How Automated Profiling Can Influence Our Perception of Others

Minao Kukita

Islamic Trust Studies International Symposium "Overcoming the Divide: Connectivity and Trust Building for Middle East Peace"
January 22-24, 2025

1 Introduction

According to a report in *+972 Magazine*, the Israeli army has been using an artificial intelligence profiling system to identify Hamas operatives in the Gaza Strip since the October 7, 2023 Hamas attack. Testimonies from anonymous sources reveal how the Israeli military is abusing this system to practice acts that violate international humanitarian law. This paper will discuss this example and how automated profiling can negatively affect our perception of others and undermine trust.

2 The Israel's Abuse of AI Profiling against Gaza People

Currently, artificial intelligence is being applied to the military in various ways. In so-called "hybrid warfare," AI is used to generate and disseminate fake news, intervening in people's perceptions. In "Lethal Autonomous Weapon Systems" (LAWS), which are believed to be under development by countries such as Russia, the United States, Israel, and China, AI plays a key role in autonomously carrying out actions from identifying targets to executing attacks. However, the focus here is not on such systems, but rather on how AI that performs profiling is applied in military contexts.

Since the September 11 attacks in 2001, the United States has waged a "war on terror," killing individuals suspected of being terrorists in South Asia, the Middle East, and Africa. In doing so, the U.S. military and CIA adopted a method called "signature strikes" during the process of identifying and attacking targets[2, 7]. This method involves monitoring people in regions where terrorists are believed to be hiding and identifying individuals highly suspected of being terrorists based on patterns in their communication records, behaviors, and social connections. Missiles are then launched at these targets from remotely operated unmanned aerial vehicles. The decision to target these individuals was approved by then-President Obama.

Since anti-tank missiles were dropped from high altitudes on targets hiding in urban areas, many civilians, including numerous children, were inevitably caught in the crossfire. According to U.S. government documents revealed through whistleblowing, all individuals killed in drone strikes were categorized under the label "Enemy Killed In Action (EKIA)"[7]. For the United States, individuals exhibiting behavioral patterns suspected of being terrorist-related were deemed terrorists, and anyone killed by U.S. attacks was considered an enemy who deserved to die.

By using artificial intelligence, the same actions as signature strikes can be carried out more "efficiently." According to Israel's online news outlet *+972 Magazine*, this has been implemented by the Israeli military against residents of the Gaza Strip since the October 7, 2023, surprise attack by Hamas[1]. The Israeli military uses an artificial intelligence system called "Lavender" to identify Hamas or Palestinian Islamic Jihad (PIJ) operatives in the Gaza Strip. Israel has constructed an evaluation model by learning from data on known Hamas and PIJ operatives.

The trained model is applied to data on the 2.3 million residents of the Gaza Strip collected through a large-scale surveillance system, evaluating each individual's likelihood of being a Hamas or PIJ combatant on a scale from 1 to 100. Residents who exhibit characteristics such as being part of WhatsApp groups with known combatants, frequently changing their phones every few months, or frequently changing their addresses are assigned higher scores. If their score exceeds a certain threshold, they are placed on a "kill list" and become targets for assassination. This system is designed to address the bottleneck posed by human cognitive limitations in the routine of identifying operatives, pinpointing their locations, and carrying out attacks.

The Israeli military explains that decisions are not entirely entrusted to machines, as humans manually verify the process. However, according to sources interviewed by *+972*, in practice, human verification was limited to confirming whether the target was male, and this had to be completed within a matter of seconds. They were under pressure to identify more new targets, and when they were unable to do so, they sometimes lowered the threshold for judgment.

According to the sources, one out of every ten targets assassinated under this system was not actually a Hamas operative. However, the Israeli military reportedly considers this margin of error acceptable. The military randomly selected samples from Lavender's predictions and had them manually verified. Once this manual verification confirmed that the system's predictions were 90% accurate, the military approved its widespread use. In other words, a roughly 10% error margin was factored in from the outset. One source remarked, "the protocol was that even if you don't know for sure that the machine is right, you know that statistically it's fine. So you go for it," adding, "Everything was statistical, everything was neat–it was very dry."

