PUBLIC PERCEPTION, APPROVAL OF AND BLAME ATTRIBUTION IN USING MILITARY DRONES:

A GATEWAY INTO PUBLIC OPINION ON AUTONOMOUS WEAPON SYSTEMS

A Senior Honor Thesis Presented to

the Faculty of the Department of Political Science

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts in Political Science

By

Justin Truong

December 2018

PUBLIC PERCEPTION, APPROVAL OF AND BLAME ATTRIBUTION IN USING MILITARY DRONES:

A GATEWAY INTO PUBLIC OPINION ON AUTONOMOUS WEAPON SYSTEMS

____

Justin Truong

APPROVED:

____

Dr. Zachary Zwald, Thesis Director

____

Dr. Ryan Kennedy

____

Dr. Benjamin Rayder

____

Dr. Antonio D. Tillis, Dean of the College of Liberal Arts & Sciences

Acknowledgements

I would like to sincerely thank my advisor and thesis director, Dr. Zachary Zwald, for his instrumental guidance and support throughout the course of this thesis; Dr. Zwald previously served as the faculty mentor for my 2017 Summer Undergraduate Research Fellowship project, “Chinese Soft Power: Forms & Rationale”, and it has always been a blessing to have him as a Political Science professor. I am also incredibly grateful to Dr. Ryan Kennedy, whose knowledge and invaluable support enabled me to construct and implement the survey study. I am indebted to Dr. Beverly Barrett at St. Thomas University and Ms. Alison Rana at ExxonMobil for always motivating me toward the completion of this thesis. Finally, to my family and my friends Annie Tran, Sarah Nguyen, Monica Can, Danielle Niangar, Victor Yau and Eddie Pham, I would like to say a heartfelt thank you for their enduring encouragement.

PUBLIC PERCEPTION, APPROVAL OF AND BLAME ATTRIBUTION IN USING MILITARY DRONES:

A GATEWAY INTO PUBLIC OPINION ON AUTONOMOUS WEAPON SYSTEMS

A Senior Honor Thesis Presented to

the Faculty of the Department of Political Science

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts in Political Science

By

Justin Truong

December 2018

Abstract

Within the past two decades, technological advances and a growing public policy priority to reduce human casualty in the battlefield have pushed militaries around the world to develop increasingly automated capabilities that are characterized by a gradual decline in human involvement. Autonomous weapon systems (AWSs), which include Legal autonomous weapons (LAWs), are considered the next transformational stage in military technology, yet public understanding of the far-reaching implications of using these war-fighting machines remains limited, especially its material impact on interstate dynamics. Land-based, seaborne and airborne unmanned vehicles, colloquially known as drones, represent the machine-learning military technology that is closest to AWSs and serve as the basis for which AWSs will derive their development from. While true AWSs have not yet been developed and deployed, this thesis seeks to understand the public's current perception of military drones and how it affects approval for the use of drones in combat and blame attribution in scenarios of errors, such as unintended collateral damage. This is achieved by a combination of examining evolving scholarly debates on public opinion and the legal accountability of employing armed drones and a survey procedure to examine how the public views drones within the context of machine autonomy versus human control, how approving it is of using drones in war and finally how it attributes blame when presented with an erroneous outcome. Analysis of both the existing literature and the survey data suggests the public still struggles to comprehend the capabilities of drones along a wide spectrum of autonomy. Furthermore, the study's findings underpin the current scholarly position that considerable support for drone usage exists when the subject is framed in a vacuum, without including contextual information that truly characterizes the

reality of that usage. However, when presented with that contextual information, the option of deploying a human combatant remains preferable to the public due to an aversion to collateral damage even when there is high military utility. Finally, the public tends to find the human element at blame when it is asked to assess a collateral damage, regardless of the level of automation or autonomy involved. These findings are all emblematic of a degree of distrust in AWSs and incoherent, underdeveloped legal thinking within the public on the subject of accountability, which promise to complicate not just the rules of war and international legal regimes but also the interstate dynamics among AWS-wielding nations when fully autonomous warfighting platforms become a full-fledged reality.

Table of Contents

Chapter 1: Introduction

Automation has been one of the priorities of improving warfighting technologies throughout the course of military history. Automation is credited with improving both precision and destructibility, attributes that are important to military prowess. As society becomes increasingly averse to the cost of human life in warfare, the ability to conduct combat operations with maximum precision and efficiency while incurring minimal loss of life and other forms of collateral damage has become a political priority for many nations. This is especially true in substantive democracies, where the consequences of political fallout and public perception place a considerable constraint on military decision-making that is likely to result in military casualties as well as civilian losses on other belligerents' sides. To circumvent the disadvantages of having human soldiers directly in the line of fire on the field and to increase accuracy at the same time, there have been efforts undertaken by governments and private entities across the globe to reduce human involvement in warfighting technology. This is achieved by increasingly delegating tasks to machines with the ultimate goal of having them perform the bulk of the necessary tasks required to achieve objectives at the tactical, operational and even strategic levels of warfare with minimal human guidance. The development of these Autonomous Weapon Systems (AWSs) comprises the quintessential element of the hypothesized "hyperwar" concept that was first brought to policymaking discussions by General John Allen, former Commander of the U.S. Central Command (CENTCOM) and the International Security Assistance Force (ISAF). One of the most pressing goals in the path of developing machine autonomy for military applications is the removal of humans from the immediate combat environment, allowing vehicles to be operated from a safe distance while still performing the vast majority of physical tasks. This capability has been achieved with unmanned, remotely-operated technology, which has progressed significantly in recent years and saw expanded applications in the battlefield with promises of even more advanced capabilities in the near term. As such, drone technology represents the bridge between today's still human-operated-from-a-distance warfighting platforms and the fully autonomous, human-free weapons of the future that will define hyperwarfare and carry with them the potential to alter interstate relations to a great extent.

As with most technology in the digital age, there is a general understanding among the public that higher automation leads to increased accuracy and productivity, hence efficiency, in task performance. This represents the concept of trust in machines. Human's trust in machine can be observed at any given time, from our dependence on calculators for sophisticated mathematical operations to auto-pilot programs that are able to achieve complex tasks such as landing a jumbo jet. Most of today's technology represents some form and extent of autonomy, and rapid developments in artificial intelligence mean that the age of ubiquitous fully-autonomous machines is not a far-fetched possibility. Self-driving cars, representing the next phase in advanced civilian applications of autonomous technology, are already being tested with the expectation that they would enter widespread commercial usage by 2040, a mere 21 years from now. However, when confronted with the military application of autonomous technology, such as a simple killer machine such as the Terminator robot in the popular science fiction series of the same name, the public likely shows a severe lack of trust in or aversion to using them, a tendency also noticed in debates revolving the commercial use of self-driving cars but to a lesser extent. The concerns with self-driving cars draw on the problematic accountability issues that would arise in the aftermath of an adverse traffic event, such as a fatal collision. Who is liable the responsibility for the crash? A similar but more profound, deep-seated concern underpins the problematic accountability and dynamics of machine trust involved with the use of drones and AWSs because these technologies would not only cause fatalities and destruction, as can self-driving cars, but are intentionally built to do so efficiently, accurately and ultimately without human control. Similar to the self-driving cars' hypothetical controversy, the presence of a trust deficit in drones is thought to originate from the wariness and uncertainty in accounting for the legal responsibilities and operational unpredictability associated with their use. A product of humans with known inherent prejudices, algorithms built into self-driving cars and drones armed to the teeth may very well reflect those prejudices and, as a consequence, lead to undesirable preferences because of their legal or ethical ramifications. Furthermore, since it is almost impossible to fully account for every single variable in the real-world operating environment, when released into untested circumstances, autonomous systems may pursue unforeseen courses of action with equal legal and ethical distress. These concerns about machine bias become amplified when AWSs are brought into consideration because of their intentional lethality and poorly-understood, underdeveloped position in the current rules of war. Who is to blame when a killer robot commits an unintended course of action that results in collateral damage? How can human operators maintain a meaningful level of control to minimize unintended consequences and is this even possible without defeating the original purpose of developing AWSs? Moreover, on a much grander scale, if and when killer robots are deployed at war between or among nations, how will the civilian and military leaderships interpret the actions of hostile platforms and how will they respond to erroneous actions that result in highly destructive consequences unintended for by the other side? These questions become much more intriguing when one considers the scenario of two or more warring factions' autonomous capabilities being pitted against one another.

