively features a double-lock: there is both an objective test and a subjective test. If the crime was a foreseeable consequence of the common purpose, but the accused had not personally apprehended the possibility of its perpetration, then there would be no liability. Equally, if the accused was aware of the risk of the spin-off crime, but that crime was not an objectively foreseeable consequence of the common purpose, there would

⁷⁷ See *supra* Part III.

⁷⁸ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 228 (Int'l Crim. Trib. for the Former Yugoslavia July 15, 1999).

⁷⁹ *Id.*

⁸⁰ GUILFOYLE, *supra* note 47, ¶ 12.6.4.

be no liability under JCE III (though there may be liability under another heading).⁸¹ This double-lock might appear to take an unusually strict stance on mens rea, but it was introduced in recognition of the fact that JCE III is a controversial basis for criminal liability given that it imposes derivative guilt on one individual flowing from the actions of another. Furthermore, as we will see, the subjective limb of the test has a low threshold. The accused need only have been aware of the risk of the crime occurring—they need not be aware of a “certainty” that the crime will be perpetrated or even of a “probability” that it will be perpetrated. For most defendants, this will not offer much protection against conviction. The remainder of this part will consider objective foreseeability in more detail, while the next part will consider subjective awareness.

Regarding the first mens rea requirement that the spin-off crime was foreseeable, the approach noted above in *Tadić* was affirmed by the Appeals Chamber in *Stakić*, which held that the spin-off crime must be a “natural and foreseeable consequence” of effecting the common purpose.⁸² In that case, the “Crisis Staff” (a wartime authority which planned, supervised, and oversaw three of the most notorious detainment centers in northwest Bosnia) had formed an “Intervention Platoon” to transport Muslim prisoners, notwithstanding the fact that the Platoon members were largely convicted criminals. The Trial Chamber’s reasoning, which the Appeals Chamber adopted, is that to be in the position of the accused with knowledge of these arrangements was “to reconcile oneself to the reasonable likelihood that those traveling on the convoy will come to grave harm and even death.”⁸³ The ICTR has adopted the same approach to foreseeability. In *Ntakirutimana*, the court observed that it must have been “foreseeable that such a crime might be perpetrated by one or other members of the group.”⁸⁴ It was held in that case that “while murder may not have been explicitly acknowledged to be part of the common purpose, it was nevertheless foreseeable that the forcible removal of civilians at gunpoint might well result in the deaths of one or more of those civilians.”⁸⁵ Of course, the list of foreseeable spin-off crimes in these cases is not

⁸¹ Ambos, *supra* note 60, at 175.

⁸² Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 87 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

⁸³ Prosecutor v. Stakić, Case No. IA-97-24-T, Trial Judgment, ¶ 600 (Int’l. Crim Trib. for the Former Yugoslavia July 31, 2003).

⁸⁴ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17-A, Appeal Judgment, ¶ 467 (Int’l Crim. Trib. for Rwanda Dec. 13, 2004).

⁸⁵ *Id.* ¶ 465.

exhaustive, and other crimes may also be foreseeable, depending on the facts at hand.⁸⁶ In short, foreseeability is approached in a fairly straightforward manner in the context of JCE III. It takes the standard of what crimes a reasonable person in the position of the accused would have been able to foresee as being natural consequences of the common purpose.

The question for present purposes is how this foreseeability requirement works in the context of autonomous weapons and software developers. Two principal difficulties arise here. Firstly, to hold someone accountable for the foreseeable results of a common purpose assumes that they are aware of that common purpose. In *Stakić*, the defendant was aware of the common purpose because he was politically affiliated with paramilitary elements who wished to see Muslims expelled from Serbian territory, and he was well situated in terms of his position to know the details of the plan.⁸⁷ Software developers may be in very different positions. Software is often created without a particular end-use in mind. It is typically developed, just as with conventional weapons, in anticipation of future conflicts. There is nothing inherently illegal about speculative development of this sort. War is a legitimate tool of state power, provided it is justified by, and conducted in line with, international law.⁸⁸ Furthermore, although states are required to keep under constant evaluation the compliance of emerging weapons technology with IHL,⁸⁹ there is no blanket provision banning the development of weapons that would breach the regime if actually used in the field (although there are specific treaties that do this in limited contexts such as the Biological Weapons Convention).⁹⁰ In situations of speculative software development such as this, where there is no common purpose at the time the coding work is being done, it will be very difficult to make a case that a reasonable software developer in the place of the accused would have foreseen the spin-off violation of IHL that ultimately unfolded.

Of course, there is a degree of nuance here in the sense that there are different types of software developers. Here we may refer again to the three categories of programmers discussed above: the generic commercial developers, the defense contractor developers, and the military developers. The generic commercial developers are the most likely to write code without any knowledge of what it might

⁸⁶ GUILFOYLE, *supra* note 47, ¶ 12.6.4.

⁸⁷ See generally *Stakić*, IT-97-24-A.

⁸⁸ THE OXFORD HANDBOOK OF THE USE OF FORCE IN INTERNATIONAL LAW (Marc Weller ed., 2015).

⁸⁹ Protocol Additional, *supra* note 57, art. XXXVI.