Individuals placed on the "kill list" are subsequently tracked using a system called "Where Is Daddy?" to estimate their location. In many cases, the timing is planned so that they are bombed along with their homes when they are most likely to be at home. This is because it is easier to estimate the probability of the target being present at their residence. As a result, family members, often women or children, are also caught in the

attack. There were also instances where the target was not actually at home.

The "statistical" and "dry" approach is consistent in estimating collateral damage (harm to non-combatants or civilian infrastructure caused during military operations). Previously, intelligence personnel took significant time to verify collateral damage, but since October 7, 2023, this verification process has been automated, significantly reducing human involvement. The margin of error in the automated model was tolerated due to the speed of processing. "The collateral damage calculation was completely automatic and statistical," one source noted, adding that the figures generated were not even whole numbers. Since October 7, the level of acceptable collateral damage has reached unprecedented levels, with the sacrifice of up to 10 civilians for a low-ranking operative and up to 100 civilians for a high-ranking operative being tolerated. For context, during the U.S. military's fight against ISIS in Iraq, collateral damage exceeding 15 civilians required special authorization from the U.S. Central Command commander.

3 What Automated Profiling Does

The case of automatic profiling of Gaza's population by the Israeli military is an extreme example, but AI-driven profiling has become commonplace across various domains. The reason for utilizing AI profiling is to more reliably secure gains or avoid losses–in other words, for the purpose of "risk management".

Here, "risk" refers to the product of the magnitude of gain or loss associated with an event and the probability of that event occurring. In everyday language, "risk" usually refers to the possibility of loss or harm, but in the field of risk engineering, even the possibility of gain is referred to as "risk."

Risk management, broadly speaking, involves (1) identifying the available options for decision-making, (2) estimating the level of risk associated with each option, and (3) choosing an action based on these estimates. Step (2) in this process is referred to as "probabilistic risk analysis", or simply "risk analysis". Artificial intelligence is primarily used for risk analysis, specifically for estimating the probability of certain events occurring.

AI used in fields such as marketing, justice, border control, and the military serves the purpose of risk management. For instance, in companies like Amazon, recommending products to users can result in profit if it leads to a purchase, or a loss of potential profit if ignored. Therefore, selecting and recommending products that users are likely to purchase is crucial. Companies achieve this by profiling users based on their data and estimating the probability that a specific recommendation will lead to a purchase.

For Israel, leaving Hamas operatives active poses the risk of attacks on its citizens, making it critical to neutralize as many Hamas operatives as possible. To achieve this, the Israeli military profiles Gaza residents and estimates the probability that a given individual is a Hamas operative.

The use of AI-driven profiling has already been criticized for perpetuating and reinforcing discriminatory stereotypes and practices present in society, promoting privacy violations by

corporations and governments, and producing self-fulfilling effects[1] (cf. [5, 8]). However, the aspect I wish to focus on here is how the use of AI-driven profiling influences our perception of others and our social relationships.

Sociologist Shoshana Zuboff argues that the proliferation of machine learning systems based on big data leads to a society where "machine processes replace human relationships so that certainty can replace trust" ([14], p. 351). Zuboff is talking about AI tools that can not only predict but also transform people's behavior, but even profiling tools alone can have a significant negative impact on trust. When tools are available to express the risks associated with engaging with an individual in clear numerical terms, it becomes difficult to interact with others without relying on those tools, or to see them as something more than a risk factor–someone with depth and individuality.

This becomes particularly concerning when such tools are used to identify individuals deemed as risk factors with the aim of avoiding or eliminating them. At the same time, I do not believe this issue is unique to artificial intelligence or something initiated by it. Rather, it is a societal trend that predates AI, one with deep roots, but which AI has the potential to greatly accelerate. This makes it an even more complex and challenging issue to address.

4 Trust and Assurance

To trust someone in a certain matter, broadly speaking, is to have the expectation that "this person will do such and such" and to base one's own actions on that expectation. For example, if an acquaintance asks to borrow money from me and I decide to lend it with the expectation that "they will probably pay it back eventually," this means I have trusted that acquaintance to return the money.

The degree of certainty in one's expectations of another can vary depending on the situation. In fields such as sociology, it is generally not considered trust if one takes for granted that the other person will not betray their expectations. Trust involves making a decision based on one's expectations, while also acknowledging to some degree the possibility that the other party might not act as expected. For instance, when I order coffee at a caf 辿, I am not said to "trust" the barista to serve me coffee because I typically do not even entertain the possibility that this expectation might be unmet. In other words, trust can be understood as "expectation accompanied by a certain degree of doubt."