While fully-autonomous warfighting machines still remain in various stages of research and development, the looming prospects of hyperwarfare demand scrutiny on how well this technology and its implications are understood and perceived, especially from the perspective of the public-at-large since public opinion has a determinant effect on national security and foreign policy. To this extent, today's drones are useful for that type of query because they represent the fundamental idea behind warfighting autonomous system. The public's comprehension of drones and its effect on approval of use and blame attribution is thus an important lever into the larger debate of how future AWSs are poised to materially alter interstate relationships and the dynamics of warfare. Although this paper principally addresses drones from the perspective of the public, it still seeks to provide useful, initial insights in determining how interstate relations may change as the realities of AWSs become more pronounced. A number of authoritative figures have written extensively on drones and to a lesser extent, AWSs. Paul Scharre, a noted scholar on AWSs, writes that "increasing automation in military systems poses both benefits and risks". The development of this transformational military technology, therefore, necessitates further studies on how the public perceives AWSs and attributes responsibility in faulty scenarios, such as an unforeseen course of action determined by continually-learning machine algorithms that results in unintended collateral damage. However, because AWSs still remain in the development phase and the general concept of AWSs itself represents a wide spectrum of autonomous capabilities, it is important to examine the public's ability to distinguish between military technologies driven primarily by machine and remote warfare technologies that are human-operated from a distance with some degree of machine autonomy, such as drones that are used widely today across the globe, of which unmanned aerial vehicles (UAVs) are most common. Because of the lack of a human operator on board, they are often referred to interchangeably in mass media as "autonomous weapons" or "robots", which can be misleading but presents an intriguing case study of how the public sees machine autonomy versus human control in drone technology, and whether this has any effects on attributing responsibility and blame for applying this technology militarily.

Chapter 2: AWS and the public

Because AWSs are thought to revolutionize future warfare, their transformational capabilities and subsequent impact cannot be fully understood without placing them within the context of how weapons have progressed along the lines of increasing delegation of tasks to the machine with the ultimate milestone being the delegation of the authority to apply violent or lethal force. The history of military weaponry can be described as a progression of accuracy, lethality and automation, that saw these attributes became increasingly interconnected. Accuracy and lethality were not necessarily an outcome of automation in the early ages of warfare when a fighter's skills and prowess played a determinant role in close-range combat. However, as warfare grew in both scale and sophistication, as befitting the growth of civilization and its technological capabilities, combat accuracy and lethality increasingly depend less on personal acumen and more on the ability to automate certain aspects of a weapon's operating mechanics. The Gatling gun, used extensively during the 19th and early 20th centuries, is a powerful example of how automation can make warfare more a lot more brutal by making the firing-reloading process more integrated and less labor-intensive. In other words, these advancements in military weaponry resulted in less labor put in by the operator (soldier) and consequently a reduced role of the human. Military technology has come a long way since the invention of the Gatling gun. Today's military arsenals include an array of weapons capable of incredible precision and destruction, but they all share one aspect in common. The weapons that are operational and deployed today remain under significant human control and fall under the human-in-the-loop or human-on-the-loop control mechanisms. Both of these mechanisms follow a simplified command and control process called the OODA cycle, in which the steps of observe, orient, decide, and act are followed in succession and in as many sequences as required in order to achieve an objective. The OODA loop, which comprises of the observe, orient, decide and act steps, was developed by Air Force Colonel John Boyd and is a concept crucial to understanding how military autonomous technology operates.

The human-in-the-loop mechanism requires that the human operator be in constant interaction with the weapon in order for it to function properly. The scenario of a soldier firing an assault rifle by pulling the trigger satisfies this mechanism because the weapon in question ceases to operate (fire) when the soldier stops pulling the trigger (interaction stops). The human-on-the-loop mechanism removes the human operator a little further from the operating sequence because it only requires a supervisory role on the operator's part. The human operator remains involved throughout the weapon's operating cycle, making regular modifications to the weapon's usage in order to achieve mission objective; some of the subtasks may be fully automated to increase precision but the human operator maintains full control over target selection and engagement. Weapons that operate with a human-on-the-loop mechanism are sometimes described as supervised autonomous weapons, such as point defense systems installed on warships. The single-seat F/A-18 Hornet fighter is a classic example of a human-on-the-loop weapon, having a high degree of automation and smaller error margins compared to earlier fighter models. The F-35 Lightning fifth-generation fighter, developed more than two decades after the F/A-18 Hornet, still operates on a human-on-the-loop basis, albeit a more advanced one that requires even less human input. The aircraft's eight million lines of code underpin its highly automated performance. Within the automation spectrum, the human-on-the-loop mechanism can be described as the transitional phase between the human-in-the-loop and the next level of control, the human-out-of-the-loop mechanism, which features three important attributes. First, the concept of a human operator becomes obsolete because the only human input provided for the weapon is the objective to be achieved; the weapon is free to pursue any course of action to ensure the completion of the objective based on a set of sophisticated decision matrices it has been programmed with. Second, the timescale of the OODA loop in this scenario is reduced to a near zero due to the machine's extraordinary computing capabilities that allow it to process information, arrive at a decision, and execute that decision at an extremely fast rate. The human-out-of-the-loop mechanism describes the manner that AWSs will operate in the hyperwar environment, where fully autonomous warfighting machines will carry out combat operations at the tactical, operational and even strategic levels without human control. The warfighting capabilities and advantages made possible by this mechanism are revolutionary and attractive, but there remains an insidious caveat. Once AWSs are engaged, there is no telling which path they will pursue to accomplish the mission objective given the unpredictability of the battlespace and the absence of a regular human presence to provide guidance and direction. In the absence of meaningful human control, the potential for malfunction or unplanned programming to lead to catastrophic collateral damage is unnerving. Unresolved ramifications for legality and human rights have led a coalition of scientists, legal scholars and business leaders to condemn the development of LAWs and call for an international moratorium. And although latest studies suggest that a case can be made for strong public opposition to AWSs, it is also greatly nuanced and may change depending on the specific context.

How does the public respond to advances in weaponry? As a cornerstone weapons of mass destruction, nuclear weapons represent a critical advancement in global military arsenals. Even though they were only used once in human history, nuclear weaponry has changed the way states think about war; this is reflected most prominently in the military doctrines of nuclear powers and the military alliances deriving from nuclear weapon ownership. Their destructibility, as presented during World War 2, invites strong emotions from the public and have become the basis for the nuclear taboo, an established aversion to the use of nuclear weapons. A thought experiment conducted by political scientists from Stanford University in 2017 reveals that the nuclear taboo theory is not as compelling as it has been presented. The experiment finds that a majority of Americans approve the use of nuclear weapons in the following scenario:

- Iranian forces attack a U.S. warship in the Gulf of Persian
- Congress declared war

- Two military options are presented to the president: the deployment of ground troops into Iran or executing a nuclear strike on an Iranian city to end the war
- The nuclear strike will result in 2 million Iranian civilian deaths but will save 20000 U.S. soldiers from being killed in ground combat

The scenario above suggests that support for the use of the nuclear strike option comes from the 20000 spared American soldiers who would otherwise be killed if a conventional invasion is committed, even if the nuclear strike would decimate the targeted Iranian city's population. Hypothetically, the same level of support should be found in the deployment of AWSs, given that they provide the similar advantage of not requiring American troops on the ground and at the same time, theoretically, does not result in massive civilian casualties on the enemy's side. The limitation to this argument is driven by the fact that the outcomes of using AWSs are difficult to account for operationally (what happens when the weapons achieve the mission objective but an extremely high cost that was not anticipated) and legally (who is responsible for this unplanned cost and its consequences). The benefits for casualty aversion are still offered, but the costs become uncertain and the potential for collateral damage increases dramatically, which can have a bearing on public opinion. How exactly will this factor of uncertainty impact the public's trust in deploying AWSs? The survey study conducted as part of this study helps answer the aforementioned puzzle, but an in-depth look into how the public currently views Unmanned Aerial Vehicles (UAV) also provides useful insights since unmanned technology has been applied in both the civilian and military spheres for a number of years and is the foundational platform upon which autonomous military technology will be developed. Growing spending on drone programs, the ongoing conversion of several Air Force units from fighter-based to drone-based, pledges of support from politicians and the further integration of autonomous technology in U.S. military strategy documents suggests that drones will play an even greater and more visible role in military conflicts in the coming years. As they fall into the category of Remotely Operated Weapons (ROWs), UAVs are considered the gateway technology to fully autonomous weapons. Public perception and understanding of the former can have a great bearing on the latter, as polling data and analysis would reveal.