⁹⁰ Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on Their Destruction, art. I, Apr. 10, 1972, 1015 U.N.T.S. 163.

be utilized for or in what context. They could be programming something as innocuous as an anti-malware application or an energy efficiency add-on that could be used in a wide array of systems and by potential clients all around the world. In such scenarios, the reasonable developer in the place of the accused would be highly unlikely to foresee a future violation of IHL involving an autonomous weapon, and it seems appropriate to insert a rebuttable presumption into the law to the effect that such individuals cannot foresee spin-off crimes.

As we move along the spectrum, we get to the defense contractor developers. The position here is the most ambiguous because they are civilians with no personal military affiliation or Stakić-like positioning to lend an overview of what is happening. However, unlike the generic developers, they know that their work has an inherently military purpose. The company they work for may also have deep links to a particular state in the way that Raytheon Technologies has ties to the United States. For example, they may work alongside military personnel or even be seconded to the military. In such cases, a reasonable programmer in the place of the accused may be positioned well enough to foresee spin-off crimes that could result from the common purpose. However, it does not seem wise to set a presumption either way regarding what this class of developers can or cannot foresee—the range of information available, the degree of interconnectedness, the span of clientele, the breadth of technology, etc., are all too wide.

At the end of the spectrum are the military programmers—personnel in the army, navy, air force, or another military (or paramilitary) branch who work on software development. It is likely that these individuals will indeed know about the common purpose. This might be anything from invading another state’s territory to conducting a “targeted killing.” Knowing the common purpose puts one in a position to determine what crimes may spin-off from that common purpose. For example, if the common purpose was to unlawfully occupy another state’s territory with the aid of autonomous weapons when that territory contained civilians, it seems that violations of IHL would be objectively foreseeable given the current inability of those weapons to comply with the principle of distinction (as they struggle to deal with dynamic factors such as surrender or the plight of a foe who has suddenly become hors de combat).⁹¹ Alternatively, if the common purpose was to carry out a targeted killing via an autonomous weapon, violations of IHL would be objectively foreseeable because autonomous weapons are not yet able to comply with the principle of proportionality (as they cannot properly balance military gain against collateral damage).⁹² For these reasons, it would be sensible to insert a presumption

⁹¹ See Winter, *supra* note 2, at 867–76.

⁹² Winter, *supra* note 3, at 193–94.

into the law to the effect that programmers based within the military should be presumed to foresee the risk of violations of IHL involving autonomous weapons resulting from the pursuit of the common purpose.

Another factor that might make it difficult to determine what the reasonably foreseeable spin-off crimes (of a common plan involving autonomous weapons) might be is the fact that autonomous weapons are an unknown element that add an additional layer of uncertainty into combat. In other words, autonomous weapons are intended to act independently from humans and are thus not as predictable as conventional weapons that are directly controlled by humans on the battlefield—a gun fired, or a missile launched. According to Jain, “the primary barrier for establishing accountability for the harm caused by an AWS is the epistemic uncertainty associated with its conduct, which is a deliberate part of its design features.”⁹³ In particular, this uncertainty is a result of the fact that artificial intelligence (albeit made by humans) has replaced human decision-making in combat. This means that the extent to which any given individual will foresee what might happen on the battlefield is diminished. This will offer software developers a potential escape route from criminal liability as they will be able to argue that the actions of the autonomous weapon were not predetermined and were thus unforeseeable. As Amoroso put it, “autonomy in weapons systems will increase the incidence of cases where the human agent can at best formulate probability assessments as to what the weapon will actually do in the theatre of war. This is likely to create serious hurdles for responsibility ascription.”⁹⁴

In previous articles, the present author has considered how autonomous weapons might function—most recently in the context of the extent to which they might comply with the principle of precaution in attack.⁹⁵ However, a recurring caveat is that there are no examples yet of offensive autonomous weapons being deployed. Consequently, we simply do not know how they might behave in practice. Jain and Amoroso may be correct—the technology may prove to be erratic and may make it difficult to argue that violations of IHL that arise through their use were foreseeable (though instability in itself could be said to present foreseeable risks). However, this instability seems unlikely. In fact, one could argue that the machine learning approach taken to developing autonomous weapons (whereby artificial

⁹³ Neha Jain, *Autonomous Weapons Systems: New Frameworks for Individual*, in *AUTONOMOUS WEAPONS SYSTEMS: LAW, ETHICS, POLICY* 324 (Nehal C. Bhuta et al. eds., 2016).

⁹⁴ Amoroso & Giordano, *supra* note 10, at 222.

