

THE ONTOLOGY OF THE ENEMY: AN INTERVIEW WITH PETER GALISON

SINA NAJAFI

In September 1940, with the German air raids over Britain at their peak, M.I.T. mathematician and physicist Norbert Wiener wrote Vannevar Bush, in charge of American war research, and volunteered his services. Over the next few years, Wiener focused on making what he called the anti-aircraft predictor, a computational device designed to improve the accuracy of ground-based gunners by calculating more precisely the location of enemy aircraft at a given point just moments in the future. The success of the project was in part dependent on the type of enemy that Wiener imagined was piloting the airplane. As Peter Galison, Mallinckrodt Professor of Physics and History of Science at Harvard University, writes, Wiener gradually came to see the predictor "not only as a model of the mind of an inaccessible Axis opponent but of the vast array of human proprioceptive and electro-physiological feedback systems." This expanded model became the basis of a new science that Wiener named "cybernetics." *Cabinet's* editor-in-chief Sina Najafi spoke to Galison about the implications that Wiener's understanding of the enemy had for the subsequent evolution of cybernetic theory.

You've written about the way cybernetics is rooted in Norbert Wiener's World War II anti-aircraft predictor, and in particular his notion of the enemy.

Yes, what interested me was that the category of "the enemy" covers many different kinds of relationships. Enemies are not alike. One version of the enemy, for example the way the Japanese were viewed by the Americans and the British, was the racialized, monstrous, sub-human other. Another distinct Allied vision made the enemy anonymous, the example here being the pilot who bombs a city several tens of thousands of feet below and has no psychological connection through empathy or fellow-feeling. The third category, which is what Wiener was interested in, is an enemy who is much more active than the anonymous invisible inhabitants of a distant city and more rational than the horde-like racialized enemy. This enemy, more complex than either of the other two perspectives, was a cold-blooded mechanized Enemy Other who made at least predictable moves that could be modeled through some kind of black box machinery. ("Black box" is a term that comes from World War II from the era of radar work and it meant that some circuit inside the box performed a certain way without you needing to understand how that circuit was laid out.) This picture of the enemy that emerged through World War II is less well-known but in many ways more powerful and enduring than the first two notions of the enemy. This is not to say that the kind of enemy that Wiener's dealing with somehow supplants those other forms, but it is different. And the distinction is important for understanding what we're considering and what the consequences are.

Wiener's ambition was to make a black box model of the enemy pilot and then use that to form an anti-airplane system that could characterize the pilot's movements and learn from past experience in order to predict where the anti-aircraft gun should aim. Some prediction was necessary because it took up to 20 seconds before anti-aircraft fire reached the airplane. So if you shoot at where the plane is now, you'll surely miss it. If you extrapolate to where it might be if it were to travel in a straight line, you'll also miss it if the pilot moves from side to side. So it was necessary to be able to predict where the pilot would go, and that's what Wiener's machine was designed to do. What Wiener did that was so unusual was that he characterized the motion of a particular plane using a primitive computer. The radar follows the plane, looks at its motion, makes a statistical characterization of what the pilot's been doing in the last several tens of seconds, and then uses that knowledge to predict where the pilot will be 20 seconds later. The underlying mathematical methods Wiener used for the AA predictor were carried over from his earlier mathematical studies of servomechanisms, that is to say, feedback devices such as thermostats and self-guided torpedoes. Feedback is the basis of any self-correcting system, and with no access to anything in the enemy plane, Wiener simply treated the pilot-plane assembly as a complex machine—a servomechanism—that once characterized could be simulated in order to predict what it would do. And then it could be blown out of the sky.

So the statistics are absolutely particularized. It's a model that begins each time with an empty data set which fills up based on that particular pilot's run.