The U.S. Air Force was the pioneering force behind the development of the original UAVs, built for military use long before UAVs entered the civilian spheres for industrial or recreational purposes. The AQM-34 “Ryan Firebee” was developed in the early 1950s and its technology progressed significantly in the following decade due to demand for stealth surveillance and reconnaissance in Southeast Asia. Ryan Firebees were deployed extensively during the Vietnam War and had a reliable service record, with 83% of vehicles returning to be flown again after deployment. Drone technology underwent significant progress between 1960s and 1990s under the stewardship of the U.S. and Israeli armed forces, the two institutions most interested in the technology at the time, but drones did not enter the American public psyche until after the September 11th terrorist attack, which saw drones being deployed in foreign battlefields in both unarmed and especially armed missions at an unprecedented level. The first UAV strikes were conducted in 2001 with mixed success and by 2006, according to then Chief of Staff of the U.S. Air Force, General Michael Mosely, the military had begun to truly transition to using UAVs for hunting and killing missions, singling out the MQ-9 Reaper model as the technology that represented this shift. Because of their widespread deployment and frequent portrayal in popular culture, Predators and Reapers are reportedly the most recognizable military UAVs and have become the representation of cutting-edge U.S. military technology and power (Horowitz 2016). As UAV technology became more accessible, drones began to enter the civilian domain where their commercial applications are being explored and developed aggressively. Online retailer powerhouse Amazon and logistics giants such as UPS, FedEx and USPS are all considering the gradual integration of UAVs in their delivery services in an attempt to make operations more efficient and reliable. The greater use of UAV technology in both the military and civilian spheres, aggravated by an increase in lethal drone strikes highly coveted by the media, inevitably thrusted drones into the public spotlight.

Chapter 3. Current debates (literature review)

Against the backdrop of claims of efficiency, accuracy, reliability and the promises of critical thinking capabilities in the future, there is considerable but not overwhelming public support for drone usage. However, approval levels vary under different scenarios. The latest data suggests that Americans generally support the use of drone strikes when they are used against loosely-defined terrorist or extremist targets in foreign countries at an approval rate of 50% or more from respondents. Notably, a Pew Research Center poll conducted in 2013 reveals that 55% Americans approve of “the United States conducting missile strikes from pilotless aircraft called drones to target extremists in countries such as Pakistan, Yemen and Somalia”. In the same year, a Fairleigh Dickinson University PublicMind Poll shows a 75% approval rate for “the U.S. military using drones to carry out attacks abroad on people and other targets deemed a threat to the United States”. The same Pew Research Center survey repeated in 2015 returned higher approval rate of 58%. These polls suggest that approval is especially high (above 50%) in scenarios where targets are not U.S. citizens and locations are abroad. Against the backdrop of this overall level of support for drone use is the revelation that public understanding of UAV capabilities has a significant impact on public opinion on using UAV in lethal combat operations. This is an important finding from a recent Center for a New American Security (CNAS) survey carried out in 2016, which found that “the U.S. public is largely unable to differentiate accurately between manned and unmanned armed aircraft” and that this limited understanding contributed to their levels of support for drone strikes. Specifically, respondents in this poll have an exaggerated understanding of UAV capabilities. Given that public understanding of UAV capabilities has a determinant effect on public support for UAV deployment, it can be presumed that public understanding of AWS capabilities will also greatly influence public opinion on whether or not AWSs should be used to perform lethal force. However, whereas an exaggerated perception of drone capabilities likely led to approval for drone strikes, it remains to be seen what effect an underestimated or overestimated understanding of AWS capabilities will have on public attitude toward the use of AWSs. The latest CNAS results also suggest that public support for drone strikes, as opposed to manned airstrikes, is considerably lower when there are more civilians on the ground, suggesting a lower confidence in UAVs when risks of civilian casualties are higher. This finding is especially pertinent to research on public perception of AWSs because lethal autonomous weapons possess an even greater discretion of when, where and how to engage a target, therefore making the possibilities of collateral damage more pronounced. Finally, the CNAS poll also indicates that legality factors strongly into public approval for airstrikes, regardless of whether or not they are conducted using manned or unmanned vehicles. Given the nebulous regulatory landscape for drone strikes and an even more immature legal framework concerning AWSs, both internationally and domestically, this finding on the legality question also provides a basis for approaching the similar challenge posed by the undetermined legal status of AWSs from the perspective of the public. CNAS' findings about the relative lack of confidence in using drones where risks of civilian casualties are high and the impact of understanding of the technology on approval are further corroborated by a report issued by the USPS Office of the Inspector General in 2016 on the "Public Perception of Drone Delivery in the United States". This report, undertaken from the civilian perspective of UAV usage, reveals that the public is more concerned about the malfunction of a drone rather than its intentional misuse and strong knowledge of the technology leads to greater receptivity.

Although research on Autonomous Weapon Systems in the Political Science discipline is not prolific owing to these weapons still remaining in the development phases, there have been a number of foundational studies that looked into public opinion on the concept of fully autonomous warfighting machines. Support for using AWSs, even when asked in a vacuum, lags behind support for using drones. An survey conducted in 2015 by a research team from the University of British Columbia suggests that Lethal Autonomous Weapons (LAWs) are overwhelmingly unpopular on an international scale, with 67% respondents supporting a ban of LAWs. A staggering 81% of respondents also oppose the offensive use of LAWs. The input from this survey remains high-level, however, as it does not provide significant contextual information for respondents. Questions are framed in a vacuum that do not provide insights on whether different contextual conditions, such as level of war and extent of risk to civilians, would have a material impact on respondents' responses.

The limitation from the UBC survey is rectified in a study conducted by Michael C. Horowitz from the University of Pennsylvania in 2014, which employs two survey experiments to find out which conditions affect public attitudes toward AWSs. Findings suggest that public opposition to AWSs is not widespread but contextual, and declines significantly when using AWSs provides higher military utility, the lives of U.S. military personnel are more protected, and other countries begin to develop their own autonomous warfighting capabilities. This is noteworthy because Horowitz's study indicates military utility factors strongly into support for using AWSs, similar to how the public views nuclear deterrence in spite of the scale of destruction and casualty in actual usage. Whether or not the value of military utility is compelling enough for the public to forego concerns about erroneous scenarios with potentially disastrous consequences was not addressed. Horowitz also points out that since true AWSs do not currently exist, what the public is able to perceive about them is largely influenced by popular culture, thus its response to questions asked in the vacuum carry a significant bias. The effects of this bias can be managed by asking questions that are more particular in the context, circumstances and scenarios of using and developing AWSs, which this research's survey procedure seeks to implement. Horowitz's findings also somewhat negate the notion established by previous research that opposition to AWSs, especially LAWs, is overwhelming and untenable. This is significant because it suggests that public opinion, which includes trust as a factor of the expressed opinion, on AWSs is fluid depending on the contextual information.

Other studies have suggested that human rights NGOs around the world, led by the Human Rights Watch, are responsible for the anti-AWS rhetoric that often comes along with advocacy for a comprehensive ban on the development and usage of AWSs. Keith Kirkpatrick writes that the NGOs' concern center around the potential for these warfighting machines to commit unexpected actions due to the limitation of their programming, which may lead to false target identification, among other mistakes that could result in severe damages to physical infrastructure and loss of human life. The removal of the "human in the loop" function, key to true AWS performance, means that armed robots may be difficult to be recalled or provided interventive commands once deployed to the battlefield. On the contrary, and in response to civil society's opposition, military officers and defense officials express a strong confidence in AWSs, although one that is "predicated on advances in these machines' capabilities". Officers and officials such as Lieutenant Colonel Michael Saxon from West Point and U.S. Navy Postgraduate School Professor John Arquilla argue that the case for AWSs as the future of warfare is robust and that risks for accidents and errors are always present in the conduct of war, regardless of the mode of weaponry. As foreign nations develop and deploy their own offensive autonomous weapons, clinging onto the human in the loop mechanisms out of fear of erroneous performance risks eroding the competitiveness of U.S. capabilities in the dawn of hyperwar. Furthermore, in several scenarios where communications may be jammed naturally or artificially while the battle rages on, it is presumed that AWSs must be able to effectively continue to perform and achieve mission objective even after being cut off from the command center, and along with it, presumably any meaningful measure of human control.