⁹⁵ Winter, *supra* note 4.

intelligence is exposed to millions of permutations millions of times to allow it to identify the consequences of different actions) would make their behavior more predictable than that of a human. As Gibney explained when discussing “AlphaGo” (a program designed by United Kingdom-based company DeepMind to play the ancient Chinese strategy game “Go”), the system was able to beat its competitor programs at 99.8% of games and had developed a “conservative” style.⁹⁶ Of course, it is impossible to be sure that autonomous weapons would reproduce this sort of stability. Nonetheless, the current picture of machine learning is not exactly painted in hues of wild unpredictability. If the actions of autonomous weapons on the battlefield are broadly predictable—about as predictable as those of human soldiers—then there would be no great break in foreseeability here. Recall from *Stakić* that allowing convicted criminals to transport prisoners was “to reconcile oneself to the reasonable likelihood” that they would be harmed and that this was so even though no specific harm to any given individual was foreseeable.⁹⁷ The same logic would apply here.⁹⁸

In summary, when it comes to foreseeability, military developers are the most likely to know about the common purpose and be adequately positioned to foresee the chance of spin-off crimes occurring. Defense contractors occupy an ambivalent position, and generic commercial developers have the lowest chance of foreseeing the risks. It will always be necessary for courts to consider the facts of a particular case—and the position of the accused in the greater scheme of events—in detail to determine what would have been reasonably foreseeable. It is unlikely that programmers could argue that autonomous weapons are erratic and that their actions are inherently unforeseeable. Indeed, there is reason to believe they may be a more predictable variable on the battlefield than humans.

D. Criterion 4: Awareness

The fourth criterion that must be satisfied for prosecution under JCE III, and the second limb of mens rea, is that the accused must have been subjectively aware that the spin-off crime was a possible consequence of the common purpose.

⁹⁶ Elizabeth Gibney, *Google AI Algorithm Masters Ancient Game of Go*, NATURE (Jan. 27, 2016), <https://www.nature.com/news/google-ai-algorithm-masters-ancient-game-of-go-1.19234>.

⁹⁷ Prosecutor v. Stakić, Case No. IA-97-24-T, Trial Judgment, ¶ 600 (Int’l. Crim Trib. for the Former Yugoslavia July 31, 2003).

⁹⁸ *Id.*

The subjective awareness test was expressed first by the Appeals Chamber in *Tadić*, which held that responsibility for a crime other than the one agreed upon in the common plan arises because “the accused willingly took [the] risk” that “such a crime might be perpetrated.”⁹⁹ In *Stakić*, the Appeals Chamber rephrased this slightly and required that the accused must have acted “in the awareness” that the spin-off crime might result from the common purpose.¹⁰⁰ In *Ntakirutimana* (before the ICTR), it was held that the accused must have been “aware that such a crime was a possible consequence” of the common purpose.¹⁰¹ The consequence of these judgments is that the “awareness” criterion is, as Guilfoyle put it, “ultimately one of subjective recklessness (*dolus eventualis*).”¹⁰² As Amoroso summarized, under *dolus eventualis*, “participants are deemed responsible . . . for crimes that they did not intend to perpetrate, nor were part of the original plan, solely on the grounds that they foresaw and accepted the possibility that those crimes would have been committed.”¹⁰³

While the theory of *dolus eventualis* is relatively simple, it is not without its controversy. Badar criticizes the adoption of *dolus eventualis* by JCE III. In particular, he notes that “if the accused had actually participated in crimes . . . as an aider or abettor they would arguably have an increased chance of acquittal, as the Prosecution would [need to] prove . . . that the accused knew that the principal perpetrator had the state of mind required for the crime at issue.”¹⁰⁴ In other words, Badar is arguing that it is perverse that the higher level of *culpa* required for aiders and abettors (i.e., knowledge of the spin-off crime) makes it easier for them to defend themselves than those accused under JCE III (who need only be aware of the *risk* of the spin-off crime occurring). However, Badar’s argument seems peculiar. Aiding and abetting is a more serious crime than JCE III. As such, it is only natural that its threshold for *mens rea* is higher. It is true that one consequence of this is that those accused of aiding and abetting have a better chance of being acquitted as the

⁹⁹ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 228 (Int’l Crim. Trib. for the Former Yugoslavia July 15, 1999).

¹⁰⁰ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 92 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

¹⁰¹ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17-A, Appeal Judgment, ¶ 467 (Int’l. Crim Trib. for Rwanda Dec. 13, 2004).

¹⁰² GUILFOYLE, *supra* note 47, ¶ 12.6.4.

¹⁰³ Amoroso & Giordano, *supra* note 10, at 220.

¹⁰⁴ Badar, *supra* note 37, at 301.

prosecution may fail to get their case over the evidential bar. However, every criminal justice system operates this way—more serious offenses carry higher thresholds. For example, in England and Wales, “negligence” used to be used as the main mens rea test for certain driving offenses, whereas murder required, and continues to require, “intent.”¹⁰⁵ By Badar’s logic, negligent drivers were unfairly treated by the criminal law because it was relatively easy for the prosecution to satisfy the low threshold of negligence, whereas murderers had an easier lot because it was harder for the prosecution to prove intent. This is an obtuse way of looking at things. Rather, we must acknowledge that the vastly different punishments in place for miscreant drivers on the one hand and murderers on the other, justifies the difference in mens rea thresholds and the attendant evidential burdens.