It's based on that particular run, with that pilot, at that time. In that sense, the machine was learning and that's what intrigued Wiener. Interestingly, he found that you couldn't take data from one pilot and use it very successfully to predict where another pilot would be in 20 seconds. In a sense, what he had done was to create a black box that learned how to behave the way a particular pilot would in the future. He thought this would have practical consequences, which, as it turned out, it didn't because he never got it to work more than a couple of seconds in advance; 20 seconds out or even 10 seconds out was too far. But the scientific administrators cleared to see Wiener's device found that its predictions for even those few seconds into the future were remarkable; more than remarkable, positively uncanny. The machine appeared to anticipate a person's intentions, to stand in for human intentionality in some fundamental way. It seemed astonishing to Wiener, astonishing enough to merit the foundation of a new science.

Because of the antagonistic relationship between attacker and defender, the anti-aircraft operator was obviously in no position to talk to or even see the pilot; he was a shrouded, hidden entity whose only features were those you could see in the control of the airplane. And since those airplane motions could be duplicated in advance with this learning black box, it seemed to Wiener and those around

him that he had created something out of circuits and gears that was practically miraculous—an electro-mechanical version of human intention. He had made a machine capable of learning to be this pilot or that.

Do you need to distinguish between the pilot's idiosyncrasies and the plane's technical limitations?

No. And this is a crucial point. All you care about is what the human-plane system does. From the point of view of the radar operator, the pilot was so merged with machinery that his human/nonhuman status was blurred. Imagine that the plane had a rocking motion from left to right every two seconds—whether this resulted from an instability in the plane's design, a mishandling by the operator, or a periodic gust of wind made strictly no difference. All Wiener was concerned with was the net total of what the input into the plane was and how the plane responded.

Would the gunner on the ground also be integrated into the system?

The gunner himself has certain properties. He has a lag time between his decision to turn his anti-aircraft gun and how fast he can issue the command; a lag between the moment he hits the controls and the time the gun actually turns, and so on. But all these decisions and delays are subject to the same sort of analysis as the plane. So Wiener wants to characterize the motion of the gun, including the firing of the shell, as an amalgamated combination of the psychological, the instinctive, and the mechanical. From the point of view of Wiener's analysis, all you want to know is this: when the gun is told to fire at a certain angle and certain elevation, how soon does the shell arrive at its aim point? This dynamic becomes part of the system, just the way the enemy pilot of the plane does. The core lesson that Wiener drew from his anti-aircraft work was that it was essential to conceptualize the pilot and the gunner as servomechanisms within a single system. In Wiener's way of thinking you treat your "own" gunner in the same way as the enemy pilot. This is not because you couldn't ask the gunner what he's thinking (under some circumstances you could), but once you've started the analysis of the "man-machine" system in this way, the pilot and the gunner seem alike—they are both servomechanisms with feedback. Ally and enemy begin to resemble each other in a war of human-machine hybrids.

Insofar as Wiener's model makes no qualitative distinction between the enemy pilot and your own gunner, it seems like a complete break from the racialized or anonymous enemy.

Yes, there is a reflexivity here: in creating an enemy of this type, a new version of ourselves comes into being. This relates to other work I have done in which I've argued that we began to conceive of our own cities as potential targets.

During World War II, the US had a bombing survey, an immense operation with thousands of people

analyzing data, planning the bombing, and experts from all the different industries. Oil industry experts were recruited, so too were consultants from the ball-bearing industry, and so on—each advised on the critical nodes to be destroyed in the corresponding German industry. After the atomic bombing of Hiroshima and Nagasaki, the bombing evaluators walked through the rubble and began to wonder in their official reports: "What if this happened to our city? What would happen if this part of the oil industry was destroyed or what would happen to the overall production of steel in the US if central Pittsburgh was laid to atomic waste." As a result, during the early Cold War, each US city began to think of itself as a target, to organize maps in such a way as to predict exactly what consequences would follow from nuclear bombs hitting in central locations. City planners learn to think of their cities as laid-out like a target in order to create planning directives to specify where people could cluster housing, factories, and other functions. Essentially they learned how to disperse the city, to use the German strategy learned in the war, back home. Of course there were many other factors driving suburbanization, but this fear of central destruction added regulatory and tax incentives to its acceleration.

