ANNALS OF INQUIRY

DO BETTER

The gospel of effective altruism.

BY GIDEON LEWIS-KRAUS

The philosopher William MacAskill credits his personal transfiguration to an undergraduate seminar at Cambridge. Before this shift, MacAskill liked to drink too many pints of beer and frolic about in the nude, climbing pitched roofs by night for the life-affirming flush; he was the saxophonist in a campus funk band that played the May Balls, and was known as a hopeless romantic. But at eighteen, when he was first exposed to "Famine, Affluence, and Morality," a 1972 essay by the radical utilitarian Peter Singer, MacAskill felt a slight click as he was shunted onto a track of rigorous and uncompromising moralism. Singer, prompted by widespread and eradicable hunger in what's now Bangladesh, proposed a simple thought experiment: if you stroll by a child drowning in a shallow pond, presumably you don't worry too much about soiling your clothes before you wade in to help; given the irrelevance of the child's location—in an actual pond nearby or in a metaphorical pond six thousand miles away—devoting resources to superfluous goods is tantamount to allowing a child to drown for the sake of a dry cleaner's bill. For about four decades, Singer's essay was assigned predominantly as a philosophical exercise: his moral theory was so onerous that it had to rest on a shaky foundation, and bright students were instructed to identify the flaws that might absolve us of its demands. MacAskill, however, could find nothing wrong with it.

By the time MacAskill was a graduate student in philosophy, at Oxford, Singer's insight had become the organizing principle of his life. When he met friends at the pub, he ordered only a glass of water, which he then refilled with a can of two-per-cent lager he'd bought on the corner; for dinner, he ate bread he'd baked at home. The balance of his earnings was reserved for others. He tried not to be too showy or evangelical, but neither was he diffident about his rationale. It was a period in his life both darkly lonesome and ethically ablaze. As he put it to me recently, "I was very annoying."

In an effort to shape a new social equilibrium in which his commitments might not be immediately written off as mere affectation, he helped to found a moral crusade called "effective altruism." The movement, known as E.A. to its practitioners, who themselves are known as E.A.s, takes as its premise that people ought to do good in the most clear-sighted, ambitious, and unsentimental way possible. Among other back-of-the-envelope estimates, E.A.s believe that a life in the developing world can be saved for about four thousand dollars. Effective altruists have lashed themselves to the mast of a certain kind of logical rigor, refusing to look away when it leads them to counterintuitive, bewildering, or even seemingly repugnant conclusions. For a time, the movement recommended that inspirited young people should, rather than work for charities, get jobs in finance and donate their income. More recently, E.A.s have turned to fretting about existential risks that might curtail humanity's future, full stop.

Effective altruism, which used to be a loose, Internet-enabled affiliation of the like-minded, is now a broadly influential faction, especially in Silicon Valley, and controls philanthropic resources on the order of thirty billion dollars. Though MacAskill is only one of the movement's principal leaders, his conspicuous integrity and easygoing charisma have made him a natural candidate for head boy. The movement's transitions—from obscurity to power; from the needs of the contemporary global poor to those of our distant descendants—have not been altogether smooth. MacAskill, as the movement's de-facto conscience, has felt increasing pressure to provide instruction and succor. At one point, almost all of his friends were E.A.s, but he now tries to draw a line between public and private. He told me, "There was a point where E.A. affairs were no longer social things—people would come up to me and want to talk about their moral priorities, and I'd be, like, 'Man, it's 10 P.M. and we're at a party!'"

On a Saturday afternoon in Oxford, this past March, MacAskill sent me a text message about an hour before we'd planned to meet: "I presume not, given jetlag, but might you want to go for a sunset swim? It'd be very very cold!" I was out for a run beside the Thames, and replied, in an exacting mode I hoped he'd appreciate—MacAskill has a way of making those around him greedy for his approval—that I was about eight-tenths of a mile from his house, and would be at his door in approximately five minutes and thirty seconds. "Oh wow impressive!" he replied. "Let's do it!"

MacAskill limits his personal budget to about twenty-six thousand pounds a year, and gives everything else away. He lives with two roommates in a stolid row house in an area of south Oxford bereft, he warned me, of even a good coffee shop. He greeted me at his door, praising my "bias for action," then led me down a low and dark hallway and through a laundry room arrayed with buckets that catch a perpetual bathroom leak upstairs. MacAskill is tall and sturdily built, with an untidy mop of dark-blond hair that had grown during the pandemic to messianic lengths. In an effort to unwild himself for reëntry, he had recently reduced it to a dimension better suited to polite society.

MacAskill allowed, somewhat sheepishly, that lockdown had been a welcome reprieve from the strictures of his previous life. He and some friends had rented a home in the Buckinghamshire countryside; he'd meditated, acted as the house exercise coach, and taken in the sunset. He had spent his time in a wolf-emblazoned jumper writing a book called "What We Owe the Future,"



"The world's long-run fate depends in part on the choices we make in our lifetimes," the philosopher William MacAskill writes.

PHOTOGRAPH BY ULYSSES ORTEGA

which comes out this month. Now the world was opening up, and he was being called back to serve as the movement's shepherd. He spoke as if the life he was poised to return to were not quite his own—as if he weren't a person with desires but a tabulating machine through which the profusion of dire global need was assessed, ranked, and processed.