Amoroso also complains about the adoption of the *dolus eventualis* standard. For her, it is “very close to a criterion of ‘guilt by association,’ in blatant contrast with the principle of culpability.”¹⁰⁶ However, Amoroso’s criticism is not much more convincing than Badar’s. Her main point of contention is that “the mere acceptance of the risk of civilian casualties” is not recognized as unlawful under IHL.¹⁰⁷ Of course, this is true. However, awareness of the risk of a crime is only one of the requirements of JCE III. The other limbs, especially the common purpose (discussed above) and the decision to continue to act despite the risks (discussed below), add texture to the analysis and can help to paint a picture of culpa. More broadly, while one can be sympathetic to Amoroso’s point and recognize the dangers of lowering thresholds for criminal liability too far, one must also recall that, in practice, the enforcement of international law against individuals has been too weak, not too strong. One need only consider the low number of ICC convictions—ten according to the Court itself—for confirmation of that.¹⁰⁸ We need to find ways to strengthen individual criminal accountability rather than undermine it. Otherwise, we tip the odds of those accused of war crimes even further in their favor and at the direct expense of those civilians who have found themselves on the other side of the equation. Further, as we will see below, the need for more robust international criminal law with fewer gaps will become all the more important in the future. This is because the rise of new technologies will prompt further diffusion of responsibility and additional complexities in identifying who has made what contribution to which crime.

¹⁰⁵ JANET LOVELESS, COMPLETE CRIMINAL LAW: TEXT, CASES AND MATERIALS 89, 129 (2020).

¹⁰⁶ Amoroso & Giordano, *supra* note 10, at 220.

¹⁰⁷ *Id.* at 221–22.

¹⁰⁸ *About*, INT’L CRIM. CT., <https://bit.ly/38mvQ1k> (last visited Sept. 6, 2021).

In short, while the adoption of subjective awareness (based on *dolus eventualis*) by JCE III is not without its critics, the adoption is warranted. It is set at an appropriately low threshold to catch most accused, while still allowing those genuinely unaware of the risk of a spin-off crime to secure an acquittal. Furthermore, the other JCE III criteria provide context and can help to build a picture of *culpa*. Thus, subjective awareness should continue to apply when assessing the guilt of common purpose participants.

Assuming the above is true, it presents us with the question of how a criterion based on the subjective awareness of a spin-off crime might apply to software developers who have contributed to the development of an autonomous weapon that was subsequently involved in the violation of IHL. Would it be possible to show that a programmer was “aware” of the risk of perpetrating a crime beyond the common purpose? For context, we can consider the position in *Stakić*. There, the accused was the Vice-President of the Prijedor Municipal Assembly and, simultaneously, President of the Crisis Staff in Prijedor. The facts established that he actively supported operations of Serbian militias in the area and was in a position to directly observe the consequences. On this basis, the Appeals Chamber agreed that he “consented to the removal of Muslims from Prijedor by whatever means necessary.”¹⁰⁹ In other words, the ICTY had sufficient information on the accused’s actions to conclude that he was personally aware of what was happening and of the risk that spin-off crimes would occur as part of the broader endeavor.

Just as in *Stakić*, in order to gauge whether a software developer was aware of the risk of the spin-off crime, one must consider their individual contribution to the common purpose—regardless of whether they work for the military, a defense contractor, or a generic software firm. On this point, Ambos notes that “while some judgments . . . try to take into account the role and function of the accused in the enterprise, there still exists a tendency to [conflate them].”¹¹⁰ However, courts should be careful not to conflate individuals’ actions in this context as it is a subjective test and depends precisely on individual awareness. Many programmers will spend their time developing “non-critical” functions (those which cannot kill or injure) such as navigation, image collection, telemetry reporting, fuel monitoring, communications, and other ancillary systems. For programmers involved in this sort of work, it would be hard in most cases to impute to them awareness of the risk of a spin-off crime as

¹⁰⁹ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 92 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

¹¹⁰ Ambos, *supra* note 60, at 173.

their personal actions were—unlike those of Stakić—inherently innocuous. Individuals will likely be unaware of any incorporation of their systems into autonomous weapons and, thus, their potential for lethal military and civilian applications. Consequently, it seems that many of the cases that might potentially be raised against software developers would fall at this “awareness” hurdle. This would remedy the concern Amoroso raised above that JCE III is over-inclusive.¹¹¹

Still, a portion of software developers will have worked on “critical” systems (those which can kill or injure) and whose only possible function is to be incorporated into an autonomous weapon. Developers with these roles are self-evidently much more likely to be subjectively aware of the risk of a further crime spinning off from the broader common purpose. More than that, they are likely to be aware of any deficiencies in an autonomous weapon that may cause it to violate IHL. For example, there may be a programmer who is working on ways to enable software to identify military uniforms so that the wearer can be marked as a potential combatant for the purposes of distinction. Great strides in “machine recognition” have been made recently in a non-military context, with artificial intelligence gaining the ability to recognize images such as human faces,¹¹² the written word,¹¹³ cancerous growths,¹¹⁴ retail products¹¹⁵ (for online shopping), and even weapons¹¹⁶ in a security guard context.¹¹⁷ Still, the systems are not perfect, and minor changes in appearance could throw it off. For example, an autonomous agricultural weed remover, RoboWeedMaps, can be confused if a beetle has eaten a leaf or if the temperature has caused the weed to change color.¹¹⁸ When it comes to the Google Health system designed to detect cancers from scanner imagery, humans remain slightly better at detecting “in situ” cancers than artificial intelligence, which is better at detecting “invasive” cancers.¹¹⁹ These sorts of issues will arise in the context of

¹¹¹ See Amoroso & Giordano, *supra* note 10, at 220.