But the underlying process I want to emphasize is this: there is a relentless cycle in which one conceives of the enemy a certain way, and then that conception begins to work back on us. The enemy as human-machine black box becomes us as human-machine black box. The enemy city targeted, bombed, dispersed becomes our city dispersed in preventive anticipation.

And this comes from imagining the enemy as a black box?

That's exactly the point. You start with an input-output approach which is suggested by the fact that you have no access to the pilot, you make this enormously successful prediction of where the plane would be, and then you apply that back to your own gunner. By now, the train of generalizations is rolling and you can see then how it gets generalized to other systems. Are there are other kinds of human responses or intentions that could be modeled in this way? Soon Wiener's thinking that intention is nothing but a certain characteristic relation between input and output, which is called feedback. During World War II, a whole world of self-correcting devices were being developed and Wiener began to think, "Maybe that's all there is to intentionality."

You have written about Wiener's distinction between the Manichean devil and the Augustinian devil. How do the three different ontologies of the enemy that you've outlined relate to the Manichean and Augustinian devils?

The distinction between the Manichean and the Augustinian devils are categories within the kind of enemy that interests Wiener. He makes a distinction between an enemy who simply calculates but who cannot change the rules (the Augustinian devil) and an enemy who is cunning, who is trying to outwit

you, and can bluff and change strategy to obtain victory (the Manichean devil). For Wiener, the Augustinian enemy can actually be nature itself. That is to say, when we are trying to figure out the laws of physics or astrophysics, trying to unlock a set of rules that are hidden to us, nature isn't trying to outwit us. (Paraphrasing Einstein, nature is subtle but not malicious.) The scientist is fighting what Wiener calls "the devil of confusion," which he keeps distinct from "the devil of willful malice." A Manichean enemy is cunning in a different way; a Manichean opponent in chess or war can fake a move left in order to deceive—and then move decisively to the right.

Going back to your discussion of intention, would the ability to predict necessarily imply intentionality? Nature, of course, would be a good example of something that we can predict, at least to a degree, but which we don't normally ascribe intentionality to.

It's not that for Wiener all predictions stand in for intentions; instead, he wants to argue that there is nothing about intention that can't be replicated by a sufficiently sophisticated predictive computer. In other words, he does not see any advantage in maintaining the sanctity of a mental state that corresponds to intention: all he cares about is whether the pilot's actions can be foretold. If the machine can do this, Wiener contends, then intention has been captured. Let me step back. What Wiener thinks of himself as doing is extending behaviorism, which for him encompassed a field of inquiry much wider than a behaviorism defined as a strict, literal replication of the doctrines of J. B. Watson and B. F. Skinner. Wiener thought that behaviorism had certain good ideas, namely the elimination of mental states, which seemed to him, as to all the behaviorists, a form of metaphysical mumbo jumbo. There were many criticisms of behaviorism, not least that it could not account for many things, including intention. But Wiener claimed to be capable of expanding the old concept of behaviorism in such a way that concepts as subtle as intentionality could be reduced to observable results. If we look at self-correcting systems that tend towards a goal, we can say, "That is intention." If I watch you as you try to walk through a room and pass through a narrow door, I see you making little corrections toward the door and I can say, "It is your intention to walk through the door." Intention would, according to Wiener, be the self-correcting motion that you exhibit. Nothing more.

But would a heat-guided missile have intentionality?

This is exactly what some of Wiener's critics worried about. For example, in 1950, Richard Taylor, a philosopher at Brown University, asked if Wiener and his collaborators were serious in proposing a definition of intentionality or purposefulness that was built purely on the culmination of a sequence of events. It's worth quoting the definition that Wiener and his colleagues published in a 1943 paper:

"The term purposeful is meant to denote that the act or behavior may be interpreted as directed to the attainment of a goal—i.e., to a final condition in which the behaving object reaches a definite correlation in time or in space with respect to another object or event. Purposeless behavior is that which is not interpreted as directed to a goal."