Social psychologist Toshio Yamagishi [9] proposed distinguishing the state in which one does not need to consider the possibility of one's expectations being betrayed as "assurance," rather than "trust." Such assurance might arise, for example, due to the enforcement of laws, strong mutual dependence, or close bonds within a group. According to Yamagishi, unlike assurance, trusting someone means acknowledging the possibility of betrayal or deception by the other party. However, trust arises when one evaluates the other person's internal characteristics–such as their intentions or emotions–and concludes that they are not likely to betray or deceive.

[1] A phenomenon where a prediction causes itself to become true. For example, rumors of a financial institution being at risk of collapse may trigger a bank run, ultimately leading to its failure.

According to Yamagishi, traditional Japanese society was one in which lies and betrayals were suppressed by fostering strong bonds (commitments) within closed, fixed interpersonal relationships. This created a society of assurance. However, in such a society, the costs of rigid relationships were significant–for example, collusion hindering the proper functioning of market mechanisms. Yamagishi argues that Japan's old, closed, assurance-based society is now irreversibly collapsing. Therefore, Japanese society must transition into a new, open, trust-based society.

Throughout history, humanity has always been compelled to address the risks posed by others. Consequently, humans have long developed and relied on various mechanisms to meet this need, continuously refining and inventing new ones. What I refer to here as "mechanisms" encompasses innate cognitive, emotional, and behavioral tendencies, culturally formed and learned customs, moral norms, social systems such as politics and law, and technology.

These mechanisms can broadly be divided into two categories. The first is those that enable the identification of individuals posing risks, allowing us to avoid or exclude them. This includes, for example, the ability to detect subtle signs of another person's emotions or thoughts from their facial expressions or behaviors. Norms surrounding etiquette and manners serve as clues to reveal the nature of individuals, while criminal investigations and judicial processes assist in identifying and isolating dangerous persons.

The second type of mechanism imposes constraints on people's actions in some way, making it difficult for them to betray or deceive others. This category includes psychological and emotional factors such as loyalty and the sense of reciprocity, as well as penalties for causing harm to others.

AI-driven profiling can be considered a technology aimed at "avoidance." Will the advancements in AI allow us to reliably identify trustworthy individuals in advance, and then will we live in a society so filled with assurance that trust becomes unnecessary?

5 The Pitfalls of Artificial Intelligence

AI-driven profiling, if sufficiently accurate, could serve as an exceptional risk-avoidance tool. By using it, one could not only detect when someone is trying to deceive in the moment and avoid immediate risks, but also eliminate potential future risks even before someone entertains malicious intention. Such AI could reveal whether someone will demonstrate the required abilities, work diligently, exhibit violent tendencies, contract certain diseases, or quit their job prematurely. It might become a silver bullet that exposes and neutralizes the "werewolves" lurking in our midst. Advocates of AI-driven profiling seem to harbor such expectations. However, I want to assert that no such silver bullet exists. The reasons for this are outlined below.

The first reason is that human personality and abilities are not fixed traits. They evolve dynamically, influenced by the environment, interactions with others, and the opportunities one is given. The difficulty of accurately predicting such changes far surpasses that of determining whether someone is lying in the present moment. While it is true that past and present environments, behaviors, and innate traits partially shape the future, they do not

fully determine it. Therefore, any predictions about the future are inherently probabilistic and cannot achieve absolute certainty.

The second reason is that profiling using artificial intelligence is not a neutral observation independent of the subject's actions. AI-driven profiling is often employed to inform decisions about whether to establish a cooperative relationship with the individual being evaluated. Such decisions frequently influence the person's subsequent behavior. Recall the concept of self-fulfilling predictions mentioned earlier. A person deemed "suitable for cooperation" by AI is likely to be given opportunities to achieve results and, by leveraging those opportunities, will be further positively evaluated. Conversely, if AI judges someone as "unsuitable for cooperation," they may be deprived of the chance to collaborate and produce positive outcomes.