He was doing his best to retain a grasp on spontaneity, and we set off on the short walk to the lake. Upon our arrival, MacAskill vaulted over a locked gate that led to a small floating dock, where he placed a Bluetooth speaker that played a down-tempo house remix of the 1974 pop hit "Magic." The water temperature, according to a bath-toy thermometer, was about fifty degrees. He put on a pair of orange sunglasses with tinted lenses, which enhanced the sunset's glow, and stripped off his shirt, revealing a long abdominal scar, the result of a fall through a skylight as a teen-ager. He reassured me, "If all you do is just get in and get out, that's great." I quickly discharged my duty and then flung myself, fingers blue, back onto the dock. MacAskill did a powerful breaststroke out into the middle of the lake, where he floated, freezing, alone and near-invisible in the polarized Creamsicle sunset. Then he slowly swam back to resume his obligations.

MacAskill, who was born in 1987 as William Crouch, grew up in Glasgow and attended a vaunted private school. He excelled at almost everything but was the first to make fun of himself for singing off-key, juggling poorly, and falling out of treehouses. Though his mother grew up in conditions of rural Welsh privation, his family had little political color—as a child, he was given to understand that all newspapers were right-leaning tabloids. From an early age, however, he demonstrated a precocious moral zeal. At fifteen, when he learned how many people were dying of AIDS, he set out to become a successful novelist and give away half of his earnings. He volunteered for a disabled-Scout group and worked at a care home for the elderly, which his parents found baffling. In his milieu, the brightest graduates were expected to study medicine in Edinburgh, but MacAskill, as class dux, or valedictorian, won a place to read philosophy at Cambridge. Robbie Kerr, MacAskill's closest schoolmate, told me, "The Glasgow attitude was best summed up by a school friend's parent, who looked at Will and said, 'Philosophy. What a waste. That boy could have cured cancer.'"

MacAskill found Cambridge intellectually and socially satisfying: he discussed meta-ethics on shirtless walks, and spent vacations at friends' homes in the South of France. But he also remembers feeling adrift, "searching for meaning." "There weren't a lot of opportunities for moral activism," he told me. He spent a summer volunteering at a rehabilitation center in Ethiopia and, after graduation, another as a "chugger," a street canvasser paid to convert pedestrians to charitable causes. "We used to say it only cost twenty pence to save a life from polio, and a lot of other stuff that was just wrong," he said, shaking his head. Nevertheless, he continued, "it was two months of just sitting with extreme poverty, and I felt like other people just didn't get it." In graduate school, "I started giving three per cent, and then five per cent, of my income," he said. This wasn't much—he was then living on a university stipend. "I think it's O.K. to tell you this: I supplemented my income with nude modelling for life-drawing classes." The postures left him free to philosophize. Later, he moved on to bachelorette parties, where he could make twice the money "for way easier poses."

He told me, "I was in the game for being convinced of a cause, and did a bunch of stuff that was more characteristically far-lefty. I went to a climate-justice protest, and a pro-Palestinian protest, and a meeting of the Socialist Workers Party." None passed muster, for reasons of efficacy or intellectual coherence. "I realized the climate protest was *against* cap-and-trade, which I was *for*. The Socialist Workers Party was just eight people with long hair in a basement talking about the glory of the Russian Revolution." He surveyed working philosophers and found that none felt like they'd done anything of real consequence. George Marshall, a friend from Cambridge, told me, "He was at dinner in Oxford—some sort of practical-ethics conference—and he was just deeply shocked that almost none of the attendees were vegetarians, because he thought that was the most basic application of ethical ideas."

When MacAskill was twenty-two, his adviser suggested that he meet an Australian philosopher named Toby Ord. In activist circles, MacAskill had found, "there was this focus on the problems—climate is so bad!—along with intense feelings of angst, and a lack of real views on what one could actually do. But Toby was planning to give money in relatively large amounts to focussed places, and trying to get others to do the same—I felt, 'Oh, this is taking action.'" At the time, Ord was earning fifteen thousand pounds a year and was prepared to give away a quarter of it. "He'd only had two half-pints in his time at Oxford," MacAskill said. "It was really hardcore." Unlike, say, someone who donates to cystic-fibrosis research because a friend suffers from the disease—to take a personal example of my own—Ord thought it was important that he make his allocations impartially. There was no point in giving to anyone in the developed world; the difference you could make elsewhere was at least two orders of magnitude greater. Ord's ideal beneficiary was the Fred Hollows Foundation, which treats blindness in poor countries for as little as twenty-five dollars a person.

MacAskill immediately signed on to give away as much as he could in perpetuity: "I was on board with the idea of binding my future self—I had a lot of youthful energy, and I was worried I'd become more conservative over time." He recalled the pleasure of proving that his new mentor's donations were suboptimal. "My first big win was convincing him about deworming charities." It may seem impossible to compare the eradication of blindness with the elimination of intestinal parasites, but health economists had developed rough methods. MacAskill estimated that the relief of intestinal parasites, when measured in "quality-adjusted life years," or QALYs, would be a hundred times more cost-effective than a sight-saving eye operation. Ord reallocated.

If Peter Singer's theory—that any expenditure beyond basic survival was akin to letting someone die—was simply too taxing to gain wide adherence, it seemed modest to ask people to give ten per cent of their income. This number also had a long-standing religious precedent. During the next six months, MacAskill and Ord enjoined their friends and other moral philosophers to pledge a secular

tithe. MacAskill told me, "I would quote them back to themselves—you know, 'If someone in extreme poverty dies, it's as if you killed them yourself,' and other really severe pronouncements—and say, 'So, would you like to sign?'" Singer said yes, but almost everyone else said no. On November 14, 2009, in a small room in Balliol College, MacAskill and Ord announced Giving What We Can. MacAskill said, "At the launch, we had twenty-three members, and most of them were friends of Toby's and mine."