For Taylor, this definition was so all-encompassing as to rule out nothing but also so devoid of content that it had no overlap with any common meaning of the term. So if a clock runs for many years and breaks down at midnight on New Year's Eve, why would that not show intentionality? Or how about a brick that falls off a building and kills a passerby? Or could a roulette wheel—the paragon of purposelessness—be made into a purposeful machine by adding some lead weight to its perimeter? In a postwar response to Taylor, Wiener and one of his colleagues made no apologies for classifying a crooked roulette wheel as purposeful but they emphasized that the weighted wheel is different from the servomechanisms of guided missiles and AA predictors because the wheel is passive but the latter are active. (The weighted wheel doesn't learn, the predictor does.) And Wiener insisted that "as objects of scientific enquiry, humans were no different from machines." In 1941-42, from a military perspective, it made sense to Wiener to see humans as not any different from machines. But by 1950, Wiener had globalized his claims so that human intentionality was not different from the self-regulation of machines, full stop. So when you're driving your PT boat in the South Pacific with a Japanese guided torpedo hot on your wake and you say, "That torpedo is trying to get me," as far as Wiener is concerned, you're actually saying something that makes perfect sense. If you think that all you have access to is the outside and you don't want to attribute mental states, it shouldn't be any different whether there's a little man inside the torpedo steering it or the torpedo is probing magnetic fields to determine which way to turn.

But Wiener did not stop with intention. Already in November 1944, the Harvard psychologist and historian of psychology Edwin Boring had approached Wiener to see if other kinds of previously construed mental categories could also be reduced to something characterized by the outside, mathematized and modeled using circuits, and so on. Wiener said, "Sure, give me the list of supposedly mental states and we'll reduce them one by one through models and circuits to something that can be measured from the outside." The list that Boring handed over to Wiener consisted of 14 psychological properties, with others like "Generalization" and "Abstraction" added later. Black-box engineering now had a more complex goal, namely to recreate the mind itself.

At the first meeting of the Teleological Society in 1945, which formed the basis for the later cybernetics conferences, it's interesting to see that philosophers seem to be missing despite the fact that basic philosophical concepts such as intentionality were being discussed.

The philosophers did get involved soon. But there's something that needs to be clarified about what's happening in philosophy during this period. Beginning in the 1920s with the Vienna circle, we see the rise of a scientific philosophy, or what becomes the philosophy of science later on with Rudolf Carnap, Ludwig Wittgenstein, and so on. When the Nazis came to power, that group shattered and some of them take refuge in the United States. Carnap goes to Chicago, Philip Frank goes to Harvard, and so on. This group begins to train a new generation of people who are interested in these issues. All of this is still fairly new, so when you get to 1946, the philosophy of science as we would understand it—that is to say, people who think of both philosophy's mode of inquiry and its object of inquiry being closely associated with the sciences—is something quite new. So there isn't an established body of philosophers of science who would be natural allies for the cybernetics group. But they did arrive; some of the first ones had worked with Carnap and others in this early generation of American scientific philosophy. There are also anthropologists, people like Gregory Bateson, Margaret Mead, and others, who begin to take it up and reformulate their ideas (from the social sciences) in terms of cybernetic notions of feedback loops and so on. John von Neumann gets involved and begins to carry over the ideas into computation and the building of the first electronic computer.

All these ideas get very closely associated and it's the cause of some friction in fact, but you have the development of computer science and what becomes cognitive science. There are also physiological approaches, like some of Wiener's collaborators who are using these ideas to model the heart and, even more ambitiously, nerve disorders. All of this begins to develop quite rapidly. Cybernetics becomes a hugely popular movement and, in a way, burns itself out fairly quickly in the US. In the Soviet Union, it has a very different trajectory: rejected in the early years but then promoted beyond all measure, almost to a state philosophy. It has had a revival recently in the US, partly through developments in computer science and partly through the interest these ideas have had in the social sciences and in the humanities for people from Donna Haraway to Niklas Luhmann. It's had a tremendous echo in sociology, literary theory, and so on as people come back to Wiener's standpoint in which the human and the non-human merge, where the human is de-essentialized and nature is de-sacralized.