¹¹² Winter, *supra* note 2, at 863.

¹¹³ *Id.* at 866.

¹¹⁴ *Id.* at 866–67.

¹¹⁵ *Id.* at 867.

¹¹⁶ *Id.* at 865.

¹¹⁷ *Id.*

¹¹⁸ *Id.* at 874.

¹¹⁹ *Id.* at 867.

battlefield distinction, and the programmers who have worked on these visual systems will be aware of the shortcomings.

Similarly, another programmer may be working on ways to incorporate a “collateral damage estimation methodology” (CDEM) into code so that a future artificial intelligence system could use it to make proportionality decisions on the battlefield by weighing the expected military gain against the anticipated collateral damage.¹²⁰ CDEMs take into account matters such as the timing of attacks, attack vectors, the effective radius of weapons, the strength of buildings, and the number of civilians in the area of operation in order to make sure that assessments about gain and harm go to the appropriate level of authority.¹²¹ In the future, they are likely to enable the full automation of such assessments, with Schmitt noting that “[t]here is no question that autonomous weapon systems could be programmed to perform CDEM-like analyses.”¹²² However, there are still many difficulties in this area. For example, the United States has admitted that the “operational environment, weapon reliability, and fidelity of intelligence data” can impact the reliability of the results.¹²³ Additionally, professionals concede that “the science is inherently limited by the quantity and reliability of collected and analyzed weapons effects data, weapon delivery uncertainties, and target information.”¹²⁴ Programmers working on such systems are at the sharp end of dealing with these limitations and thus are likely to be aware that the use of autonomous weapons employing a CDEM might trigger spin-off crimes beyond the ambit of the common purpose.

In summary, courts need to consider an individual’s personal involvement in a common purpose to determine whether they were aware of the risk of spin-off crimes. In the context of software developers and autonomous weapons, this means identifying the systems to which individuals contributed. If they helped to develop a non-critical system, it would be very difficult to show that the requisite awareness existed. When it comes to critical systems, given the current limitations inherent in the technology, this will be much easier to establish.

¹²⁰ Winter, *supra* note 4, at 257.

¹²¹ *Id.*

¹²² Schmitt, *supra* note 15, at 20.

¹²³ U.S. DEP’T OF DEF., CJCSI 3160.01A, NO-STRIKE AND THE COLLATERAL DAMAGE ESTIMATION METHODOLOGY (2012), <https://bit.ly/2Zzzggj>.

¹²⁴ *Id.*

E. Criterion 5: Action in Furtherance of the Common Purpose

The fifth and final criterion that must be satisfied for prosecution under JCE III is that, as the ICTY Appeals Chamber held in *Stakić*, the accused “acted in furtherance of the common purpose.”¹²⁵ This requirement was confirmed by the ICTR in *Ntakirutimana*, where it was held that JCE III requires the accused to “participate in and further the common criminal purpose of a group.”¹²⁶ For the purposes of the present analysis, and based on a synthesis of case law in this area, three facets of the requirement to act in furtherance of a common purpose are particularly relevant when it comes to the accountability of software developers for violations of IHL involving autonomous weapons. First, the test captures legally neutral forms of participation; second, the threshold for the level of participation required is set low; finally, evidence must be adduced to establish participation as a matter of fact. Each facet will be considered in turn.

Regarding the criminalization of otherwise neutral acts, it was held in *Tadić* that “participation need not involve commission of a specific crime . . . (for example, murder, extermination, torture, rape, etc.), but may take the form of assistance in, or contribution to, the execution of the common plan or purpose.”¹²⁷ This language was echoed verbatim by the ICTR in *Ntakirutimana*.¹²⁸ This is helpful from a prosecutorial point of view because it precludes the defense that, when taken in isolation, the defendant’s actions were lawful. When it comes to software developers, as we saw above, their activities will almost always fall into this category. Developing code is not an inherently illegal activity. Even in the most egregious potential scenario where the programmer works within a state’s military (such that the spin-off crime is likely to be objectively foreseeable) and where they are working on critical systems (such that they are likely to be subjectively aware of the spin-off crime), recall that IHL does not render weapons development illegal per se and merely requires review of their compatibility with international law.¹²⁹ JCE III

¹²⁵ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 87 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

¹²⁶ Prosecutor v. Ntakirutimana, Case No. ICTR-96-17-A, Appeal Judgment, ¶ 467 (Int’l Crim. Trib. for Rwanda Dec. 13, 2004).

¹²⁷ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 227 (Int’l Crim. Trib. for the Former Yugoslavia July 15, 1999).

¹²⁸ *Ntakirutimana*, ICTR-96-17-A ¶ 466.

¹²⁹ Protocol Additional, *supra* note 57, art. XXXVI.

effectively circumvents this stumbling block by criminalizing neutral behavior, which nonetheless contributes to a broader criminal purpose.