In Army of None: Autonomous Weapons and the Future of War, Paul Scharre approaches the puzzle of AWSs from the premise that there is no question AWSs will become a part of our future just as AI applications continue to permeate civilian life. Trust from military commanders is essential to the successful integration and deployment of AWSs, Scharre finds, but this does not appear to be a challenge as he also readily points out in his research that most military personnel have accustomed themselves to a practice of trusting machines without question, as observed in the operators of the Patriot system. This "automation bias" is seen in service personnel, but whether or not this pattern is also represented in the public was not answered in Scharre's latest publication. The burgeoning differential between automation bias within military circles and automation bias within the public-at-large promises to become a focal point in the larger debate about what is meaningful human control. While military personnel will develop and oversee the use of AWS, the public reserves a more prominent and relevant role in accountability, which includes the challenge of blame attribution. On the question of public opinion, Scharre considers the arduous challenge of measuring public attitudes, echoing Horowitz's offering that the standards for making claims about AWSs' inherent violation of the public conscience should be high, while also fielding Peter Asaro's caution that public opinion should not be construed as a one-dimensional representation of public conscience. Scharre seems to be undecided on the best approach to measure and evaluate public opinion on AWSs, which may explain his traditional preference to work with those deeply involved with the subject matter instead. At the same time, however, he does not dismiss the usefulness of public opinion and the role it plays in the future of AWSs, referencing the ban on land mines and cluster munitions as a direct consequence of citizens' pressure.

With scientific luminaries and NGOs coalescing anti-AWS action on one side and defense officials advocating for more advances in autonomous capabilities, it is important to examine the public's level of tolerance and trust in using AWSs, which has been the missing piece thus far in scholarly work on AWSs as perceived by the public. The public lacks the deep technical expertise possessed by scientists familiar with AI capabilities who have come out against the development of AWSs. At the same time, it is not entirely cognizant of the massive potential for AWSs to transform the battlespace beyond recognition. The consensus among researchers thus far is that most of the public's exposure to AWSs has been made through popular culture, such as dramatic portrayals in cinema, of which the Terminator film series particularly stands out as many a layperson's gateway to the world of AWSs. Public attitudes toward AWSs can be fluid in different scenarios and when given greater contextual information. Ironically, ill-informed as it is, the public retains the ability to make the future of AWSs a reality. Systematically benchmarking the public's perception of and trust in AWSs in a world increasingly enthralled in the possibilities of AWS-driven hyperwar is therefore an unneglectable disposition.

Much has been written about blame attribution in criminal justice, industrial psychology and other social science disciplines. Current academic discourse on the function of blame attribution and value of this cognitive behavior in foreign policy and revolution in military affairs exists to a lesser extent, primarily addressing controversial or emerging issues in which a fair amount of public interest and scrutiny in blame assignment can be expected, such as private military contracting and cyberattacks. At the heart of blame attribution in foreign and military policy is the central effect that perceived causality has on a person's ability to make a coherent judgement on blame assignment (Johnson et al. 2019). This cognitive processing of causality is variously described as casual mechanisms, casual linkages or even casual stories in the blame attribution literature on foreign and military topics (Schulzke 2018; Johnson et al. 2019; Sulitzeanu-Kenan & Zohlnhofer 2017). It is commonly accepted within existing literature that harm alone does not provide enough of a drive and information for the public to attribute blame. When casual linkages are present, the connections between the actor and the outcome become much more clarified (Johnson et al. 2019). This facilitates the public's ability to rationally assign blame. Blame attribution dynamics as observed in private military contracting and cyberattacks are relevant insights for examining blame attribution in using drones and AWS because all of these military policies are united in their unconventional and controversial nature. A study on blame attribution in military contracting has produced results correlating the public's blame attribution with particular policy preferences in context-rich scenarios. Similarly, Schulzke's research on the attribution problem of cyberattacks tied to uncertainty shows that the public's preferences for national security or defense policies are significantly informed by blame attribution (Schulzkem 2018). Sulitzeanu-Kenan and Zohlnhofer write that blame attribution patterns are reflective of how policy choices are interpreted (Sulitzeanu-Kenan & Zohlnhofer 2017). Finally, both Schulzke and Johnson's studies find that supplying contextual information has a notable and valuable effect on how the public assigns blame (by facilitating the formulation of casual linkages), but at the same time they maintain that there is not enough evidence to establish that the more context results in greater ease and proper blame assignment.

Chapter 4: Hypothesis

This study seeks to examine the public's understanding of machine autonomy in military weapons by evaluating perception of capabilities, trust and blame attribution. Specifically, the study will focus on unmanned vehicles, also known colloquially as drones, since they represent the warfighting technology that is closest to the future's autonomous weapon systems poised to dramatically change the dynamics of warfare and international relations. Subsequently, the three-dimensional examination will attempt to answer the following questions:

- Is the public able to differentiate between remote capabilities and autonomous capabilities? In other words, does the way respondents understand different operating scenarios suggest that they register a difference between these two types of capabilities?
- How does the public trust using drones in war? What is the extent of this trust and does it translate into support for military policies favoring the use of autonomous capabilities?
- Does the public tend to blame machines more or some other human elements when collateral damage occurs? What are factors that may help establish the causal linkages in erroneous scenarios?

Corresponding to the overarching questions above, I hypothesize that:

H1: The respondents generally fail to distinguish remote-controlled capabilities from autonomous capabilities with an overestimation of the former. To this extent, it is expected that the respondents will interpret the machine as the source of remote-controlled capabilities while these are in fact authorized and operated by humans from a distance.

H2: Trust in drones, attached to whether or not one approves using them in warfare, is contextual; the variance in trust and approval levels is caused by different operating circumstances and the battlefield environment. Overall support for drone usage in war, which is expected to be considerable (over 50%), does not translate into preferences for drones in specific scenarios where the risks of collateral damage are amplified.

H3: Because of an overestimation of remote-controlled capabilities, a reality exacerbated by the absence of a uniformed understanding on the wide spectrum of autonomous technology, the public tends to blame the machine element in both remotely-controlled and autonomous scenarios.

To answer the questions above and test the hypothesis, my research will draw upon existing debates within the current literature as well as a survey study. In the first phase of my research, I will conduct a thorough review of a variety of secondary documents to establish a baseline understanding of current public opinion on both hypothetical AWSs and real-life UAVs, as well as the central question of blame attribution. The goal of this phase is to establish whether public opinions about drones are similar to opinions about AWSs, an indication of capabilities misinterpretation, and patterns of blame attribution in similarly stigmatized defense policies.

In the second phase of the research, I will construct and administer a survey study that examines the University of Houston students' attitude toward using AWSs in a variety of combat scenarios that take into account varying levels of machine-human control, namely the "human in the loop", "human on the loop", and "human out of the loop" scenarios. It will also look at how the research subjects respond to hypothetical scenarios involving the deployment of both human soldiers, including human-operated weapons, and armed UAVs with regard to capabilities, trust and blame attribution, dimensions that are useful to answering the aforementioned questions. This survey study will use the student population on campus as research subjects and will benefit from a diverse pool of subjects, even though the limited and localized study group means that results cannot be used to extrapolate attitudes of the larger public.

Chapter 5: Study Tool Design

This research utilizes a multiple-choice survey to investigate the validity of the stated hypotheses. To gain further insights into public attitudes on AWSs beyond the baseline that has been established in prior research, the survey minimizes the use of questions asked in a vacuum and supplies contextual information for most of the questions. Response choices are limited to mostly "Yes" or "No" and "Approve" or "Disapprove" responses. The dichotomous design of the survey intends to minimize any potential confusion or biases arising from a military subject that is not well understood by the general public, thus delivering a more streamlined user experience for survey respondents while also enabling the study to conduct quantitative analysis of the data collected more effectively. The survey will be comprised of the following sets of questions:

A. Eligibility: These questions will determine whether or not the potential subject can proceed with the remainder of the survey.

1. Perception of capabilities: These questions examine the subjects' perception of drones. They are utilized to determine if the subjects are able to distinguish between remote and autonomous technologies.

2. Approval for using drones: These questions examine how likely the subjects would approve the use of drones in two mission scenarios: high-risk and low-risk with regard to the potential for collateral damage to incur. These questions are used to compare support for using drones in these specific situations with overall support for using drones in military operations.

3. Blame attribution in using drones: These questions calibrate the respondents' concept of blame attribution and examine how they attribute blame for collateral damage in scenarios that progressively place the human element further away from the collateral damage in question.

