



## Do We Want Dystopia?

*Stefan Beck*

**I**nasmuch as there are canonical texts of American education, Aldous Huxley's *Brave New World* is one of them. But students may wonder why their teacher presents as “dystopian” a text that reads, in 2020, like an operating manual for the technocratic American Dream. The taming of reproduction and heredity by science; the banishment of boredom, discomfort, and sorrow by entertainment and pharmacology; the omnipresent availability of attachment-free sex; the defeat of death, sort of, by blissed-out euthanasia: Huxley foresaw not our fears but some of our deepest aspirations.

To read and teach *Brave New World* as dystopia is at best an oblivious atavism, at worst a piece of deluded self-flattery. As a character (not even an especially bright one) observes in Michel Houellebecq's *The Elementary Particles* (1998), “Everyone says *Brave New World* is supposed to be a totalitarian nightmare, a vicious indictment of society, but that's hypocritical bullshit.” The only thing Huxley got wrong, the character adds, is society's acceptance of genetic caste stratification. In reality, we expect “advances in automation and robotics” to render such attine division of labor as obsolete as the

sundial, the cotton gin, and the dot matrix printer.

It's easy to look back at Huxley's novel and attribute the radiant, meaningless future toward which it so fearfully looked as the realization of the dreams of scientists—including Huxley's own brother, the eugenicist Julian Huxley—with their Promethean curiosity and procrustean “solutions.” But Huxley fretted about the machinations of industry as much as he did about scientists: *Brave New World* is peppered with the surnames of Henry Ford, Sir Alfred Mond, and Maurice Bokanowski. Huxley seemed convinced that when the last irregularity was removed from the human condition, and the last inconvenience stripped from the human experience, it would be scientists' and industrialists' hands wielding the plane. But where the scientists pursue knowledge for its own sake, or in service of the good as they see it, the tech titans pursue it the better to sell us what we want. How well the would-be Aldous Huxleys of our day understand that—and how much blame they place on us and our appetites—is the subject of this essay.

**I**n Ray Bradbury's much-anthologized 1950 short story “There Will Come

Soft Rains,” the only “character” (apart from an unfortunate dog) is also the setting: a fully automated house, equipped with every modern convenience, left vacant in the aftermath of a nuclear war. It is the morning of August 4, 2026. The house goes about its customary business, preparing breakfast, reading reminders from a PA system, cleaning its carpets, watering its lawn. Much of the housework is performed by robotic mice that emerge, adorably, from what we are given to picture as cartoon-style holes in the baseboards.

A series of accidents results in a fire, and the house burns down in spite of its safety features—though not before it has read to its absent mistress the Sara Teasdale poem from which the story takes its name. That poem, about nature reclaiming a battlefield, underscores Bradbury’s pointed message, that Earth will get along just fine without man. Today we see the same type of paradoxical self-aggrandizing prostration in rhetoric about climate change, about the planet one day “self-correcting”

and spitting out its wayward children. Yet, to crow that nature does not need us is itself to anthropomorphize nature. Nature does not have conscious needs or aims; we do. The story’s staying power lies not in its self-flagellating homily but in its illustration of man’s deep wish—which, the story suggests, ultimately amounts to a death drive—to have it easy, to live inside of one big labor-saving device.

A good example of this desire is the rise of the autonomous vehicle, which is the subject of the new novel *The Passengers* by British author John Marris. The book is neither hard science fiction nor literary fiction, but rather the kind of thriller that used to give itself away with a gold-embossed title. That said, if the portentous gimmick line on the cover—*Who Lives? Who Dies? You Decide.*—fails to scare readers off, they will be rewarded with a provocative take on twenty-first-century attitudes about safety, privacy, celebrity culture, ethics, justice, and rational thought itself. What it lacks in depth

*Zed*

<i>The Passengers</i>	By Joanna Kavenna	<i>Machines Like Me</i>
By John Marris	Doubleday ~ 2020 ~ 352 pp.	By Ian McEwan
Berkley ~ 2019 ~ 340 pp.	\$27.95 (cloth)	Nan A. Talese/Doubleday
\$26 (cloth)		2019 ~ 338 pp.
		\$26.95 (cloth)



or plausibility, *The Passengers* makes up for by encouraging debate.