This facet of JCE III has proven to be controversial because it risks diluting the individual's level of involvement with the crime too much. For example, Amoroso complains that the test is satisfied "regardless of whether [the individual programmers] were physically involved in the commission of the crime(s) and of whether their participation took place through a criminal or a legally neutral act."¹³⁰ Indeed, one can imagine a situation where a military developer has contributed to the code used to teach a piece of artificial intelligence how to recognize enemy combatants, which is subsequently incorporated into an autonomous weapon that kills a civilian in the field. The individual may feel they have done nothing wrong, but JCE III would take a different view. However, if courts were to accept Amoroso's criticism and return us to a position where only programmers who directly violated IHL are held accountable, this would be a backward step. It would leave us in the position we were in before the ICTY and ICTR came into existence by re-opening the accountability gap—with particular consequences for situations where programming made by software developers displaces the decision-making power of military commanders. This cannot be correct. As Ambos noted, the whole point of JCE III is that it endeavors to take into account "the collective, widespread and systematic context of such crimes and, thus, helps to overcome the typical difficulty in proving the . . . contributions of individual participants."¹³¹

Surely, by proper construction—and bearing in mind the object and purpose of the law (including the ICTY and ICTR Statutes) as required by the Vienna Convention—it was not the intention of international law to tolerate such gaps.¹³² No one, including software developers, should be able to escape the reach of the law. As the Appeals Chamber stated in *Tadić*, "an interpretation of the Statute based on its object and purpose leads to the conclusion that . . . responsibility for serious violations of international humanitarian law is not limited merely to those who actually carry out the actus reus of the enumerated crimes but appears to extend also to other offenders."¹³³ Further, "all those who have engaged in serious violations of international humanitarian law, whatever the manner in which they may have

¹³⁰ Amoroso & Giordano, *supra* note 10, at 220.

¹³¹ Ambos, *supra* note 60, at 159–60.

¹³² Vienna Convention on the Law of Treaties, art. XXXI(1), May 23, 1969, 1155 U.N.T.S. 331.

¹³³ Prosecutor v. Tadić, Case No. IT-94-1-A, Appeal Judgment, ¶ 189 (Int'l Crim. Trib. for the Former Yugoslavia July 15, 1999).

perpetrated, or participated in the perpetration of those violations, must be brought to justice.”¹³⁴ It justified its interpretation on the basis that:

[a]lthough only some members of the group may physically perpetrate the criminal act (murder, extermination, wanton destruction of cities, towns or villages, etc.), the participation and contribution of the other members of the group is often vital in facilitating the commission of the offense in question. It follows that the moral gravity of such participation is often no less—or indeed no different—from that of those actually carrying out the acts in question.¹³⁵

This logic applies to software developers just as much as it applies to politicians like Tadić, Stakić, and their ilk. If they satisfy the other criteria of JCE III, the legality of their actions when, taken in artificial isolation, should offer no defense.

The second interesting facet of the requirement that the accused “acted in furtherance” of the common purpose is the threshold required for participation. On this point, the position of the international tribunals has evolved somewhat over time. In *Tadić*, the rule was simply that it was not necessary for the accused’s participation to “be a *sine qua non*, or that the offence would not have occurred but for his participation.”¹³⁶ In other words, the accused’s actions did not need to be the “legal cause” of the spin-off crime, and the existence of intervening acts by third parties would not break the chain of causation. In *Kvočka*, the Appeals Chamber considered the level of participation needed in more detail and held that, although the accused’s participation or contribution need not be “substantial,” a substantial contribution might help prove the mental element required.¹³⁷ In *Mpambara*, the court stated explicitly that “there is no minimum threshold of significance or importance” and “the actus reus [of JCE] may be satisfied by any participation, no matter how insignificant.”¹³⁸ Later cases revealed *Mpambara* to be the low-water mark. In *Milutinovic*, the Trial Chamber asserted that while the contribution “need not be

¹³⁴ *Id.* ¶ 190.

¹³⁵ *Id.* ¶ 191.

¹³⁶ *Id.* ¶ 199.

¹³⁷ Prosecutor v. Kvočka, Case No. ICTY-98-30/1-A, Appeal Judgment, ¶¶ 97, 104, 187 (Int’l Crim. Trib. for the Former Yugoslavia Feb. 28, 2005).

¹³⁸ Prosecutor v. Mpambara, Case No. ICTR-01-65-T, Trial Judgment, ¶¶ 13–14 (Int’l Crim. Trib. for Rwanda Sept. 11, 2006).

necessary or substantial, it should at least be significant.”¹³⁹ This seems to be a sensible compromise as it ensures that even those who have made a modest personal contribution to the common purpose will be held accountable for their actions while retaining, in effect, a de minimus rule whereby negligible contributions are allowed to pass by.

Assuming the “significance” test is correct as a matter of doctrine, the challenge then becomes determining what it means in practical terms. Helpfully, the Trial Chamber in *Milutinovic* shed some light on this issue. In essence, the action of the accused should contribute to the efficiency, effectiveness, and smooth running of the plan with relevant factors including: “the functions performed by the accused and his efficiency in performing them, and any efforts made by the accused to impede the efficient functioning of the joint criminal enterprise.”¹⁴⁰ Further, “an accused’s leadership status and approving silence . . . militate in favor of a finding that his participation was significant.”¹⁴¹ Guilfoyle has since observed that a “very low standard applies to leaders (approving silence may be a significant contribution), but a higher level of contribution will be required before a readily replaceable subordinate is considered a participant in a JCE.”¹⁴² It is probably fair to say that most software developers will be “readily replaceable” in this context and so their failure to speak out against the common purpose or any other omission will not be enough—it will be necessary to show a positive, if limited, contribution on their part. This leads us neatly to the final aspect of the “acted in furtherance” criterion—evidence.