When AI evaluation systems become pervasive, there is a risk that certain individuals could be consistently judged as "unsuitable for cooperation" in all areas of life, continuously stripped of opportunities. This is the situation constitutional scholar Tatsuhiko Yamamoto [11] describes as a "virtual slum." Placed in a virtual slum and repeatedly denied opportunities for cooperation, such individuals might lose the will to collaborate with others. Even when given opportunities, they might behave selfishly, unable to trust the system or others. In the worst-case scenario, they might fall into despair, harbor resentment against society, and turn to crime. This, in turn, could lead to even more negative evaluations of them, perpetuating a vicious cycle. Thus, the werewolf-exclusion system, which was expected to act as a silver bullet, may ironically end up creating new werewolves in the process.

Third, if one aims to make artificial intelligence appear flawless, it would be wise to design it as a highly "suspicious" system. This becomes clear when considering the circumstances under which errors in AI evaluations are revealed. Errors in AI risk assessments fall into two categories. The first is when AI judges someone as "low risk," leading to a decision to cooperate with them, but the individual fails to meet expectations (a false negative). The second case is when AI judges someone as "high risk," resulting in a decision not to cooperate, even though cooperation with that person would have been successful (a false positive).

As is immediately clear, in the case of false positives (where the system incorrectly labels someone as "high risk"), it is usually difficult to determine that the AI made a mistake. This is because it is hard to verify whether cooperation with the excluded individual would indeed have succeeded. On the other hand, in the case of false negatives (where someone is deemed "low risk" but turns out to be a risk), it becomes evident that the AI's judgment was incorrect. Therefore, from the perspective of those designing or operating the system, it seems natural to prioritize avoiding false negatives at all costs, even if it means tolerating some false positives. In other words, the system is likely to be designed with the principle of sacrificing specificity to achieve high sensitivity[2]. This is especially true when

[2]The terms "specificity" and "sensitivity" are commonly used in medical testing. Specificity refers to the proportion of true negative results out of all the actual negative cases. Sensitivity refers to the proportion of true positive results out of all the actual positive cases. Typically, there is a trade-off between specificity and sensitivity: a highly sensitive test (one that minimizes missed diagnoses) often sacrifices specificity, resulting in more false positives for individuals who do not actually have the disease.

the consequences of being betrayed or failing to detect a risk are severe.

In such a scenario, a werewolf-exclusion system may appear to work well on the surface but comes with hidden costs: the rejection of potential partners who could have cooperated successfully. However, this cost is conveniently obscured by claiming that everyone excluded was truly a werewolf. This is akin to how the United States labeled all individuals killed in drone strikes during the "war on terror" as "Enemy Killed in Action."

This issue is not unique to artificial intelligence. When we avoid or exclude members of certain groups based on stereotypes about those groups, we are doing the same thing. Such actions, in certain contexts, constitute severe discrimination. For instance, unfavorable treatment by police or governments based on gender, race, or illness is considered impermissible discrimination. Similarly, companies making hiring or promotion decisions based on these attributes engage in illegal discrimination.

The problem with artificial intelligence lies in its potential to inadvertently learn and reinforce societal stereotypes without developers or users realizing it and in its ability to create newly marginalized groups. Suppose AI, analyzing vast amounts of data, learns that "50% of individuals exhibiting certain characteristics are werewolves" (these "certain characteristics" might be too complex for humans to comprehend). This would lead to individuals matching those characteristics being deemed "dangerous" with a 50% probability and excluded.

But is it fair to exclude someone based on a probabilistic assumption about their danger? Such an action raises profound ethical and social concerns about fairness, bias, and the unintended consequences of relying on AI-driven systems.

6 The Ecosystem of Trust

When we trust someone, we often place our faith not only in who they are at present but also in their potential for future growth and development. Positive change, growth, and development often begin precisely with that expectation because many people strive to live up to the trust placed in them. While this resonates with our everyday experiences, a study involving young children has empirically demonstrated that trusting a child can foster their honesty[13].

In this sense, trust is also self-fulfilling. By trusting someone, we cast positive expectations upon them and offer them opportunities for growth. In doing so, we help elicit positive actions from them. Conversely, distrust also carries its own consequences. By choosing not to trust someone, we communicate negative expectations and simultaneously distance them from opportunities to act positively. In doing so, we risk stifling their potential for growth.

Unlike artificial intelligence, humans do not always base their trust on solid evidence. We sometimes trust individuals with no proven track record or even strangers. Moreover, we often forgive those who have betrayed our expectations in the past and continue to place trust in them. This unique human ability to trust beyond evidence reflects a forward-looking hope rather than a reliance on certainty.