The premise takes a cue from that gem of mid-nineties auteur cinema, *Speed*, only with a whopping *eight* doomed vehicles and minus a hero as inspiring as Keanu Reeves. The book introduces the occupants of a number of self-driving cars, their backgrounds, personal problems, and intended destinations. (A handful are vaporized before we get to make their acquaintance.) We also meet Libby, a principled woman “with a profound hatred for all things driverless” who has been selected for a very specific kind of jury duty: adjudicating fault in accidents involving such vehicles. Driverless cars are ubiquitous in the United Kingdom of the future.

We soon discover that Libby’s hatred of autonomous vehicles is not without reason. Not only are they fallible, and not only has she been traumatized by their destructive failure, but also the inquests are a sham, consistently blaming “human error” by the circular logic that only humans are capable of error. The villain here is a callous government transportation minister, the consummate mansplainer, a symbol of how reckless and unimaginative a technocrat can be in his impatience to fix society.

This is a melodrama, and Marrs’s true Snidely Whiplash is the hacker, imaginatively codenamed “The Hacker,” who seizes control of eight autonomous vehicles, broadcasts video of their terrified passengers

over the Internet, and makes the world vote on which of them will be spared death by explosives or remotely controlled collision. One of them, to supply a romantic subplot and a sequence of successively more ludicrous twists, is Libby’s lost soul-mate, a guy she met once in a bar but failed to track down in weeks of determined e-stalking. Will the world reunite them, or will it save the pregnant woman, or the aging but beloved actress, or...will everyone end up in a fireball?

Most of the book is spent introducing the “innocent” passengers and then complicating our view of them. One woman is revealed to have her dead husband in the trunk. The beloved actress turns out to have been helping her husband cover up his sickening child abuse. Libby’s romantic interest is disqualified from the world’s mercy when it comes to light that he is suicidally depressed: If he doesn’t want to live, why should anyone else want him to? The book harbors an Augustinian sense that there is no such thing as an innocent person, and makes the reader wonder how many variables an ethical equation needs in order to have a reasonable claim upon being just.

At its crudest level, *The Passengers* points up the fact that a new technology’s dangers are often downplayed or ignored. Nothing is hack-proof, and especially nothing that connects to the Internet. (One wonders at first why these vehicles

are not susceptible to that low-tech attack, kicking the windows out; Marrs eventually remembers to mention that the glass is *really strong*.) It also encourages us to ponder the several rationales for, as Yakov Smirnoff might say, letting car drive you: increased safety, and with it decreased insurance rates and medical costs; increased productivity, as more of our time is freed up for our employers; increased laziness, as more of our time is freed up for Netflix and Pornhub, as Huxley foresaw.

The title *The Passengers* suggests passivity, but the passivity the book most successfully skewers is intellectual, not mechanical. When the Hacker forces each Passenger to justify his existence in a *Bachelorette*-like interview, and then reveals the dark secret each Passenger hopes to withhold, he is implicitly arguing that such limited, albeit sensational, input is enough to output a morally sound life-or-death judgment. His diabolical game is just the trolley problem writ large. When it turns out that the cars' own onboard decision-making software is using an economic rather than a moral calculus—it decides whether to sacrifice a driver or to remove him from jeopardy based on his relative status as a “producer”—the reader may recognize queasily that neither is just, merely ugly in different ways. Human error may be, for humans, the most palatable of several evils.

The “driverless car” represents nothing if not our attempt at a post-human legal apparatus. How difficult would such a system be to devise? What would it look like? What Marrs seems to understand, for all the infelicities of his prose and all the face-palming contortions of his plot, is that we will never be able to abdicate moral responsibility. We will never comfortably cede the gavel to an algorithm, an infinitely populous jury, a superintelligent judge, or even a robotic deity. When it comes to the complex interplay of memories, emotions, biases, priorities, and laws that govern moral decision-making, someone will always wind up with blood on his hands.