The final facet of the requirement that the accused “acted in furtherance” of the common purpose concerns evidence. It is all very well to say that a software developer must have made a significant contribution to the common purpose in the form of positive action, but stating the test and proving that it is satisfied in practice are two different things. Again, we can return to *Stakić* for context. In that case, in the course of establishing the facts, the Trial Chamber held that “as the highest representative of the civilian authorities, Dr. Stakić played a crucial role in the coordinated cooperation with the police and army in furtherance of the plan to

¹³⁹ Prosecutor v. Milutinovic, Case No. ICTY-05-87-T, Trial Judgment, ¶ 104 (Int’l Crim. Trib. for the Former Yugoslavia Feb. 26, 2009) (quoting Prosecutor v. Brđanin, Case No. IT-99-36-A, Appeal Judgment, ¶ 430 (Int’l Crim. Trib. for the Former Yugoslavia Apr. 3, 2007)).

¹⁴⁰ *Id.* ¶ 105.

¹⁴¹ *Id.*

¹⁴² GUILFOYLE, *supra* note 47, ¶ 12.6.3.

establish a Serbian municipality in Prijedor.”¹⁴³ The Trial Chamber also noted that Stakić “actively participated in and threw the full support of the civilian authorities behind the decision to establish the infamous Keraterm, Omarska, and Trnopolje camps”¹⁴⁴ and that he was “one of the main actors in the persecutorial campaign.”¹⁴⁵ The Appeal Chamber was content to adopt these factual findings and so agreed that “the Appellant acted in furtherance of the common purpose and played an important role in it.”¹⁴⁶

Of course, demonstrating that a particular software developer was responsible for furthering a common purpose will not be quite as simple as it was in *Stakić*. This is because software packages, especially complex ones, can be developed by hundreds or even thousands of people. More particularly, as Nissenbaum noted, software is typically created by individuals working for large organizations and then passed on to new, similarly large organizations or governments to be further honed by more people before finally being utilized.¹⁴⁷ This makes the existence and extent of any given individual’s involvement difficult to identify.¹⁴⁸ To compound the problems still further, software is rarely developed as one monolithic structure. Rather, many packages are often stitched together to form a final product, often with the different components developed in different countries, at different times, and sometimes even for completely different end uses.¹⁴⁹ Added to the mix is the fact that software is then often paired with hardware which, again, is likely to have been developed by different people in a different setting.¹⁵⁰ This complex account of software development will be replicated in the context of autonomous weapons. Thus, Heyns observed that the production of autonomous weapons “will invariably involve a vast number of people,”¹⁵¹ and McFarland opined that “a sophisticated

¹⁴³ Prosecutor v. Stakić, Case No. IT-97-24-T, Trial Judgment, ¶ 822 (Int’l Crim. Trib. for the Former Yugoslavia July 31, 2003).

¹⁴⁴ *Id.* ¶ 595.

¹⁴⁵ *Id.* ¶ 823.

¹⁴⁶ Prosecutor v. Stakić, Case No. IT-97-24-A, Appeal Judgment, ¶ 76 (Int’l Crim. Trib. for the Former Yugoslavia Mar. 22, 2006).

¹⁴⁷ Helen Nissenbaum, *Accountability in a Computerized Society*, 2 SCI. ENG’G ETHICS 25, 29 (1996).

¹⁴⁸ *Id.* at 30.

¹⁴⁹ *Id.* at 29–30.

¹⁵⁰ *Id.* at 30.

¹⁵¹ Heyns, *supra* note 5, ¶ 77.

autonomous weapons system will not be developed by a single individual, but by many teams of developers in many organizations.”¹⁵² This “group effort” nature of autonomous weapon software design will result in a permutation of the “many hands problem” first identified by Thompson, who noted its existence in the context of decision-making by public officials. Thompson observed that so many people are involved in government processes that it can be almost impossible to identify whether a specific individual has contributed to negative outcomes.¹⁵³ In turn, this has led to fears that determining who has contributed to an autonomous weapon’s development “may simply be too difficult.”¹⁵⁴

However, there is a potential solution to the problem of not knowing who has made what contribution to a piece of software—what might be known as the “opacity problem.” Marchant has highlighted that an additional layer of technology could be added to autonomous weapons to allow them to produce a precise reconstruction of what occurred during lethal operations through the incorporation of black-box-style monitoring devices and the mandatory review of the telemetry collected.¹⁵⁵ This would indeed be useful and, arguably, the approach could be taken a step further by requiring that autonomous weapons come with lists, kept by the relevant states, of the names of everyone who made a contribution to their development and what the nature of that contribution was—akin to the end credits of a film. Where parts of the software have been purchased from private corporations, whether defense contractors or generic commercial developers, those companies should be required to retain similar lists for their records. The combination of these two innovations would allow future courts to determine whether an autonomous weapon did indeed violate IHL and, if it did, to match up that machine with those who contributed to its development. For each of those people, assuming the four criteria above are also satisfied, they will be guilty under JCE III if they made a significant contribution to the software.