To borrow the words of philosopher Nobutoshi Nagamori, trust has the "power to bind the other to 'become' a subject worthy of trust," and it "seeks to fulfill its expectations

by aiming toward the future" ([3], p. 6). Nagamori illustrates this idea by referencing the actions of Bishop Myriel toward the protagonist Jean Valjean in Victor Hugo's novel *Les Mis* 辿 *rables*.

When Valjean, newly released from prison, meets Bishop Myriel, he harbors deep resentment toward society. Despite Valjean's theft of his silverware, the bishop forgives him, even giving him additional gifts. This act of compassion inspires Valjean to transform his life. However, such outcomes are not universal. The neuroeconomist Paul J. Zak's research using the "Trust Game" indicates that around 5% of individuals are "unconditional non-reciprocators," highlighting the inherent risks of extending trust.

In JoJo's Bizarre Adventure, a similar act of forgiveness occurs when Lord Joestar pardons Dario Brando for stealing his ring. However, unlike Valjean, Dario does not repent. Instead, his selfishness leads to further harm, culminating in his son Dio bringing tragedy to the Joestar family. This contrast underscores the unpredictable outcomes of trust and forgiveness. Trust may transform some werewolves into virtuous villagers, but not all werewolves will change. Therefore, we must strive to reduce the uncertainty in others' behavior using every mechanism available to us, carefully evaluate whether someone is worthy of trust, and, ultimately, act by placing a measure of faith in our judgment despite lingering uncertainties.

What if the world were populated only by people like Bishop Myriel, Jean Valjean, and Lord Joestar? Such a society would undoubtedly be remarkable, but it would not be stable or enduring. Eventually, someone like Dario Brando would emerge, taking advantage of people's goodwill and tolerance to reap significant benefits. The Brando lineage would thrive, while the descendants of the Joestars would face adversity. However, over time, a new type of individual–different from both the Brandos and the Joestars–would emerge. These individuals would not trust everyone unconditionally, but neither would they refuse to trust anyone. They would carefully evaluate whether others are trustworthy, and if they judged someone to be reliable, they would place their trust in them and establish strong cooperative relationships. They would feel a sense of obligation to reciprocate favors and would not hesitate to respond appropriately, even with anger, when betrayed. In terms of the characters from *JoJo's Bizarre Adventure*, this type would resemble Speedwagon.

The Joestars and Speedwagon-like individuals would ultimately fare better than those like the Brandos, who are constantly scheming and betraying others. This is because the Joestars and Speedwagons can achieve the benefits of cooperation that the Brandos, with their untrustworthy nature, cannot obtain. While they may occasionally suffer losses due to betrayal, the overall outcome would likely favor the Speedwagons and Joestars, as their cooperative efforts would lead to greater long-term success compared to the self-serving Brandos.

Traits such as tolerance, quickness to anger, persistence, gullibility, skepticism, caution, openness, strong group loyalty, a sense of obligation, cunning, and others are personality characteristics related to trust. Society is composed of individuals who exhibit these traits to varying degrees. Moreover, even within a single person, these traits are not fixed. They may manifest differently depending on the situation or evolve under the influence of others. This dynamic interplay of trust-related traits forms the ecosystem of trust that humanity

has developed over a long period of time.

This ecosystem has likely been forced to adapt repeatedly in response to changes in social structures and technology. I believe that the widespread use of AI-driven profiling has the potential to significantly disrupt this ecosystem. The consequences of such a disruption are likely beyond anyone's ability to predict. Moreover, it will likely constitute a "transformative experience" (cf. [6]) that alters our values and preferences, making it difficult to evaluate in advance whether it will be ultimately good or bad.

However, I fear the realization of the following scenario: the designers and users of werewolf detection systems are motivated to prioritize sensitivity at the expense of specificity (as in the case of Lavender), leading the system to make more "suspicious" judgments. This not only incurs the hidden cost of obstructing potential cooperative relationships but also risks turning innocent villagers into actual werewolves after they are falsely identified and excluded. This transformation, in turn, reinforces the perceived necessity and justification of the werewolf detection system itself, creating a self-perpetuating cycle where society becomes increasingly dependent on such systems.