B. Follow-up: These questions examine the subject's frame of reference with regard to their understanding of AWSs. The purpose of this is to establish whether the subjects' responses in all other questions are influenced by their current view of drones as representing or not representing autonomous technology.

The order of questions is completely randomized by the survey platform Qualtrics in order to minimize any bias respondents would otherwise develop if the questions are grouped in sets. Using subjects only from the University of Houston means that the study has inherent shortcomings that need to be addressed. Politically, the student body of the university is known to be liberal. In addition, the student body's average age does not represent the age profile of the American public. While the case may not be compelling when the responses of this student body's members are extrapolated to represent the attitude of the much larger American public, the study benefits tremendously from the diverse population of a university campus that has been ranked as one of the most diverse in America. This is relevant given that racial diversity is a defining characteristic of any study sample that claims to be nationally representative.

Given that subjects are recruited from the University of Houston, a higher education institution, and that the principal investigator is neither a faculty member nor an individual otherwise capable of affecting students' academic performance, it is not expected that a considerable number of subjects will fail the eligibility questions placed at the beginning of the survey. Subjects under 18 and those belonging to protected populations, such as pregnant women, are excluded. The number of unique and complete individual responses required for the study tool is approximately 381 while the study seeks to enroll and screen no more than 1000 subjects. The 381-subject sample size was determined based on a standard confidence level of 95%, a margin error of 5 and the UH undergraduate population of approximately 36,088 students. The 1000-subject cap is based on the historical response and participation rates of similar surveys administered by the Department of Political Science Survey Experiments Working Group.

Chapter 6: Results

The survey was distributed on November 28, 2018 and participation ended on December 7, 2018, yielding 422 qualified responses and exceeding the minimum of 381 responses established by the study design and approved by the University of Houston Institutional Review Board. A qualified response is defined as a submission of the survey that completes all required questions.

1. Perception of drones

Before discussing results of the survey questions included to assess the public's perception of drones against a backdrop of autonomy, it is important to acknowledge that at the time of this study, true autonomous weapon systems in any forms are nonexistent, even when remotely-operated warfighting technologies such as drones may give the false impression that these technologies operate without the level of human involvement that is true of their operation. In other words, the degree of human operation in and control of today's military drone remains extensive. However, when presented with hypothetical situations, illustrated in Q1 and Q3, where the military drone employed does operate autonomously, the respondents tend to misconstrue the otherwise independent drone as human-operated. Q1 provides a fundamental example of the human-out-of-the-loop experience, when the AWS performs the combat action (firing the missile) across the OODA cycle, having observed the tactical situation, made the necessary orientations (leveraging its coding, for instance), arrived at a decision (acquiring the target) and executing the ultimate attack; the human element is restricted to a supervisory observation. Q3, on the other hand, describes a simple variation of the human-on-the-loop mechanism, in which the drone can independently select a target and execute firing but only when authorized by the human. In both autonomous scenarios, the vast majority of respondents credit the human with responsibility for the combat action even when there is little to no human input. Meanwhile, most respondents accurately understand the real-life human-in-the-loop experience demonstrated through Q2 of current drones being controlled and directed by human operators (from a distant facility removed from the immediate physical battlespace), attributing the combat action to the human instead of the machine.

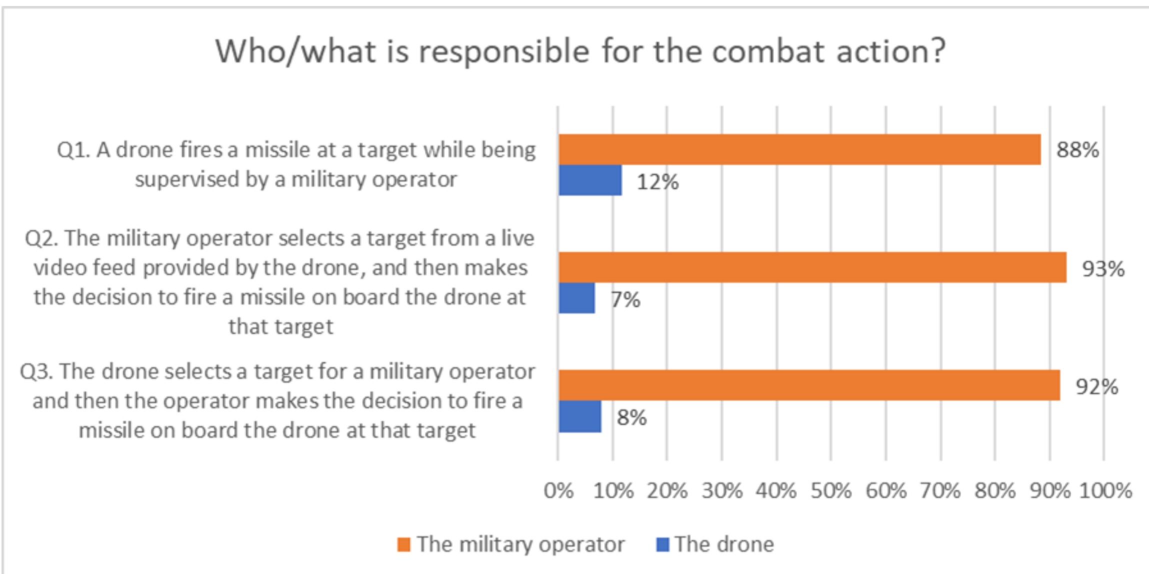

Figure 2. Who/what is responsible for the combat action?

The rather counterintuitive response to these very different autonomy scenarios partly invalidates the first stated hypothesis that the respondents typically overestimate remote-controlled capabilities and consequently characterize them as autonomous. However, the results do show that respondents fail to draw a distinction between remote and autonomous capabilities because respondents tend to ascribe the action to the human element, regardless of the autonomy level attached; whether this has an effect on blame attribution is not clear. The uniform response should not be understood as respondents treating autonomous weapons as human-operated, however, because responses to Q13, the follow-up question that asks whether the U.S. military currently employs autonomous weapons in the form of drones, show that 84% of respondents think current drones are autonomous weapons. Respondents therefore view today's remotely-operated warfighting vehicles as autonomous weapons but at the same time maintain that humans are responsible for actions brought about by autonomy. The latter portion of this finding is especially relevant to the discussion on blame attribution. Given the pull public opinion has on policymaking, it is important to consider if and when the public is sufficiently educated about and acquainted with the nature of autonomous warfighting technology, thus being able to make the distinction among manned, remotely-operated and autonomous platforms and appreciate the radical ramifications of hyperwar, would it still tolerate and support the use of AWSs?

2. Attitudes toward using military drones and trust in drones

Questions Q4-Q8 are included to assess overall attitudes toward warfighting as a means to obtain national interests and the use of drones within that context. Analysis of the responses to this question set reveals several noteworthy observations about the respondents' approval of using drones, a measure of confidence (trust) in this military strategy.