**T**he incompatibility of machine learning and machine reasoning with human morality is also at the heart of Ian McEwan's novel *Machines Like Me*, a quiet but engrossing exploration of artificial intelligence. The book is set in an alternate-history 1980s Britain, primarily so that it can include as a character an alternate-history Alan Turing, a Turing who did not commit suicide in 1954 but continued to contemplate the mysteries of the “thinking machine.” In this parallel reality, the United Kingdom is thumped in the Falklands War; Margaret Thatcher's popularity dissipates, and life is uncertain for our narrator, Charlie Friend. He is an archetypal “mediocre white guy”

who makes a modest living as a day trader and pursues a tepid romance with his upstairs neighbor, Miranda.

Charlie, like many young men of an inquisitive, neophilic temperament, is obsessed with technology, and approaches the grave challenge of artificial intelligence with all the caution of a kid in a candy store. He has just spent—it is too soon to say squandered—his inheritance on a lifelike £86,000 robot named Adam. If McEwan borrowed Miranda's name from *The Tempest* ("O brave new world, That has such people in 't!"), his Adam is not unlike Shakespeare's Caliban, both human and not quite human, first a friend and then a servant, or slave. Charlie initially dislikes thinking of himself as Adam's "user." Later, after Adam has sex with Miranda—further shades of *The Tempest*, in which Caliban makes an attempt on Miranda's virtue—jealous Charlie does try to assert his authority.

Adam is what Turing, in our own historical timeline, called a "child machine." As Turing wrote in his 1950 essay "Computing Machinery and Intelligence" (in which he also describes the "imitation game," now better known as the Turing Test):

Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of

education one would obtain the adult brain....Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child.

Adam accumulates content at breakneck speed. His working memory is superhuman: "Every moment of his existence, everything he heard and saw, he recorded and could retrieve." Improvements to his "mechanism," in this case the development of a distinctive, human-like personality, are made at the outset by his programmers, Charlie and Miranda. "I would fill in roughly half the choices for Adam's personality," Charlie decides, "then give her the link and the password and let her choose the rest."

I wouldn't interfere, I wouldn't even want to know what decisions she had made. She might be influenced by a version of herself: delightful. She might conjure the man of her dreams: instructive. Adam would come into our lives like a real person, with the layered intricacies of his personality revealed only through time, through events, through his dealings with whomever he met. In a sense he would be like our child.

Set aside how creepy this is—not least because Adam does not begin

existence as a baby but as a “handsome dark-skinned young man,” who “was capable of sex and possessed functional mucous membranes.” Why not just *have* a child? The answer lies in Charlie’s conviction that artificial people will become “more than us,” that as they integrate into our world, “tragedy was a possibility, but not boredom.” He is not wrong, but, like that other “modern Prometheus,” Victor Frankenstein, he should have been careful what he wished for. Further complicating things, circumstances place a neglected *human* child named Mark in Charlie and Miranda’s care. This surprise triggers Miranda’s maternal instinct, Adam’s jealousy of Miranda, and Adam’s envy of Mark, who is in certain respects an even more sophisticated “learning computer” than Adam is.

Adam is unpredictable: first when he sleeps with Miranda, again when he rats out Miranda to Charlie for having committed a serious crime in her past, and a third time when he turns Miranda over to the police. (Adam also writes insufferable poetry, McEwan’s way of suggesting that when the robot revolution comes, *his* job, at least, is secure.) Miranda’s crime was an attempt to serve justice where the legal system failed to do so, by framing a guilty man. The book’s central moral question is whether punishing her vigilantism is preferable to letting her live her future to the good, as the adoptive parent of

a flesh-and-blood, non-robotic child. McEwan knows that many humans would err on the side of what they perceive as compassion. He knows that many humans would rather see a wrong repaid with a right than with a talionic punishment. Whether this tendency belongs to man’s goodness, his sentimentality, or his desire to ‘scape whipping, McEwan doesn’t say. Having placed the question in our path, he lets us work out our own complicated feelings about artificial intelligence.