7 Conclusion

This paper has discussed the issues surrounding the rapidly advancing field of artificial intelligence, particularly systems that use big data to perform human profiling.

The increasing use of AI systems that profile humans based on big data may encourage us to treat others as machine–readable bundles of data and inferred attributes. Additionally, AI may promote an attitude of evaluating others primarily as sources of risk–that is, from the perspective of how much benefit or harm they could potentially bring. This phenomenon could be described as a "culture of risk analysis," which may negatively impact trust in human relationships.

However, this trend is not an issue unique to AI. Instead, it is an extension of the modern tendency to measure and quantify human performance as much as possible. Given this trajectory, the growing ubiquity of AI in various domains seems inevitable. Organizational managers may increasingly see it as their duty to use AI to estimate and mitigate risks. Just as management consultants promote various performance metrics to administrators, vendors of AI systems may come to dominate a significant portion of the "bullshit job market" in organizational management.

On the other hand, those being managed will likely strive to earn favorable evaluations from AI systems. Much like how web administrators aim for search engine optimization (i.e., improving their ranking in search engines), individuals may aim to optimize themselves for AI-driven evaluations. This focus risks obscuring the original goals that these evaluations and improvements were supposed to address. People labeled as "high risk" by widely used AI systems may find themselves barred everywhere they go, effectively relegated to a "virtual slum."

Perhaps I am being overly pessimistic. Yet, as AI becomes increasingly prevalent in various spheres of life, history suggests that it will inevitably reinforce our already excessive obsession with measurement, while fostering cultures of prediction and risk analysis. In

this context, it is vital to ensure that a humanistic perspective on humanity–one that sees humans not merely as sources of bias but as sources of value (cf. [4])–is not entirely lost. We must also remember that believing in the potential for human growth, even in the absence of evidence or in the face of contradictory evidence, is never an irrational stance.

References

[1] Yuval Abraham, " 'Lavender': The AI machine directing Israel's bombing spree in Gaza", *+972 Magazine*, April 3, 2024. https://www.972mag.com/lavender-ai-israeli-army-gaza/

[2] Grégoire Chamayou, *Théorie du Drone*, Fabrique, 2013.

[3] 永守伸年、『信頼と裏切りの哲学』、慶應義塾大学出版会、2024 年。(Nobutoshi Nagamori, *Philosophy of Trust and Betrayal*, Keio University Press, 2024)

[4] 隠岐さや香、『文系と理系はなぜ分かれたのか』、星海社新書、2018 年。 (Sayaka Oki, *Why Were the Humanities and Sciences Divided?*, Seikai-Shinsho, 2018)

[5] Cathy O'Neil, *Weapons of math destruction : how big data increases inequality and threatens democracy*, Crown, 2016.

[6] L. A. Paul, *Transformative experience*, Oxford University Press, 2014.

[7] J. Scahill and The Staff of The Intercept, *The Assassination Complex: Inside the Government's Secret Drone Warfare Program*, Simon & Schuster, 2017.

[8] Carissa Véliz, *Privacy Is Power: Why and How You Should Take Back Control of Your Data*, Transworld, eBooks edition, 2020.

[9] 山岸俊男、『信頼の構造: こころと社会の進化ゲーム』、東京大学出版会、1998 年。(Toshio Yamagishi, *The Structure of Trust: Evolutionary Game of Mind and Society*, The University of Tokyo Press, 1998)

[10] 山岸俊男、『安心社会から信頼社会へ——日本型システムの行方』、電子版、中央公論新社、2013 年。(Toshio Yamagishi, *From Assurerance to Trust: Where Japanese System Will Go*, digital edition, Chuoukouronn Shinsha, 2013)

[11] 山本龍彦、『おそろしいビッグデータ』、朝日新聞出版、2017 年。(Tatsuhiko Yamamoto, *Big Data Is Terrible*, Asahi Simbun Press, 2017)

[12] Paul J. Zak, *The moral molecule : the new science of what makes us good or evil*, Transworld Publishers, 2012.

[13] Li Zhao, Haiying Mao, Paul L. Harris, and Kang Lee, "Trusting young children to help causes them to cheat less", *Nature Human Behaviour*, 2024 Apr;8(4):668-678. doi: 10.1038/s41562-024-01837-4. Epub 2024 Feb 20.

[14] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, Profile Books, 2019.