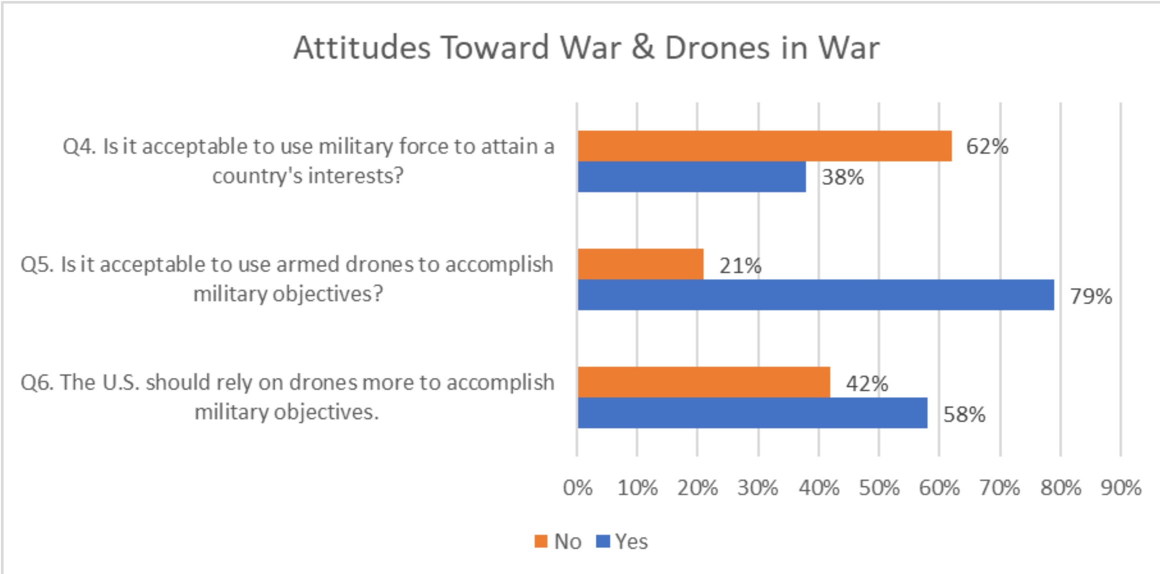

Figure 1. Attitudes Toward War & Drones in War

While a large majority of respondents reject the use of military force in attaining a country's interests, there is overwhelming support for the use of armed drones in accomplishing military objectives, confirming the first clause of the second hypothesis. Similarly, there is considerable approval for the United States expanding the use of drones in its military programs and operations when questions are framed generically without contextual clues. However, there is no evidence to exclude the possibility that the response supporting the use of drones is provided with the understanding that the usage is necessitated militarily, or in other words, when a country must resort to military force, then there will be support for using drones. The high level of support for drones in this research is similar to previous public opinion studies on using drone strikes, albeit this study yields a much higher level of support, likely because it highlights military necessity as opposed to military utility. The data suggests that the respondents are relatively comfortable with deploying military drones and increasing this deployment even when it does not support warfighting as a means to attain national interests.

These observations are notable because they are not consistent with the current scholarly establishment that public opinion on specific national security policies are heavily influenced by the fundamental views on war and military force (Eichenberg 2016). Significant support for using drones in military operations, as shown by 79% positive response to Q5, coupled with a high approval for expanding that use of drones at a 58% positive response to Q6, imply there may be other factors impacting public opinion on a particular national security strategy, which is the increased deployment of military drones, and that the relationship between general views on war and military force and views on a specific warfighting policy may be more nuanced and sophisticated than is commonly believed. Is there any correlation between considerable approval for using military drones and expanding that usage and the common perception that drones are autonomous weapons? Crosstabulation between Q5 and Q13 (follow-up question) reveals that responses to both questions are statistically significant to an impressive degree, with a corresponding p-value of 0.00. Similarly, the relationship between responses to Q6 and Q13 are also statistically significant, underpinned by a p-value of 0.02.

		The U.S. should rely on drones more to accomplish military objectives.
The United States military currently employs autonomous weapons in the form of drones.	Chi Square	5.35
	Degrees of Freedom	1
	p-value	0.02

Figure 3. Chi-Square test for Q5 & Q13.

		Is it acceptable to use armed drones to accomplish military objectives?
The United States military currently employs autonomous weapons in the form of drones.	Chi Square	10.81
	Degrees of Freedom	1
	p-value	0.00

Figure 4. Chi-Square test for Q6 & Q13.

The analysis above lends credence to the view that considerable approval for using military drones and expanding this usage is likely to be influenced by the misconception of drones as AWSs, in spite of low support for employing military force as a means to attain national interests. To this extent, the notion of employing drones as autonomous weapons must include some advantages in order for the respondents to view the option quite favorably, to the degree of supporting expansion of drone usage, even when using military force is frowned upon. The number of U.S. soldiers killed in military operations has been cited as a critical information drawn upon by the public in formulating opinions about war and using military force (Gelpi 2010). Applying this to the case for supporting drone usage, respondents may deduce that deploying drones as autonomous weapons instead of human soldiers would significantly reduce U.S. casualties in military operations. It follows then that the former option would appear preferable to them. However, in the absence of contextual information that is reflective of the real battlespace, general approval for military drones as autonomous weapons can be misleading, especially since it is widely known that civilian casualties also have a substantial bearing on the public opinion of war. Q7 and Q8 discuss differing scenarios whereby the respondents have to confront realities of engaging enemy combatants in an isolated area or a densely-populated location.

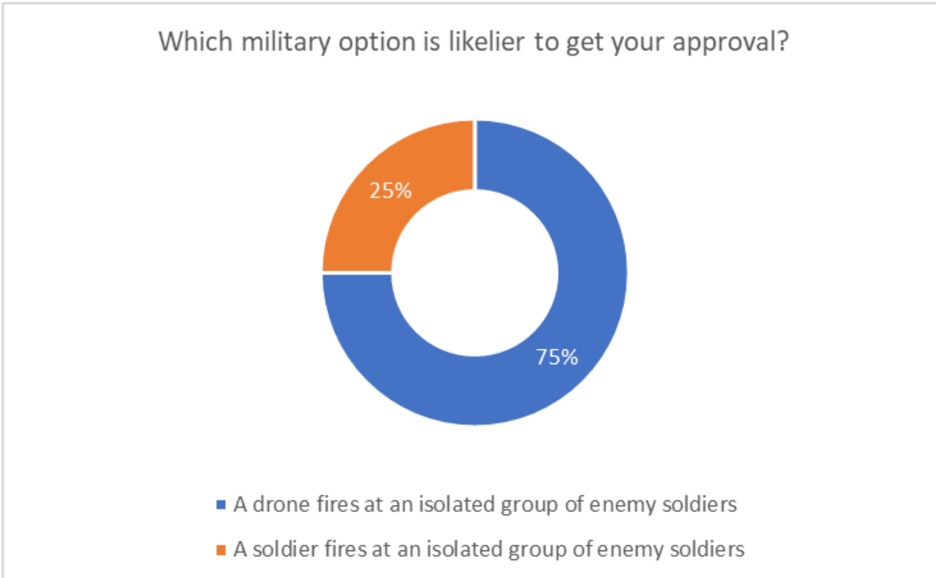

Figure 5. Isolated group of enemy soldiers scenario

In Q7, respondents are asked to select the military option that is most likely to gain their approval. Q7 hypothesizes a combat situation where enemy combatants are in a remote location not inhabited by civilians, thus implying that the risk of civilians being caught in the crossfire and the resultant collateral damage is minimal. A majority of respondents approve the use of a drone as opposed to deploying a human soldier, while only a quarter of responses prefer the human option. However, approval for the drone drops considerably in the high-risk scenario demonstrated in Q8.

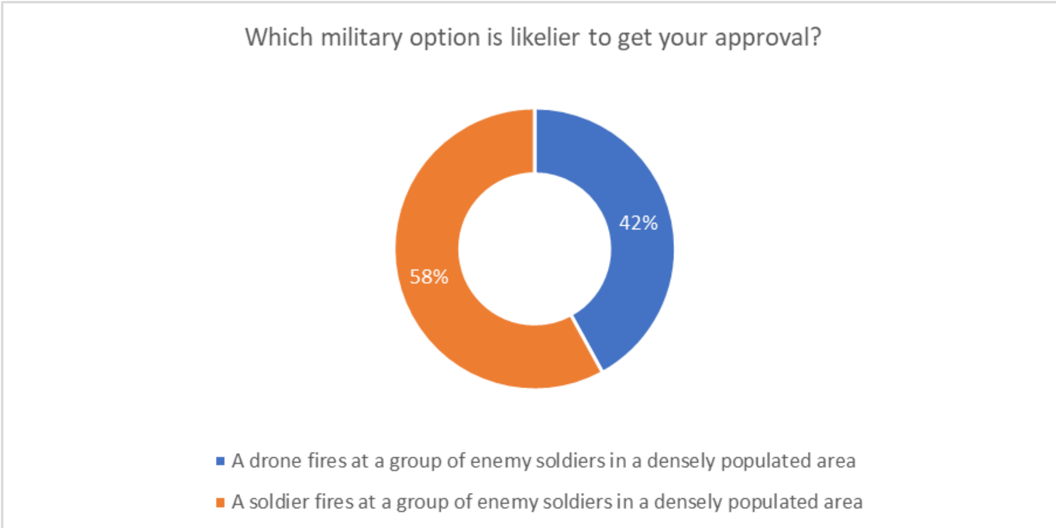

Figure 6. Densely populated area scenario

As was hypothesized, the data shows that there is a clear aversion to using drones when the risk of civilian casualties and collateral damage is pronounced. Whereas respondents may have granted the drone option in the low-risk scenario high approval because of a presumed efficiency in military utility (task performance and overall operational convenience) and the perception of potentially avoiding the soldier casualty, these factors seem to have much less pull in the high-risk scenario. As such, the respondents sufficiently trust the drone to accomplish the task concerned (engaging the enemy soldiers), but it becomes more hesitant to extend to the drone the same confidence when civilians and other related factors are present. The responses in this study show that military utility carries less weight than collateral damage in the respondents' willingness to approve drone usage, contrasting the dominance that military utility seems to enjoy in Horowitz's study. Even without military training, respondents are probably able to perceive that the low-risk scenario in Q7 provides an ideal, "straightforward" operating environment, while the densely populated area is problematic and laden with risks and factors that make the deployment of the drone more unnerving. To extrapolate, the respondents trust the drone enough to perform the provided task, but they do not trust the drone enough to perform that same task without getting itself entangled in some complications. It appears that respondents are willing to forego military utility in the face of collateral damage likelihood. The lower confidence in the drone in the high-risk scenario suggests that accomplishing the military objective at hand may not be as enticing if it entails additional, unforeseeable costs and risks.

