

SHLOMO BRODY

# Jewish Ethics in the Age of Autonomous Warfare



MAGINE YOU ARE a head of state with the opportunity to kill in one strike the entire political and military leadership of an enemy terrorist group. Every. Single. One. All the leaders and commanders who have launched repeated attacks on buses, cafes, and shopping centers would be gone in a flash. Along with terrorists, however, many non-combatants would inevitably be killed by the massive bomb that would be necessary to topple the building.

In September 2003, Israeli Prime Minister Ariel Sharon faced this dilemma. Hamas leader Ahmed Yassin had gathered with all of his senior men in a three-story Gaza apartment building. This was Yassin's dream team. Intelligence officials, led by Shin Bet head Avi Dichter, saw a historic opportunity to cause irreparable damage to the terrorist group.

Yet Israel didn't strike. Fearful of dozens of civilian casualties and

the local and international protests that would ensue, Sharon, at the urging of IDF Chief of Staff Moshe Yaalon, called off the bomb. An alternative plan was hastily proposed and approved: to fire a smaller missile that would destroy the third floor, where intelligence officials believed the meeting was taking place. They were wrong. The meeting was on the first floor. Immediately after impact, the Hamas men fled. Israel could have utilized drones to blast every screeching car. The defense minister, Shaul Mofaz, ruled out that option. Civilians were likely to be hurt, he said later.

It wasn't just the "CNN effect" that guided Yaalon. Yaalon was weighed down that day by a previous assassination of a Hamas leader, Salah Shahada, in which over a dozen non-combatants were also killed. In an interview with the *Washington Post*, Yaalon asserted that two moral factors guided his thinking. First, any action taken had to pass the 'mirror test': At the end of the day, will he be able to look at himself in the mirror? Second, he learned from his mother, the sole survivor of the Holocaust from her family, that "Jews shouldn't be killed, but it also means that we don't kill others. You need strength to defend Israel, and on the other hand, to be a human." Dichter, by contrast, thought that given the targets, the strike was proportionate and ethically justified. The collateral damage would be extensive but not excessive. Dichter, whose father was the lone Holocaust survivor in *his* family, countered with a different moral lesson from the Holocaust: "I'm not going to let anyone kill a Jew just because he's a Jew."

Who was right: Yaalon or Dichter? The bombing would have wiped out the enemy leadership, but the collateral damage would have been extensive. Would it have been excessive, given the targets? Perhaps not. On the other hand, would new Hamas leaders—or some other terrorist group—have popped up to replace them, anyway? My guess is that SAPIR readers will be conflicted on this question because sound arguments can be made for either side.

Now, suppose, in these days of advancing AI applications, the strike *and* the decision could be made by an autonomous drone system. The system would evaluate the probability of a successful strike, estimate the extent of collateral damage and public outrage, and decide whether to shoot. There would be no last-minute decision-making scramble by security and political leaders, and no emotional baggage from the Holocaust in the background. Would our decision-making be any worse for it? Might it even be better?

From the perspective of Jewish ethics, the broad utilization of autonomous weapons systems would be a terrible moral mistake. Even if we could develop such systems in a way as to result, routinely, in morally reasonable outcomes as reliable as those made by humans, we'd lose a critical component of military ethics. It's not just a problem of legal responsibility, i.e., a problem of who is responsible for decisions made by the autonomous system. That can be solved, as I discuss below. The irreducible problem is that the machine's decisions would lack an ethical reckoning—a moral accounting—critical to the moral life.

To understand why this is critical, it's important to appreciate that Jewish ethical discourse is driven by a plurality of voices and values. As I show in *Ethics of Our Fighters*, my forthcoming book on Jewish military ethics, several types of moral appeals are found in the Biblical canon, Talmudic discourse, and later Jewish legal and ethical writings. These include the following factors:

1) *Dignity of mankind*. All humans, friend and foe alike, are created in the image of God. “Whoever sheds the blood of man, by man shall his blood be shed; for in His image did God make man” (Gen. 9:6). This requires us to grant basic dignity to any person and not cavalierly treat people as a means toward some desired end.

2) *Inherent wrong of illicit bloodshed*. The commandment “Thou shall not murder” is reflective of this deep theological principle and

demands that we do not take a life lightly. In fact, the ability to avoid unnecessary bloodshed is one of the factors that make the Jews worthy of settling the Land of Israel, according to Deuteronomy 19:10.

3) *Individual responsibility.* Individuals bear primary responsibility for their actions and should ideally bear the sole weight of that responsibility. “The person who sins — he alone shall die” (Ezekiel 18:20).

4) *Vision of world peace.* The ultimate Biblical vision is for the cessation of all warfare, and represents a goal toward which humanity must aspire. “And they shall beat their swords into plowshares and their spears into pruning hooks. Nation shall not take up sword against nation; they shall never again know war” (Isaiah 2:4).

5) *Warfare in pursuit of justice.* Until such time, the Bible calls upon its followers to take up arms for the sake of justice. This can be to defend oneself, to settle the homeland, or to rid the world of evil.

6) *Warfare, by its nature, is a collective affair.* This entails citizens and soldiers endangering themselves for their nation, alongside a willingness to kill members of the enemy nation. Accordingly, warfare creates a form of communal identity and responsibility. “When the Lord your God delivers them to you and you defeat them, you must doom them to destruction.... For you are a people consecrated to the Lord your God: of all the peoples on earth, the Lord your God chose you to be His treasured people” (Deuteronomy 7:2,6).