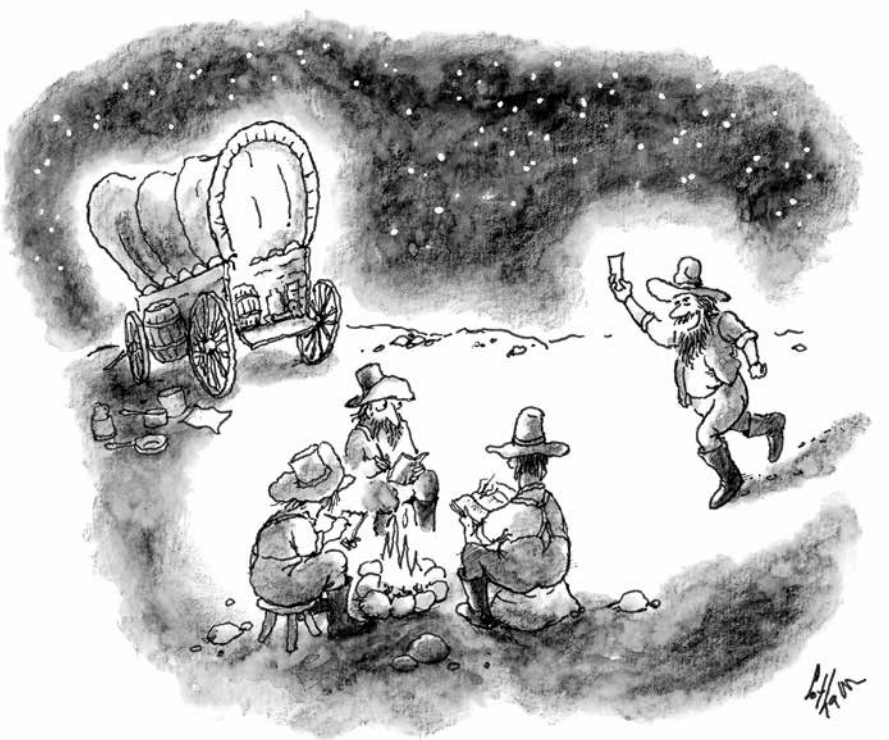
When MacAskill took his vow of relative poverty, he worried that it would make him less attractive to date: "It was all so weird and unusual that I thought, Out of all the people I could be in a relationship with, I've just cut out ninety-nine per cent of them." This prediction was incorrect; in 2013, he married another Scottish philosopher and early E.A., and the two of them took her grandmother's surname, MacAskill. Later, a close relative found out what MacAskill had been doing with his stipend and told him, "That's unethical!" If he wasn't using his scholarship, he should return it to the university. He loves his family, he told me, "but I guess if I'd spent that money on beer it would have been O.K."

Like agriculture, echolocation, and the river dolphin, the practice that would become effective altruism emerged independently in different places at around the same time. Insofar as there was a common ancestor, it was Peter Singer. Holden Karnofsky and Elie Hassenfeld, young analysts at the hedge fund Bridgewater Associates, formed a club to identify the most fruitful giving opportunities—one that relied not on crude heuristics but on hard data. That club grew into an organization called GiveWell, which determined that, for example, the most cost-effective way to save a human life was to give approximately four thousand dollars to the Against Malaria Foundation, which distributes insecticide-treated bed nets. In the Bay Area "rationalist" community, a tech-adjacent online subculture devoted to hawkish logic and quarrelsome empiricism, bloggers converged on similar ideas. Eliezer Yudkowsky, one of the group's patriarchs, instructed his followers to "purchase fuzzies and utilons separately." It was fine to tutor at-risk kids or volunteer in a soup kitchen, as long as you assigned those activities to a column marked "self-interest." But the pursuit of a warm glow should be separate from doing the most impartial good.

When I asked Singer why the late two-thousands were a time of great ferment for applied consequentialism, he cited the Internet: "People will say, 'I've had these ideas since I was a teen-ager, and I thought it was just me,' and then they got online and found that there were others." Julia Wise, then an aspiring social worker, had been giving for years to the point of extreme personal sacrifice; she met Ord in the comments section of the economist Tyler Cowen's blog, and made the Giving What We Can pledge. She told me that she was attracted on a "tribal basis" to the movement's sense of "global solidarity."

Proto-E.A. attracted people who longed to reconcile expansive moral sensibilities with an analytical cast of mind. They tended to be consequentialists—those who believe that an act should be evaluated not as it conforms to universal rules but based on its results—and to embrace utilitarianism, a commitment to the greatest good for the greatest number. Their instinct was to see moral interventions as grand optimization problems, and they approached causes on the basis of three criteria: importance, tractability, and neglectedness. They were interested in thought experiments like the trolley problem, in part because they found such exercises enlivening and in part because it emphasized that passive actors could be culpable. It also made plain the very real dilemma of resource constraints: if the same amount of money could save one person here or five people there, there was no need for performative hand-wringing. As the rationalists put it, sometimes you just had to "shut up and multiply." A kind of no-hard-feelings, debate-me gladiatorialism was seen as a crucial part of good "epistemic hygiene," and a common social overture was to tell someone that her numbers were wrong. On GiveWell's blog, MacAskill and Karnofsky got into a scrap about the right numbers to assign to a deworming initiative. Before Wise met MacAskill, she had e-mailed to say that he had got some other numbers wrong by an order of magnitude. "A few months later, here I was having a beer with Will," she said.