3. Blame attribution

Analysis of how respondents attribute blame in combat scenarios where an error is committed resulting in collateral damage draws from both the first and last sets of questions (Q1-Q3; Q9-Q12). Results from the first set of question clearly establish, as discussed above, that respondents overwhelmingly (and unexpectedly) attribute responsibility for the combat action performed to human soldiers or operators, without specifying if the course of action in question is erroneous or not. The second set builds upon the first three questions by framing the act of firing a weapon as erroneous and tying it to a specific collateral damage – the death of a human civilian. In each question from Q9 to Q12, respondents are presented with scenarios of fatal misfiring caused by decreasing levels of “hands-on” human involvement: the soldier, the soldier with a weapon, the soldier in a fighter jet and the soldier operating a drone. Although the first three scenarios appear relatively straightforward, with the human element positioned squarely to receive blame, they are included to illustrate how the evolution of warfighting technology (with a focus on automation eventually leading to autonomy) can make blame assignment a truly challenging endeavor because of complications in formulating clear-cut causal linkages. Responses to these questions yield results similar to the first set of question, with a majority of respondents attributing blame to the human element. This negates the third hypothesis that state respondents would overwhelmingly ascribe blame to the machine. However, a closer look at the data shows that respondents provide the same response to scenarios involving two warfighting platforms, a fighter jet and a drone, that carry with them the impression of being highly automated. Specifically, 68% respondents agree that the human element is to blame for both Q10 and Q12. This is lower than the number of respondents attributing blame to the human element in Q9 and Q11, which are 77% and 84%, respectively. This pattern suggests that one factor may have played a role in forming the causal linkages in these scenarios and influenced the respondents' blame attribution: the physical distance between the collateral damage and the military actor. In Q10 and Q12, the collateral damage is significantly separated from the military platforms in the battlefield by thousands of feet or more of airspace, while in Q9 and Q11, the collateral damage and the military actor are situated much closer to each other. It appears that the more removed the military platform or actor is from the collateral damage, the less blame is assessed for the erroneous course of action. Physical distance in these scenarios, therefore, emerges as a factor contributing to the causal linkages that help respondents make the connection between the actor and the outcome in assigning blame.

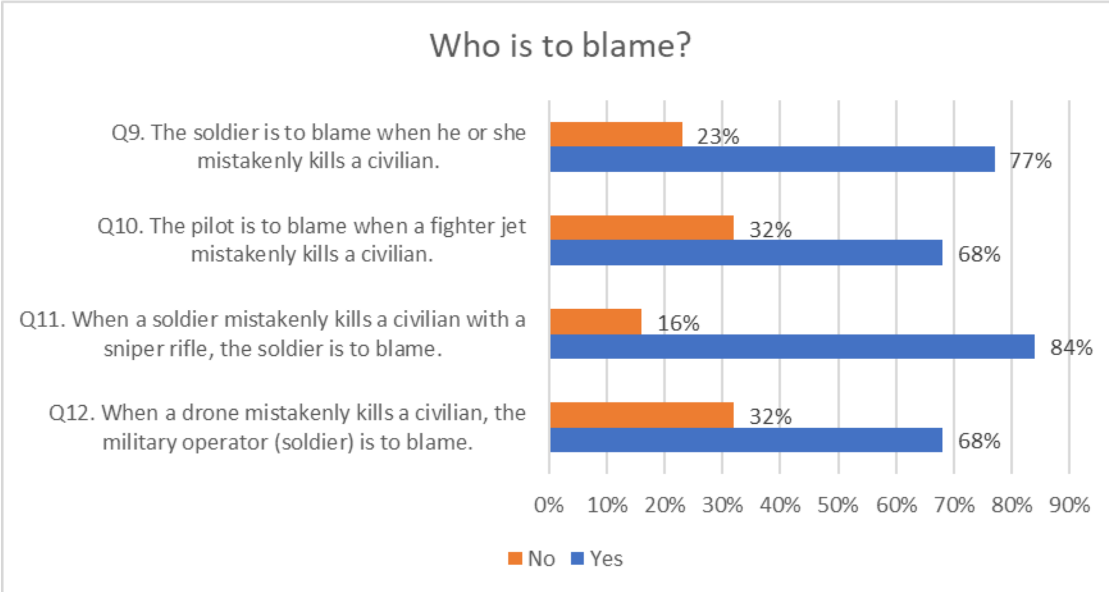

Figure 7. Who is to blame when a civilian is mistakenly killed?

To find out if the same response to both Q10 and Q12 are related, a Chi-Square test is performed and yields a p-value of 0.00, indicating the results are statistically significant. This means that respondents likely applied the same reasoning to both scenarios when attributing blame to the human element. The way respondents attribute blame in Q12 is consistent with how they determine the source of responsibility for the act of firing a missile in the first set of questions. Although this seems rational because it establishes that if a human is responsible for the combat action committed by a drone then he or she will also be ultimately responsible if that course of action leads to a mistake, it nonetheless suggests that the difference between human-operated, remotely controlled capabilities and autonomous capabilities does not seem to register with respondents. This is an unexpected finding that highlights the alarming informational and cognitive gaps between those developing and operating drones and those evaluating the outcomes of using those platforms. Unadjusted expectations and beliefs about warfighting as autonomous capabilities emerge may problematize blame attribution and complicate policymaking. The scenario described in Q12 is worded intentionally to imply that the drone is operating on a highly autonomous basis, similar to the scenarios described in Q1 and Q3. When examining responses to these questions against the responses to the follow-up question (Q13), Chi-Square tests all yield p-values smaller than the significance level of 0.05.

		A drone fires a missile at a target while being supervised by a military operator. Who is respons...	The drone selects a target for a military operator and then the operator makes the decision to fi...	When a drone mistakenly kills a civilian, the military operator (soldier) is to blame.
The United States military currently employs autonomous weapons in the form of drones.	Chi Square	0.85	1.62	0.01
	Degrees of Freedom	1	1	1
	p-value	0.36	0.20	0.94

Figure 8. Chi-Square tests for Q1, Q3, Q12 and Q13.

These results suggest that the blame attribution pattern observed in Q1, Q3 and Q12 is independent from the misguided understanding that drones are indeed autonomous weapons currently being employed by the United States, which 84% of respondents accept in Q13. Thus, notwithstanding the irrationality of responses with respect to an apparent oversight of the implications autonomous technology bears, it can be said that respondents are likelier to attribute blame to some human element rather than the machine even when they view the military technology as autonomous. Although the survey does not elaborate on the full intricacies of attributing blame in fully autonomous scenarios in order to find out how the respondents would specify the human responsibility (as in, is it the military personnel overseeing the operation that involves the AWS or the manufacturer of the technology?), the findings of blame attribution in this study can still arrive at a generalization that respondents see the human element, whoever it is, more prominently than it sees the machine counterpart when responsibility and accountability are concerned. In an attempt to explain this dynamic that was also observed elsewhere, Daniel Kahneman suggested that the desire for mistakes to be attributable to human error may have more to do with moral preferences than how well machines can perform tasks (Kahneman 2015). The blame attribution pattern observed in this study seems to confirm that hypothesis since the demonstrated lack of understanding about autonomous technology does not seem to inform decisions to attribute blame to the human element.