To what extent do you think the original context for these ideas coming out of the war marks the later development of cybernetics? To what extent is it possible to develop this kind of black box system of the mind without that original set of parameters being necessarily present?

As a historian, I would never say that something couldn't have happened otherwise, because there are a lot of ways things can happen. But I think what you see in World War II that's so important is that when the war ends, there are thousands of engineers and scientists who have

had the experience of working with black box systems, with feedback systems, and with electro-mechanical means of organizing goal-directed behavior for objects. You suddenly have a gigantic group of people who understand exactly what Wiener is talking about. It seems a kind of artificially-produced worldview—though I don't know what artificial or natural means in this context—because in the radar program alone suddenly \$2 billion dollars go into producing systems that have an immense impact. At a first approximation, every single citizen involved in the war effort has worked on one of three projects: the Rocket project, the Radar project, or the Atomic bomb. Black box and feedback ideas are everywhere; they too have become part of the lingua franca of engineers.

How easy is it to shift the cybernetic model into another territory, one that's distanced from these notions of the enemy, of combatting and controlling the other?

Well, I think that the originary experience of trying to understand an inaccessible but calculating enemy as associated with these technologies is very deep. But then, of course, the whole second part of the story is that cybernetics does get expanded in all sorts of ways. Bateson is not trying to kill people with torpedoes; he's trying to use it to understand the dynamics of society and re-describe some of his ideas about the double-bind and other things in terms of feedback.

And the question of the recognition of the other is presumably also not important anymore once cybernetics enters the social sciences and so on.

For reasons that are not the same as the reasons of warfare, the people in the humanities and social sciences have also wanted to achieve a certain scientificity, a blurring of the boundaries between the human and the non-human. They've wanted to get away from the intentional subject in philosophy, to get away from the Sartrean picture of the human as defined by the self-establishment of projects and goals. One shouldn't attribute all this to cybernetics, but cybernetics enters the scene at a moment when there are lots of reasons for people to move away from a certain concept of what appears as a more and more romantic ego-based humanism. In France, that typically takes the form of a move toward structuralism, but one of the ways it is manifested in the Anglo-American world is through these models of cybernetics precisely because of this shared experience of warfare, technology, and so on. That's also true in the Soviet Union. What a feminist American theorist like Donna Haraway is reacting against is the romanticized Mother Earth picture where the masculine is identified with the artificial-technological and the feminine with some sort of eternal natural and there is a radical distinction between them. She sees in cybernetics a way of describing the world without that form of what she considers to be romantic feminism.

It's not that I want to resurrect this romantic conception of the eternal masculine and feminine. Not at all. Instead, I want to emphasize that when you invoke cybernetics, you're

invoking more than simply a lack of differentiation between the human and the non-human. You're participating in a wider set of moves and associations and apparatuses that go much beyond that. In a way, this is a problem for understanding historical developments of culture more generally. At one extreme, you can say that you can always appropriate this or that element of a movement, disregarding how it was configured in its original context. At the other extreme, you can say a movement has infinite inertia, that the moment you talk about some element of Nazism, you buy the whole package. Both extremes seem to me unhelpful. The Nazis developed life jackets in concentration camps but it's clear that by using a life jacket you're not invoking the totality of the Nazi system of annihilation. On the other hand, there is a problem with the belief that we can simply invoke C³I, Command-Control-Communication-Intelligence, or cybernetics as if they were not embedded in a military way of organizing things or a concept of the enemy. My cautionary note is that when you buy into cybernetics as a model for some sort of post-humanism, you get a lot more with it than simply the difficulty of distinguishing between human and non-human.

Jean-François Lyotard's *The Postmodern Condition* is one of the most famous appropriations of cybernetics within critical theory, but you point out that his reading of cybernetics is at heart a radical misreading.