In late 2011, in the midst of the Occupy movement, MacAskill gave a talk at Oxford called "Doctor, NGO Worker, or Something Else Entirely? Which Careers Do the Most Good." That year



"Any of you boys interested in fresh gossip for your diaries?"

saw the launch of 80,000 Hours, an offshoot of Giving What We Can designed to offer "ethical life-optimisation" advice to undergraduates. His advice, which became known as "earning to give," was that you—and the "you" was pretty explicitly high-calibre students at élite institutions—could become a doctor in a poor country and possibly save the equivalent of a hundred and forty lives in your medical career, or you could take a job in finance or consulting and, by donating intelligently, save ten times as many.

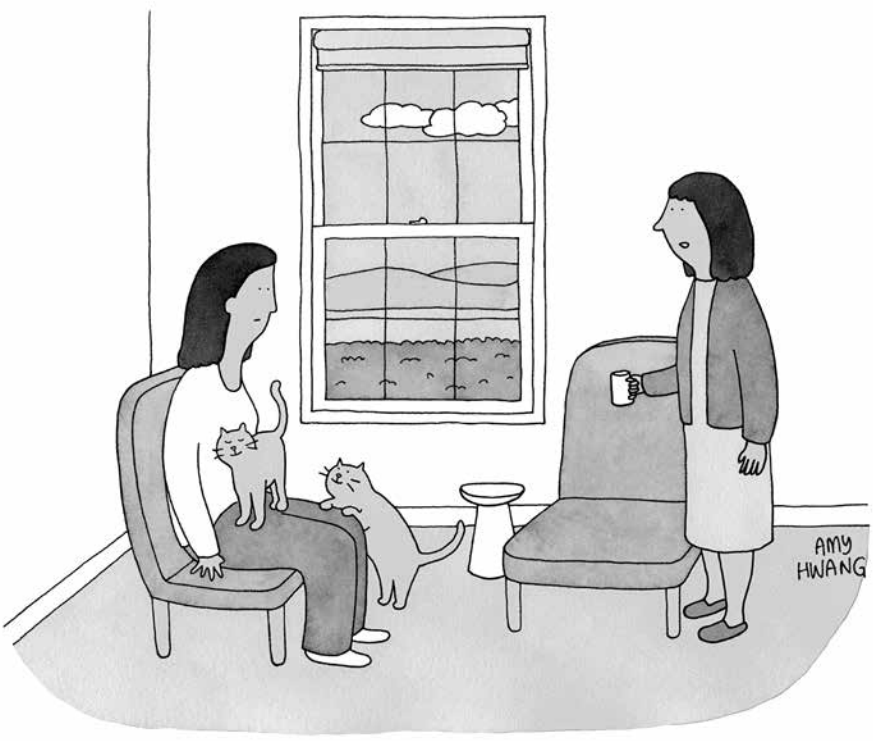
A young Oxonian named Habiba Islam was at that talk, and it changed her life. "I was the head of Amnesty International at university, I was volunteering at the local homeless shelter in Oxford—that kind of thing," she told me. "I know people who were committed to climate change as their thing—a pretty good guess for what's important—before getting involved in E.A." Islam was considering a political career; 80,000 Hours estimated that an Oxford graduate in Philosophy, Politics, and Economics who embarked on such a path has historically had about a one-in-thirty chance of becoming an M.P. She took the pledge, agreeing to give away everything above twenty-five thousand pounds a year, and became a consultant for PwC. She told me, "It was just obvious that we privileged people should be helping more than we were."

Matt Wage, a student of Singer's at Princeton, decided that, instead of pursuing philosophy in grad school, he would get a job at the trading firm Jane Street. If you rescued a dozen people from a burning building, he thought, you would live out the rest of your days feeling like a hero; with donations, you could save that many lives every year. "You can pay to provide and train a guide dog for a blind American, which costs about $40,000," Wage told the reporter Dylan Matthews, for a Washington *Post* piece called "Join Wall Street. Save the world." "But with that money you could also cure between 400 and 2,000 people in developing countries of blindness from glaucoma." Matthews, convinced by his sources in the movement, went on to donate a kidney. "You go to an E.A. conference and things feel genuinely novel," he said. "Someone will give a talk about how, if we regulate pesticides differently, we can reduce suicide. I don't have to agree with everything they say, or think Will is the Pope, for it to be a very useful way to think about what deserves attention."

In the movement's early years, MacAskill said, "every new pledge was a big deal, a cause for celebration." As E.A. expanded, it required an umbrella nonprofit with paid staff. They brainstormed names with variants of the words "good" and "maximization," and settled on the Centre for Effective Altruism. Wise donated thousands of dollars; it was the first time her money was going not to "object-level" work but to movement-building. She said, "I was an unemployed social worker, but I felt so optimistic about their work that I gave them most of my savings."

In 2015, MacAskill was hired as an associate professor at Oxford and, at twenty-eight, was said to be the youngest such philosophy professor in the world. It should have been a moment of vindication, but he felt conflicted. "It was easy and high-status," he said. "But I didn't want to be too comfortable." The same year, his marriage deteriorated—he kept the surname; his ex didn't—and he published his first book, "Doing Good Better," an extended case that Westerners were in a situation akin to "a happy hour where you could either buy yourself a beer for five dollars or buy someone else a beer for five cents." That summer, 80,000 Hours was accepted into Y Combinator, a prestigious startup incubator. The Effective Altruism Global summit was held at the Googleplex, and Elon Musk appeared on a panel about artificial intelligence. (MacAskill told me, "I tried to talk to him for five minutes about global poverty and got little interest.") By then, GiveWell had moved to San Francisco, and the Facebook cofounder Dustin Moskovitz and his wife, the former journalist Cari Tuna, had tasked its new project—later known as Open Philanthropy—with spending down their multibillion-dollar fortune. Open Philanthropy invested in international development and campaigns for broiler-chicken welfare, and expanded into causes like bail reform. For the first time, MacAskill said, the fledgling experiment felt like "a force in the world."