7) *National partiality.* The primary responsibility of political leaders and citizens is to protect their own people. This is part of a general ethos that people have particularistic obligations to their family, comrades, community, or nation. These “associative commitments” create a moral obligation not to shirk one’s responsibility to fight on behalf of the collective.

8) *Bravery and courage.* In warfare, bravery is a virtue and fearful-

ness is a vice. It is virtuous to worry about killing someone illicitly, but nonetheless, one must still fight courageously.

9) *National honor*. As with all actions, the honor of both God and His people is a factor. This requires not acting in an unethical manner that will disgrace our reputation, and not becoming a downtrodden people subjected to mass ridicule.

It pays to take a second look at this list. These values are readily comprehensible and will undoubtedly appeal to many people in various contexts. Several of them clearly played a role in the debate between Yaalon and Dichter, including the dignity of mankind, individual responsibility, national partiality, and national honor. Do you think that some should always take precedence over others? Or might you argue that it depends on the variables of any given circumstance? If the latter — as I think most people would claim — then the challenge for ethicists and leaders is to determine which moral appeals take precedence in any given case.

The methodology for sorting this out is sometimes called “casuistry,” a case-based process for applying ethical principles to resolve moral dilemmas. Here, we are dealing with what my late father, the philosopher Baruch Brody, called “pluralistic casuistry,” i.e., the process of determining which of multiple values should be most prominent in any given circumstance. Some ethicists use multi-value frameworks to balance different values in determining which moral claim should outweigh others in a particular circumstance. Other ethicists express doubt as to whether we can create a hierarchy among competing values; after all, values are difficult to quantify. Instead, they suggest, we should deliberate intensively — and then make a judgment call as to which value or values should take priority. Either way, pluralistic casuistry leads one to take *all these moral claims into consideration* when making an ethical assessment in a given case, as opposed to prioritizing a single factor such as national victory (favored by

ultra-nationalists), or a meta-value, such as human rights, favored by international law jurists.

Pluralists, as the philosopher John Kekes has argued, believe that there is *no absolute hierarchy* of principles that is operative in all situations. All these important values are conditional. No matter how precious a given value might be, it may be violated when it conflicts with another value with a stronger claim in a particular situation. A moral judgment call must be made based on a debate about the relative strength of all competing values — strengths that will vary based on the political, military, and social context of the situation in question. It follows that no algorithm can be relied upon to determine the right answer in all situations. Here lies the primary problem with autonomous killing machines: the inability to create and defend a moral argument for the decisions it makes.

As I suggested, the legal difficulty — who can you hold liable for an action no one planned or performed — is challenging but surmountable: for instance, we might agree, as a matter of convention, that the last human decision-maker bears responsibility. Yet, this legal dilemma reflects another moral problem. In the absence of human control, it may not be possible to explain, after as well as before a decision is reached, exactly what happened and why. To give a moral account of decision making is the ultimate act of ethical discourse. It forces the actor to justify, before and after the act, why he or she prioritized certain values over others in any particular circumstance. It further forces them to learn from those experiences and apply them to future occasions. This form of continuous accounting, which includes but goes beyond Yaalon’s “mirror test,” is a critical part of the moral life.

Autonomous weapons systems are not so much immoral as amoral. That is to say, they don’t allow for the type of moral deliberation and reflection necessary to pass ethical judgment. The actions taken

by such systems may, overall, be as defensible as decisions made by humans, whose judgments can be deeply flawed. Yet, by replacing human deliberation with a machine, we stop using the moral compass that distinguishes our humanity. “To know good and evil,” as Genesis 3:22 puts it, is to be human. Machine decision-making threatens us with the ultimate form of digital dehumanization.

In this respect, it is useful to compare autonomous weapons systems with autonomous driving vehicles. The latter technology remains far from perfect, as recent news reports have highlighted. Human driving, however, is also flawed, both on a technical and moral level. Nevertheless, many criticize automated cars for replacing human judgment when unexpected roadblocks emerge, and accidents are imminent. Some algorithm, the critics suggest, will make a moral decision about who lives or dies—something algorithms should not be doing. And yet, in these sudden, panicked circumstances, little human moral deliberation takes place as drivers make split-second, knee-jerk decisions. Algorithms built into automated cars might actually increase the degree of moral deliberation taking place in these frenzied moments. Accordingly, an autonomous driving model may be morally appropriate. Even if this is so, the same cannot be said for deliberations over whether to kill terrorists meeting in a crowded residential building.

Artificial intelligence can play a critical role in assisting our moral deliberations in such a situation. It can help us identify the right targets, clarify the number of non-combatants in an area, and estimate the level of collateral damage. AI-controlled drones can be utilized for early, high-risk surveillance, and play a major role toward disabling enemy air defenses. These are cheaper ways to knock out missile targets that, critically, don’t run the mortal risks of piloted planes. In these ways and more, technology can help us fight more efficiently, safely, and even ethically.

But moral decision-making with life-and-death consequences must ultimately remain in human hands. Otherwise, there is no moral accountability. And retaining moral accountability is essential for retaining our humanity. \*