Adam surprises his “users” in an unsurprising way. He is superintelligent but inflexible. He understands the letter of the law but not the spirit, the force of law but not the force of custom. He is loyal to Charlie and Miranda only in the sense that he holds them to a standard he has, quite arbitrarily, taught himself to value. The manner in which he ruins their shared life demonstrates the most severe limitations of artificial—that is, *simulated*—intelligence. A machine can be taught when to bend a rule only by supplying it with more rules, *ad infinitum*. This is not unpredictability, nor is it thought, nor, most importantly, does it result in compassion or love. It is a complex simulation of human personality. A “person” thus constituted may be a technological marvel—at least, as with any technological marvel, until its novelty wears off. Yet a really persuasive simulation is still just a simulation. What makes a diamond is how it

came about, not what it looks like beneath the loupe.

To regard an artificial human as “more than us” is a symptom of a kind of profound spiritual fatigue. There is something unmistakably masochistic in the dream of AI, something akin to self-loathing at the species level. The fear that machines will become self-aware sometimes seems to mask the wish that they will do so—and find us lacking. Influenced by emotion, morality, culture, etiquette, and so on, human beings are messy and unpredictable in ways no machine can properly mimic. So, to modify Samuel Johnson, man makes a bot of himself to get rid of the pain of being a man. When human beings speak of “playing God,” when they put something as banal as AI on par with the appearance of matter, life, and consciousness *ex nihilo*, the aim may not be to elevate man so much as to demystify him and whatever brought him about. To possess true consciousness is the biggest and indeed the only responsibility in the known universe. AI promises a break, as it were, from the colossal burden of being the only show in town.

The many traces of present-day technological development converge in Joanna Kavenna’s maximalist, brilliant, maddening new novel *Zed*. Social media, automation, surveillance, data mining, cryptocurrency, AI, facial recognition, biometric data processing, augmented

and virtual reality, transhumanism, advanced predictive algorithms, even the indignities of predictive text: everything and more is in the cross-hairs here. But the book is, at a level more abstract and discomfiting than McEwan’s, about the quixotic quest to conquer the irrational. In Kavenna’s world, also a parallel Britain, “Zed” is a term of art used at Beetle, a corporation like a combination of Google, Amazon, and Facebook, but with an even more scandalous helping of state power. Zed, an employee explains, “just means the stuff that doesn’t quite fit within every paradigm. Or, the anomalies that prove the system. It’s no big deal. Every system, however immaculate, has a few little glitches. We lump them together under this category term.”

Technology has invaded and colonized every corner of the social and domestic space. There is the VIPA or Veep (Very Intelligent Personal Assistant, like Siri, Alexa, or HAL 9000), the VIADS (Very Intelligent Automated Driving System), and the BeetleBand, something akin to an Apple Watch or Fitbit if it combined the bland lifestyle branding of Gwyneth Paltrow with the tyranny of Carrie Nation, the American Prohibitionist who tore through bars with a hatchet. The world depicted in *Zed* is, indeed, the twenty-first-century version of the society destroyed in Bradbury’s “There Will Come Soft Rains”:

The Custodians Program tracked people from the moment they woke (having registered the quality of their sleep, the duration), through their breakfast (registering what they ate, the quality of their food), through the moment they dressed, and if they showered and cleaned their teeth properly, if they took their DNA toothbrush test, what time they left the house, whether they were cordial to their door, whether they told it to f\*\*\*ing open up and stop talking to them....It was sometimes difficult to determine if BeetleBand readings were good or bad; for example, a high pulse rate could indicate exercise, stress, or passionate sex....For greater accuracy [the Custodians] combined these readings with recorded visuals as well.

Into each life some rain must fall, and even in this exhaustively monitored and quantified society, Zed lurks. Beetle's real business is in the "lifechain," its all-encompassing predictive algorithm, which relies for its authority on never being wrong. One morning a Beetle employee confounds the lifechain and does the unexpected, murdering his wife and two sons and then vanishing. This touches off what techies call cascading failure, as an innocent man mistaken for the killer is "neutralized" by a hulking, headless robocop called an ANT (Anti-Terror Droid), and public confidence in Beetle is jeopardized.