MacAskill, too, was newly a force in the world. For all E.A.'s aspirations to stringency, its numbers can sometimes seem arbitrarily plastic. MacAskill has a gap between his front teeth, and he told close friends that he was now thinking of getting braces, because studies showed that more "classically" handsome people



"They like you because you're allergic to them."

were more impactful fund-raisers. A friend of his told me, "We were, like, 'Dude, if you want to have the gap closed, it's O.K.' It felt like he had subsumed his own humanity to become a vehicle for the saving of humanity."

The Centre for Effective Altruism now dwells, along with a spate of adjacent organizations with vaguely imperious names—the Global Priorities Institute, the Forethought Foundation—in Trajan House, an Oxford building that overlooks a graveyard. Nick Bostrom, a philosopher whose organization, the Future of Humanity Institute, also shares the space, disliked the building's name, which honors a philanthropic Roman emperor, and proposed that it be called Corpsewatch Manor. The interior, with shelves of vegan nutritional bars and meal-replacement smoothies, resembles that of a pre-fiasco WeWork. MacAskill struggles with the opportunity cost of his time. He told me, "It's always been an open question: What weight do I give my own well-being against the possible impact I can have?" Many evenings, he has a frozen vegan dinner at the office. (He wasn't sure, when feeding me, whether the microwave time for two dishes would scale linearly.) He schedules his days down to the half hour. He and one of his assistants recently discussed reducing some slots to twenty-five minutes, but they ultimately decided that it might seem insulting.

In a rare free moment, MacAskill, who wears tight-fitting V necks that accentuate his lack of sartorial vanity and his biceps, took me on a tour of the movement's early sites. We passed Queen's Lane Coffee House. "That's where Bentham discovered utilitarianism," he commented. "There should be a plaque, but the current owners have no idea." Most of the colleges were closed to visitors, but MacAskill had perfected the flash of an old I.D. card and a brazen stride past a porter. He paused reverently outside All Souls College, where the late moral philosopher Derek Parfit, one of the guiding lights of E.A., spent his life in a tower.

Parfit believed that our inherited moral theories were constructed on religious foundations, and aspired to build a comprehensive secular moral framework. Effective altruism, in that spirit, furnishes an all-encompassing world view. It can have an ecclesiastical flavor, and early critics observed that the movement seemed to be in the business of selling philanthropic indulgences for the original sin of privilege. It has a priestly class, whose posts on E.A.'s online forum are often received as encyclicals. In the place of Mass, E.A.s endure three-hour podcasts. There is an emphasis on humility, and a commandment to sacrifice for the sake of the neediest. Since its inception, GiveWell has directed the donation of more than a billion dollars; the Against Malaria Foundation alone estimates that its work to date will save a hundred and sixty-five thousand lives. There have been more than seven thousand Giving What We Can pledges, which total almost three billion dollars. In an alternate world, a portion of that sum would presumably have been spent on overpriced tapas in San Francisco's Mission District.



As effective altruism became a global phenomenon, what had been treated as a fringe curiosity became subject to more sustained criticism. A panel convened by the *Boston Review* described E.A.s as having cast their lot with the status quo. Though their patronage might help to alleviate some suffering on the margins, they left the international machine intact. As hard-nosed utilitarians, they bracketed values—like justice, fairness, and equality—that didn't lend themselves to spreadsheets. The Stanford political scientist Rob Reich wrote, "Plato identified the best city as that in which philosophers were the rulers. Effective altruists see the best state of affairs, I think, as that in which good-maximizing technocrats are in charge. Perhaps it is possible to call this a politics: technocracy. But this politics is suspicious of, or rejects, the form of politics to which most people attach enormous value: democracy." The Ethiopian American A.I. scientist Timnit Gebru has condemned E.A.s for acting as though they are above such structural issues as racism and colonialism.

Few of these appraisals were new; many were indebted to the philosopher Bernard Williams, who noted that utilitarianism might, in certain historical moments, look like "the only coherent alternative to a dilapidated set of values," but that it was ultimately bloodless and simpleminded. Williams held that the philosophy alienated a person "from the source of his actions in his own convictions"—from what we think of as moral integrity. Its means-end rationality could seem untrustworthy. Someone who seeks justification for the impulse to save the life of a spouse instead of that of a stranger, Williams famously wrote, has had "one thought too many."

The Oxford philosopher Amia Srinivasan, whom MacAskill considers a friend, wrote a decidedly mixed critique in the *London Review of Books*, calling MacAskill's first book "a feel-good guide to getting good done." She noted, "His patter is calculated for maximal effect: if the book weren't so cheery, MacAskill couldn't expect to inspire as much do-gooding." She conceded the basic power of the movement's rhetoric: "I'm not saying it doesn't work. Halfway through reading the book I set up a regular donation to GiveDirectly," one of GiveWell's top recommended charities. But she called upon effective altruism to abandon the world view of the "benevolent capitalist" and, just as Engels worked in a mill to support Marx, to live up to its more thoroughgoing possibilities. "Effective altruism has so far been a rather homogenous movement of middle-class white men fighting poverty through largely conventional means, but it is at least in theory a broad church." She noted, encouragingly, that one element was now pushing for "systemic change" on issues like factory farming and immigration reform.