In summary, to satisfy the “action in furtherance” criterion of JCE III, the accused must have personally participated in the common purpose. This participation may have taken the form of a legally neutral act. Further, it is enough if the accused has made a “significant” contribution to the common purpose. However, it will be difficult for the prosecution to provide evidence of the participation of individual

¹⁵² McFarland & McCormack, *supra* note 19, at 384.

¹⁵³ Dennis F. Thompson, *Moral Responsibility of Public Officials: The Problem of Many Hands*, 74 AM. POL. SCI. REV. 905, 905 (1980).

¹⁵⁴ McFarland & McCormack, *supra* note 19, at 384.

¹⁵⁵ Gary E. Marchant et al., *International Governance of Autonomous Military Robots*, 12 COLUM. SCI. & TECH. L. REV. 272 (2011).

programmers in the coding of any given system. The cooperation of states and corporations would be needed here to ensure that autonomous weapons are outfitted with black boxes to help the court to determine their involvement in the spin-off crime, and these boxes would need to have associated name lists of all those who participated in the production of the code behind the operation of the machine. While states are unlikely to adopt these measures unilaterally, they may be open to them on a multilateral basis.

CONCLUSION

The rise of autonomous weapons will give software developers an increasingly important role in warfare because the programs they produce will gradually replace human decision-making on the battlefield. To ensure that no responsibility gap arises, it is important that we consider how these individuals might be held to account when their work results in violations of IHL. Individual accountability is unlikely to be of much utility as the requisite *mens rea* will be lacking in most if not all cases. For example, software developers are unlikely to intentionally code a machine to target civilians or recklessly program a collateral damage estimation model. Likewise, the aiding and abetting doctrine is unlikely to offer much recourse as the developer's actions must directly support a particular crime, whereas one anticipates most violations of IHL involving autonomous weapons will arise far beyond the reckoning of the software developer. JCE III has the potential to fill the lacuna as it enables conviction for a contribution to a common criminal purpose which led, in a foreseeable and foreseen way, to a violation of IHL and which operates even where the actions of the accused are not inherently criminal when taken in isolation. While controversial, JCE III has a solid foundation in international law and is settled in jurisprudence, having been subjected to detailed and prolonged consideration by both the ICTY and ICTR. That said, applying the doctrine to software developers and autonomous weapons would be a novel use of the concept and would represent a departure from its hitherto prevalent function of meting out justice to participants in ethnic cleansing operations that later morphed into genocide.

This Article has distilled JCE III into five component parts and has shown the extent to which software developers might satisfy each of them. Regarding the "common purpose" criterion, the Article established that a plan to perpetuate any war crime will suffice, and that although there must be an affiliation between the accused and a plurality of persons, that group need not have any specific structure, and the accused need only be loosely integrated into it. Regarding the requirement for a "crime beyond the common purpose" (or "spin-off crime"), again, any war crime will suffice. In the Yugoslavian and Rwandan cases, spin-off crimes included, for example, genocide, murder, assault, and sexual violence—they were all beyond the ambit of the original plan, but they were consequences nevertheless. In the context of autonomous weapons, the spin-off crimes are most likely to be violations

of principles such as distinction and proportionality because they are highly context-sensitive concepts that artificial intelligence has struggled to master. Regarding the “foreseeability” criterion, it is critical that the reasonable person in the accused’s position would have known about the existence of the common purpose and about enough of its detail to foresee the risk of any spin-off crimes. This Article has suggested that adopting a three-tiered approach to software developers would help here—with military personnel, defense contractors, and generic commercial developers as the categories. Military personnel should be presumed to have adequate knowledge of the common purpose to foresee spin-off crimes, while generic commercial developers should be presumed not to have adequate knowledge. In the middle, the defense contractors should not labor under, or benefit from, a presumption either way. Regarding the “awareness” criterion, courts will need to consider an individual’s personal involvement in the common purpose to determine whether they were aware of the risk of spin-off crimes. For software developers, this means determining what the individual’s coding work involved—it will be much easier to establish awareness if they worked on critical systems (e.g., weapons targeting) than if they worked on non-critical systems (e.g., navigation). Finally, the accused must have “acted in furtherance” of the common purpose in the sense of making a significant contribution to it which need not have been inherently illegal. Many software developers will satisfy this final criterion, but the problem will be proving that this is the case owing to the “many hands” problem and the opacity of software development. State and corporate cooperation in the form of implementing black box technology and in maintaining lists of contributors to each system would be key to ensuring that justice does not fall at this final hurdle.

Ultimately, JCE III offers a viable mechanism by which to hold software developers to account for the operations of autonomous weapons. It is an established doctrine that offers a balanced and nuanced approach to a difficult problem. It will catch those developers who had sufficient appreciation of the big picture to perceive the risk of spin-off violations of international law and who nevertheless continued to contribute to the common enterprise. It will leave in peace those developers who could not, or did not, foresee the risk of the violation occurring or whose contribution was so minor as to be *de minimis*. As such, the adoption of JCE III will ensure that justice does not evaporate if or when machines replace humans on the battlefield.