Lyotard sort of parodies cybernetics and then talks about what's supposed to succeed cybernetics in terms that really are not just part of but *central* parts of the conception of cybernetics that existed in the 1950s and 1960s. He says they have no feedback system and he also claims that cybernetics did not have an agonistic theory of society, so at both the theoretical level and the historical level, he seems to miss what cybernetics is about and how it developed. The interest in agonism over the last 25 years has been, in someone like Foucault for example, a way of de-romanticizing human actions. In the tradition that goes back to Nietzsche, there is an argument that you can't think of morality, sexuality, or any other category of inquiry as eternal. Instead, we have learned more and more to see our categories as the outcome of a balance of forces. There is a sense that, in de-romanticizing in this way, one is removing an unacknowledged theological grounding to various aspects of the human condition. This agonistic conceit finds a perfect home in cybernetics. But I think there's a danger in thinking that the alternative is either a kind of romantic association with some sort of eternal and sacred qualities or the constant agonism of opponents struggling for goods. There are forms of interaction that are of another kind.

Jean Baudrillard's contention that the Gulf War did not happen—that it was a simulation at least as far as the Western experience of it goes—seems to rely on a notion of an abstracted enemy. Can you comment on how the enemy is being viewed in the current conflicts?

Our war-fighting apparatus is organized around the Cold War and it's uneasily trying to figure out how to be in a world that's not structured that way. The kind of world in which Wiener developed in his ideas—where you have German engineers who basically think just like you, only they're trying to kill you instead of you trying to kill them—is very much like the Cold War. After the Cold War had more or less ended, American and Russian nuclear weapons designers began having conferences together, joking about how they had, for so many years, tried to guess what the other one was doing. They employed many of the same methods to get their political apparatuses to give them the resources they needed. They were, so to speak, on the same page. I think one of the enormous shocks of September 11—it's certainly been said before—was that no one had box cutters in mind. Nothing was organized to defend airplanes against box cutters and against attackers who did not want to live. The cybernetic model in some way grew up around an enemy that was symmetrical, which is why you could move easily from the enemy taking evasive action to you taking evasive action or to the gunner taking anti-evasive action, but it was all in some sense a symmetric piece. One of the problems that arises in the current situation, in addition to all the immense political and moral complexities of race and religion (not to speak of oil), is that the war machines of Western countries are not organized to do what they're being asked to do.

Ultimately, in the recent wars, elements from all the different kinds of enemies are present. The anonymous enemy is there in the pictures of automatic weapons and self-guided missiles sending back television images as they go down a smokestack. The racial element is clearly present and shows no signs of abating. It's quite personal and delights in the Other's suffering. What of the Wienerian enemy? In the last few years, the game-playing opponent has migrated to new places. Cyberwar has become a catchphrase; but it no longer refers simply to hot shots hacking into the enemy's infrastructure or military nets. Data mining—the culling and coordination of data from the vast archive of police, travel, and financial records—has become ever more routine. We construct images of the enemy and then deploy vast searches to find the pattern in the noise, the terrorist message in a website. We hope to counteract asymmetric foes with computational dominance. But as so often has happened in the past, there is a danger that the picture we paint of the enemy becomes a normative portrait of ourselves. A risk that this apparatus of surveillance, pattern recognition, and coordination puts us all in the target position of our own data trawling. Do we really want to become correlated and open data sets as we search for an enemy data set lodged somewhere in the archive? The threat of such a reduction to sorted data is immediate to privacy, but also—not far behind—lies a risk to the fabric of the democratic civil society that has taken so many years to construct and so many lives to defend.

For an extended article on the issues addressed here, see Peter Galison, "The Ontology of the Enemy: Norbert Wiener and the Cybernetic Vision," *Critical Inquiry*, no. 21 (Autumn 1994), pp. 228-266.

PHOTO COURTESY OF THE UNIVERSITY OF CHICAGO PRESS