Some E.A.s felt that one of the best features of their movement—that, in the context of near-total political sclerosis, they had found a way to do *something*—had been recast as a bug. The movement's self-corrections, they believed, had been underplayed: a high-paying job at a petrochemical firm, for example, was by then considered sufficiently detrimental that no level of income could justify it. But others found Srinivasan's criticisms harsh but fair. As Alexander Berger, the co-C.E.O. of Open Philanthropy, told me, "She was basically right

that early E.A. argued for the atomized response—that you as an individual should rationally and calculatedly allocate a portion of your privilege to achieve the best outcomes in the world, and this doesn't leave much space for solidarity." During the next few years, however, the movement gained a new appreciation for the more sweeping possibilities of systemic change—though perhaps not in the ways Srinivasan had envisioned.

One of the virtues of effective altruists—which runs counter to their stereotype as mere actuaries—is that, when they feel like it, they're capable of great feats of imagination. A subset of them, for example, has developed grave concern about the suffering of wild animals: Should we euthanize geriatric elephants? Neutralize predator species? What should be done about the bugs? The prime status marker in a movement that has abjured financial reward is a reputation for punctilious (and often contrarian) intelligence. The community has a tendency to overindex on perceived cerebral firepower, which makes even leading lights like MacAskill feel a perennial sense of imposture. This means that genuinely bizarre ideas, if argued with sufficient virtuosity, get a fair hearing. Holden Karnofsky told me, "If you read things that E.A.s are saying, they sound a lot crazier than what they're actually doing." But the movement—constrained by methodological commitments rather than by substantive moral ones—has proved vulnerable to rapid changes in its priorities from unexpected quarters.

In retrospect, "Doing Good Better" was less a blueprint for future campaigns than an epitaph for what came to be called the "bed-nets era." During the next five years, a much vaster idea began to take hold in the minds of the movement's leaders: the threat of humanity's annihilation. Such concerns had been around since the dawn of the nuclear age. Parfit had connected them to an old utilitarian argument, that the protection of future lives was just as important as the preservation of current ones. The philosopher Nick Bostrom contended that, if humanity successfully colonized the planets within its "light cone"—the plausibly reachable regions of the universe—and harnessed the computational power of the stars to run servers upon which the lives of digital consciousnesses might be staged, this could result in the efflorescence of approximately ten to the power of fifty-eight beings. For any decision we made now, an astronomical number of lives hung in the balance.

In the first month of the pandemic, Toby Ord published a book called "The Precipice." According to Ord's "credences," the chances of human extinction during the next century stand at about 1–6, or the odds of Russian roulette. The major contributor to existential risk was not climate change—which, even in a worst-case scenario, is unlikely to render the planet wholly uninhabitable. (New Zealand, for example, might be fine.) Instead, he singles out engineered pathogens and runaway artificial intelligence. DNA editing might allow a scientist to create a superbug that could wipe us out. A well-intentioned A.I. might, as in one of Bostrom's famous thought experiments, turn a directive to make paper clips into an effort to do so with all available atoms. Ord imagines a power-hungry superintelligence distributing thousands of copies of itself around the world, using this botnet to win financial resources, and gaining dominion "by manipulating the leaders of major world powers (blackmail, or the promise of future power); or by having the humans under its control use weapons of mass destruction to cripple the rest of humanity." These risks might have a probability close to zero, but a negligible possibility times a catastrophic outcome is still very bad; significant action is now of paramount concern. "We can state with confidence that humanity spends more on ice cream every year than on ensuring that the technologies we develop do not destroy us," Ord writes. These ideas were grouped together under the new heading of "longtermism."

When Ord first mentioned existential risk, MacAskill thought that it was a totally crackpot idea. He was uneasy about how it related to his own priorities, and remembers attending a meeting about A.I. risk and feeling frustrated by the vagueness of the potential impacts. But profound improvements in the past half decade (DeepMind's AlphaGo, OpenAI's GPT-3), combined with arguments about the exponential gains in computational power compared to biological benchmarks, cited by Ajeya Cotra, of Open Philanthropy, brought him around. Ord believed that if we made it through the next century or two we would have about even odds of achieving the best possible long-haul future—a universe filled with the descendants of humanity, living lives of untold, unimaginable, and unspecified freedom and pleasure. MacAskill worries in his new book that annihilation per se might not be the only risk. He believes in the radical contingency of moral progress; he argues, for example, that, without the agitation of a small cohort of abolitionists, slavery might have lasted much longer. Even a benign A.I. overlord, by contrast, might produce "value lock-in": a world governed by code that forever stalls the arc of moral progress. (On the other hand, if we don't avail ourselves of the possibilities of A.I., we might face technological stagnation—in over our heads on a deteriorating planet.)

MacAskill understands that worries about a sci-fi apocalypse might sound glib when "there are real problems in the world facing real people," he writes. The distant future, however, is likely to be even more crowded with real people. And if spatial distance is irrelevant to our regard for starvation overseas, temporal distance should be an equally poor excuse. "I now believe the world's long-run fate depends in part on the choices we make in our lifetimes," he writes. This amounts to nothing less than a "moral revolution."