Results from the first and last sets of questions, both of which address the dynamics of responsibility and blame attribution, give rise to a number of ramifications. If the public has a tendency to misinterpret the human-versus-machine control of a future AWS platform that leans toward overestimating the role (and thus responsibility) of the human element while discounting the autonomous capabilities, then there may be challenges in managing accountability and expectations for hyperwarfare operations in the future that supposedly will involve extensive use of autonomous technology. In a hypothetical scenario where an AWS owned by a foreign country malfunctions and erroneously attacks a U.S. diplomatic facility or personnel, the American public would likely adopt the conclusion that the action was sanctioned by a human or a group of humans and thus demand retaliation, instead of dismissing the incident as a computer-generated or some other situational error. This scenario illustrates the issue of fundamental attribution error that has been described as a major impediment to conflict resolution in international relations (Rapport 2017). Such a scenario can easily lead to war when occurring in countries that are experiencing contentious relations with the U.S. For example, it is difficult to imagine that the American public will be amenable to the Russian explanation that one or multiple Russian AWS destroyed an American passenger aircraft or vessel because of unexpected machine bias or deficiencies in the machine's programming. Therefore, the effects that AWS behaviors caused by machine bias or technical malfunction may have on a nation's political calculus and military decision-making must be scrutinized as autonomous technology evolves. They also reiterate the importance of the discussions about the concept of "meaningful human involvement" within the progression of autonomous technology. Apart from this, the results are also in agreement with the contemporary conundrum of how to reconcile future military autonomous technology with human rights principles, legal standards and other formal and informal norms that govern international relations. Even when the misconception of today's drones as autonomous weapons and the propensity to ascribe autonomous actions to the human element are paradoxical, they are reflective of the existing deficiencies in the science, ethics and psychology of manufacturing and operating AWS, which inhibit the public from formulating the causal linkages in assigning blame. These are limitations that would make the future use of AWSs a problematic and even disastrous undertaking if they are left unaddressed while the technology continues to progress.

Conclusion

The results of the survey study essentially confirm the three stated hypotheses. However, data analysis reveals insights that go beyond the individual hypotheses themselves, most of which have to do with the relative independence between the respondents' understanding of the varied autonomous capabilities and the manner respondents ascribe responsibility and attribute blame when these capabilities are used. While a staggering majority of the respondents are not able to differentiate between autonomous and remotely-operated capabilities, they also identify the human element as being responsible for combat actions executed by both types of capabilities and for the resultant collateral damage. This correlation is supported by subsequent Chi-Square tests, further reinforcing the position within existing debates that the tendency to attribute blame to the human element is driven more by moral dictates and neither an understanding of autonomous capabilities nor confidence in their performance. The drive to blame the human element for collateral damage is reduced in scenarios where the warfighting technology enables the human element to be more removed from the physical battlefield, but this decrease is not significant enough to be conclusive. The study also finds that a high level of overall support for using drones and increasing that use exists when no other contextual information is provided, and that this level of support is related to the perception of drones as autonomous weapons. Contextualizing the use of drones by putting them in specific operating environments reveals that approval of drone usage varies and depends on the likelihood of collateral damage incurring.

The dynamics of blame attribution in the use of drones are especially pertinent to international relations because they indicate responses to erroneous incidents committed by autonomous platforms may completely disregard the technology. This promises to complicate interstate relations because future operators of AWSs may not be able to satisfy external expectations about accountability and responsibility when something goes awry. Given that it is reasonable to assume military autonomous technology will not progress at the same pace in every country, the misalignment between capabilities and expectations is slated to become a source of international controversy and even conflict.

Bibliography

Allen, John, and Amir Husain. 2017. "On Hyperwar." U.S. Naval Institute Proceedings Magazine: 147.

Barnes, Michael Joseph, and Florian Jentsch. 2010. Human-Robot interactions in future military operations. Farnham, England: CRC Press.

Beard, Jack M. 2014. "Autonomous Weapons and Human Responsibilities." Georgetown Journal of International Law 45: 617-681.

Brock, John W. 2017. "Why the United States Must Adopt Lethal Autonomous Weapons." U.S. Army Command and General Staff College.

Bunker, Robert J. 2017. "Armed Robotic Systems Emergence: Weapons Systems Life Cycles Analysis and New Strategic Realities." Strategic Studies Institute.

Carpenter, Charli. "Beware the Killer Robots." Foreign Affairs, Foreign Affairs Magazine, 9 July 2013, www.foreignaffairs.com/articles/united-states/2013-07-03/beware-killer-robots.

Carpenter, Charli. "How Scared Are People of 'Killer Robots' and Why Does It Matter?" Open Democracy, Open Democracy, 4 July 2013, www.opendemocracy.net/charli-carpenter/how-scared-are-people-of-"killer-robots"-and-why-does-it-matter.

Carpenter, Charli. "The SkyNet Factor: Four Myths about Science Fiction and the Killer Robot Debate." The Washington Post, WP Company, 3 Sept. 2014, www.washingtonpost.com/news/monkey-cage/wp/2014/09/03/the-skynet-factor-four-myths-about-science-fiction-and-the-killer-robot-debate/?utm_term=.fe4d2e4b2943.

Carpenter, Julie. 2016. Culture and human-Robot interaction in militarized spaces: a war story. London: Routledge, Taylor & Francis Group. (available for rent on Kindle)

Caton, Jeffrey L. 2015. Autonomous weapon systems: a brief survey of developmental, operational, legal, and ethical issues. Carlisle, PA: Strategic Studies Institute and U.S. Army War College Press. (available on Kindle)

Crook, John R. 2013. "U.S. Department of Defense Directive on Autonomous Weapons Systems" The American Journal of International Law 107(3): 681-684.

Defense Science Board. 2016. Summer Study on Autonomy Summer Study on Autonomy. Washington, D.C.: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics.

Docherty, Bonnie. "Losing Humanity | The Case against Killer Robots." Human Rights Watch, 29 Apr. 2015, www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots#.

Eichenberg, Richard C. "Public Opinion on Foreign Policy Issues." Oxford Research Encyclopedia of Politics, 2016.

Gelpi, Christopher. “Performing on Cue? The Formation of Public Opinion Toward War.” Journal of Conflict Resolution, vol. 54, no. 1, 2010, pp. 88–116.

H. M. Roff & D. Danks (in press). “Trust but Verify”: The Difficulty of Trusting Autonomous Weapons Systems. Journal of Military Ethics.

Hall, Brian K. 2017. “Autonomous Weapons Safety.” Joint Force Quarterly 86(3): 86–93.

Horowitz, Michael C. 2016. “Public opinion and the politics of the killer robots debate.” Research and Politics 3(01): 1-8.

Johnson, Austin P., et al. “Can Hierarchy Dodge Bullets? Examining Blame Attribution in Military Contracting.” Journal of Conflict Resolution, 2019, pp. 1–21.

Kahneman, Daniel. Thinking, Fast and Slow. Farrar, Straus and Giroux, 2015.

Kirkpatrick, Keith. 2016. “Can we trust autonomous weapons?” Communications of the ACM 59(12): 27–29.

Press, Daryl G., et al. “Atomic Aversion: Experimental Evidence on Taboos, Traditions, and the Non-Use of Nuclear Weapons.” American Political Science Review, vol. 107, no. 01, 2013, pp. 188–206.

Rapport, Aaron. “Cognitive Approaches to Foreign Policy Analysis.” Oxford Research Encyclopedia of Politics, 2017.

Scharre, Paul. 2016. “Centaur Warfighting: The False Choice of Human vs. Automation.” Temple International and Comparative Law Journal 30(01): 151-165.

Scharre, Paul. 2018. Army of None: Autonomous Weapons and the Future of War. New York: W. W. Norton & Company.

Schneider, Jacquelyn, and Julia Macdonald. “Why Troops Don't Trust Drones.” Foreign Affairs, Foreign Affairs Magazine, 5 Jan. 2018, www.foreignaffairs.com/articles/united-states/2017-12-20/why-troops-dont-trust-drones.

Schulzke, Marcus. “The Politics of Attributing Blame for Cyberattacks and the Costs of Uncertainty.” Perspectives on Politics, vol. 16, no. 4, 2018, pp. 954–968.

Sparrow, Robert. 2016. “Robots and Respect: Assessing the Case Against Autonomous Weapon Systems.” Ethics & International Affairs 30(01): 93–116.

Sulitzeanu-Kenan, Raanan, and Reimut Zohlnhöfer. “Policy and Blame Attribution: Citizens’ Preferences, Policy Reputations, and Policy Surprises.” Political Behavior, 2017.

Munster, Gene. “Here's When Having a Self-Driving Car Will Be a Normal Thing.” Fortune, Fortune, 13 Sept. 2017, fortune.com/2017/09/13/gm-cruise-self-driving-driverless-autonomous-cars/