At an inquest, a Beetle employee fumbles through an explanation of the ANT's superiority to a human in reaction time and decision-making. "The actual process," he says, "is too complex to relay, as the ANT can process millions of possible variables every second." His skeptical examiner asks, "Millions? Are there even millions of variables?" Later, when an attempt is made to blame the error on an unforeseeable "perception ellipsis" caused by a ray of light, the same examiner asks, "Are shafts of sunlight really to be described as 'unexpected'? When sunlight is a fundamental property of life on earth?" Certainly a set of millions of variables ought to contain *that* one.

The "Zed events" driving Kavenna's plot proliferate like pop-ups on a fake news site. Veeps begin to go disconcertingly off-script. Human hackers and dissidents attack Beetle and its benign image, while a rival Chinese megacorp patiently chisels away at Beetle's supremacy. Beetle's CEO, Guy Matthias, so disdainful of risk that he uses lifechain analysis to rate romantic encounters before they occur, finds himself buffeted by the sudden unpredictability of those around him—not only well-paid employees and bought-off journalists but also his fed-up ex-wife. Loyalty, conscience, and love do not, apparently, obey any obvious laws.

Zed is astute about the slippery nature of free will, about the ways in which human behavior is amenable



to forecasting and the ways in which it is not. A major expansion of computing power, the ability to consider an exponential profusion of variables, would not increase predictive accuracy but rather drown every question in possible answers—much as real life already does. If programming a convincing AI is in some sense the reverse of predicting an actual human's behavior, then bad news: The AI will become too unruly to *use* for anything long before anyone is foolish enough to wonder if it has become “self-aware.”

One can only hint at the disturbing pleasures of *Zed*; the Dickensian richness and vitality of its characters and prose; the way it becomes fractured and surreal as it asymptotically approaches a conclusion. But it succeeds in part by being fair to technology. It is not the work of a dyed-in-the-wool Luddite. It explores the connection between language and consciousness, ridiculing not only political Newspeak and corporate cant but also the literal-minded utili-talk of simulated minds. Yet it never stoops to the tedious sci-fi motif that pits Shakespeare and the soaring, indomitable human soul against cold, unfeeling science. On the contrary, it depicts the soul, with its fear, weakness, and love, as the very thing that reaches for a refuge from danger and uncertainty. *Zed* simply adds this question as a stinger: If man is saved from contingency, from danger, from struggle, why satisfy

his needs at all? Is he not, in a sense, already dead?

What these books reveal, sometimes explicitly, sometimes unwittingly, is that man cannot stray too far from peril, chaos, and the irrational without feeling cowardly, enervated, undignified, *less than human*. Robert Louis Stevenson made this observation in his essay “The Day After Tomorrow”:

It is certain that man loves to eat, it is not certain that he loves that only or that best. He is supposed to love comfort; it is not a love, at least, that he is faithful to. He is supposed to love happiness; it is my contention that he rather loves excitement. Danger, enterprise, hope, the novel, the aleatory, is dearer to man than regular meals. He does not think so when he is hungry, but he thinks so again as soon as he is fed.

Science, technology, and the free market ensure that man is fed—surfeited, at times—and then the satirists assemble to remind him that what he *really* wants, what he needs, in the depths of his soul, are “the glow of hope, the shock of disappointment, furious contention with obstacles.” He needs to go hungry from time to time.

In this, our satirists perform a quasi-priestly function. They rail against some big red devils—totalitarian governments, out-of-control scientists, greedy corporations—only

to wink between the lines, to remind us that the Devil has always been a symbol and that all of our temptations come from within. These books are an urgent reminder to scan the horizon and ask: Do I want what's coming? Why do I want it? Why is it so important to me to pretend that I didn't ask for it? What might

I be betraying in myself? From what would I derive meaning in the absence of struggle, of mystery, of the unknown?

And, mostly importantly, who's driving this thing, anyway?

*Stefan Beck is a writer living in Hudson, New York.*