In 2012, while MacAskill was in Cambridge, Massachusetts, delivering his earning-to-give spiel, he heard of a promising M.I.T. undergraduate named Sam Bankman-Fried and invited him to lunch. Bankman-Fried's parents are scholars at Stanford Law School, and he had been raised as a card-carrying consequentialist. He had recently become vegan and was in the market for a righteous path. MacAskill pitched him on earning to give. Bankman-Fried approached an animal-welfare group and asked its members whether they had more use for his volunteer time or for his money, and they strongly preferred the money. The next year, Bankman-Fried invited MacAskill to stay at his coed nerd frat, where everyone slept in the attic to preserve the living area for video and board games.

In 2014, Bankman-Fried graduated

with a degree in physics, and went to work at Jane Street. He says that he donated about half his salary, giving some to animal-welfare organizations and the rest to E.A. movement-building initiatives. In 2017, he started Alameda Research, a crypto-trading firm that sought to exploit an arbitrage opportunity wherein bitcoin, for various reasons, traded higher on Japanese exchanges. The scheme was elaborate, and required that his employees spend a lot of time in bank branches, but he made a ten-per-cent profit on every trade. One crypto impresario told me, "We all knew that was possible in theory, but S.B.F. was the one who actually went and did it."

In 2019, Bankman-Fried founded a user-friendly crypto exchange called FTX. One of the exchange's most profitable products is not yet legal in the United States; he shopped for more congenial jurisdictions and set up in the Bahamas. By the time Bankman-Fried was twenty-nine, *Forbes* estimated his net worth at about twenty-six billion dollars, making him the twenty-fifth-richest American. At least three of his co-workers, depending on the fluctuating price of crypto assets, are also E.A. billionaires. Nishad Singh had been working an earning-to-give job at Facebook when Bankman-Fried invited him to join. Singh told me, "I had been somewhat dishonest with myself. I might have been picking a path that let me lead the life I wanted to lead, but I was not picking the path of maximal good."

Bankman-Fried has refined the persona of a dishevelled, savantlike techno-fakir. He has been widely advertised for his fiscal chastity—he drives a Toyota Corolla and, on the rare occasion that he leaves the office, lives with nine roommates. Even when beds are ready to hand, he pitches down on a beanbag. According to the *Times*, visitors are sometimes scheduled to arrive for meetings during his naps; they watch from a conference room as he wakes up and pads over in cargo shorts. But his marketing efforts have been splashy. FTX spent an estimated twenty million dollars on an ad campaign featuring Tom Brady and Gisele Bündchen, and bought the naming rights to the Miami Heat's arena for a hundred and thirty-five million dollars.

Last year, MacAskill contacted Bankman-Fried to check in about his promise: "Someone gets very rich and, it's, like, O.K., remember the altruism side? I called him and said, 'So, still planning to donate?'" Bankman-Fried pledged to give nearly all his money away; if suitable opportunities are found, he's willing to contribute more than a billion dollars a year. Bankman-Fried had longtermist views before they held sway over MacAskill, and has always been, MacAskill remembers, "particularly excited by pandemics"—a normal thing to hear among E.A.s. Bankman-Fried set up a foundation, the FTX Future Fund, and hired the longtermist philosopher Nick Beckstead as C.E.O. This past December, MacAskill finished the manuscript of his new book, and hoped to spend more time with his partner, Holly Morgan, an early E.A. and the biggest single input to his stability. Instead, Bankman-Fried enlisted him as a Future Fund adviser. (He offered MacAskill a "generous" six-figure salary, but MacAskill replied that he was just going to redistribute the money anyway.)

Overnight, the funds potentially available to E.A. organizations more than doubled, and MacAskill was in a position not only to theorize but to disburse on a grand scale. The Future Fund's initial ideas included the development of early-detection systems for unknown pathogens, and vast improvements in personal protective equipment—including a suit "designed to allow severely immunocompromised people to lead relatively normal lives." With the organization's support, someone might buy a large coal mine to keep the coal in the ground—not only to reduce our carbon footprint but to insure that humanity has available deposits should some desperate future generation have to reindustrialize. The foundation was keen to hear proposals for "civilizational recovery drills," and to fund organizations like ALLFED, which develops food sources that could, in a nuclear winter, be cultivated without sunlight. (So far, it's mostly mushrooms, but seaweed shows promise.) Inevitably, there were calls for bunkers where, at any given time, a subset of humanity would live in a sealed ark.

Along with the money came glamorous attractors. Last week, Elon Musk tweeted, of MacAskill's new book, "This is a close match for my philosophy." (For



"You can't just put on the uniform whenever you don't want to have a conversation, Barry."

• •

a brief period, Musk reportedly assigned responsibility for the charitable distribution of nearly six billion dollars to Igor Kurganov, a former professional poker player and a onetime housemate of MacAskill's; in MacAskill's book, Kurganov is thanked for "unfettered prances round the garden.") MacAskill has long been friendly with the actor Joseph Gordon-Levitt, who told me, "Last year, Will called me up about 'What We Owe the Future,' to talk about what it might be like to adapt the book for the screen." MacAskill felt that such movies as "Deep Impact" and "Armageddon" had prompted governments to take the asteroid threat more seriously, and that "The Terminator"'s Skynet wasn't a bad way to discuss the menace of A.I. Gordon-Levitt said, "We've started figuring out how it could work to build a pipeline from the E.A. community to my creative one, and seeing if we can't get some of these ideas out there into the world."

Bankman-Fried has made an all-in commitment to longtermism. In May, I spoke with him over video chat, and he seemed almost willfully distracted: he didn't bother to hide the fact that he was doing things on several monitors at once. (As a child, his brother has said, Bankman-Fried was so bored by the pace of regular board games that it became his custom to play multiple games at once, ideally with speed timers.) He told me that he never had a bed-nets phase, and considered neartermist causes—global health and poverty—to be more emotionally driven. He was happy for some money to continue to flow to those priorities, but they were not his own. "The majority of donations should go to places with a longtermist mind-set," he said, although he added that some intercessions coded as short term have important long-term implications. He paused to pay attention for a moment. "I want to be careful about being too dictatorial about it, or too prescriptive about how other people should feel. But I did feel like the longtermist argument was very compelling. I couldn't refute it. It was clearly the right thing."

The shift to longtermism, and the movement's new proximity to wealth and power—developments that were not uncorrelated—generated internal discord. In December, Carla Zoe Cremer and Luke Kemp published a paper called "Democratising Risk," which criticized the "techno-utopian approach" of longtermists. Some E.A.s, Cremer wrote in a forum post, had attempted to thwart the paper's publication: "These individuals—often senior scholars within the field—told us in private that they were concerned that any critique of central figures in EA would result in an inability to secure funding." MacAskill responded solicitously in the comments, and when they finally had a chance to meet, in February, Cremer presented a list of proposed "structural reforms" to E.A., including whistle-blower protections and a broad democratization of E.A.'s structure. Cremer felt that MacAskill, the movement leader who gave her the most hope, had listened perfunctorily and done nothing. "I can't wear the E.A. hoodie to the gym anymore," she told me. "Many young people identify with E.A. as a movement or a community or even a family—but underneath this is a set of institutions that are becoming increasingly powerful."

Last year, the Centre for Effective Altruism bought Wytham Abbey, a palatial estate near Oxford, built in 1480. Money, which no longer seemed an object, was increasingly being reinvested in the community itself. The math could work out: it was a canny investment to spend thousands of dollars to recruit the next Sam Bankman-Fried. But the logic of the exponential downstream had some kinship with a multilevel-marketing ploy. Similarly, if you assigned an arbitrarily high value to an E.A.'s hourly output, it was easy to justify luxuries such as laundry services for undergraduate groups, or, as one person put it to me, wincing, "retreats to teach people how to run retreats." Josh Morrison, a kidney donor and the founder of a pandemic-response organization, commented on the forum, "The Ponzi-ishness of the whole thing doesn't quite sit well."

One disaffected E.A. worried that the "outside view" might be neglected in a community that felt increasingly insular. "I know E.A.s who no longer seek out the opinions or input of their colleagues at work, because they take themselves to have a higher I.Q.," she said. "The common criticism thrown at the Tory Party here is that they go straight from Oxford to a job in Parliament. How could they possibly solve problems that they themselves have never come into contact with? They've never been at the coalface. The same criticism could be said of many E.A.s." The community's priorities were prone to capture by its funders. Cremer said, of Bankman-Fried, "Now everyone is in the Bahamas, and now all of a sudden we have to listen to three-hour podcasts with him, because he's the one with all the money. He's good at crypto so he must be good at public policy . . . what?!"

The bed-nets era had been chided as myopic, but at least its outcomes were concrete. The same could not be said of longtermism. Among the better objections was a charge of "cluelessness," or the recognition that we have trouble projecting decades down the line, let alone millennia. It does, in any case, seem convenient that a group of moral philosophers and computer scientists happened to conclude that the people most likely to safeguard humanity's future are moral philosophers and computer scientists. The movement had prided itself on its resolute secularism, but longtermist dread recalled the verse in the Book of Revelation that warns of a time when the stars will fall from the sky like unripe figs. Rob Reich, the Stanford political scientist, who once sat on the board of GiveWell, told me, "They are the secular apocalypticists of our age, not much different than Savonarola—the world is ending and we need a radical break with our previous practices." Longtermism is invariably a phenomenon of its time: in the nineteen-seventies, sophisticated fans of "Soylent Green" feared a population explosion; in the era of "The Matrix," people are prone to agonize about A.I. In the week I spent in Oxford, I heard almost nothing about the month-old war in Ukraine. I could see how comforting it was, when everything seemed so awful, to take refuge on the higher plane of millenarianism.

Longtermism also led to some bizarre conclusions. Depending on the probabilities one attaches to this or that outcome, something like a .0001-per-cent reduction in over-all existential risk might be worth more than the effort to save a billion people today. (In the

literature, this argument is called "fanaticism," and, though it remains a subject of lively scholastic debate in E.A. circles, nobody openly endorses it.) Referring to such regrettable episodes as all previous epidemics and wars, Nick Bostrom once wrote, "Tragic as such events are to the people immediately affected, in the big picture of things—from the perspective of humankind as a whole—even the worst of these catastrophes are mere ripples on the surface of the sea of life." Nick Beckstead, the philosopher at the helm of the Future Fund, remarked in his 2013 dissertation, "Richer countries have substantially more innovation, and their workers are much more economically productive. By ordinary standards—at least by ordinary enlightened humanitarian standards—saving and improving lives in rich countries is about equally as important as saving and improving lives in poor countries, provided lives are improved by roughly comparable amounts. But it now seems more plausible to me that saving a life in a rich country is substantially more important than saving a life in a poor country."

Beckstead's comment may formalize what many philanthropists already do: the venture capitalist John Doerr recently gave a billion dollars to already over-endowed Stanford to bankroll a school for studying climate change. But such extreme trade-offs were not an easy sell. As Holden Karnofsky once put it, most people who sit down to reason through these things from a place of compassion don't expect to arrive at such conclusions—or want to. E.A. lifers told me that they had been unable to bring themselves to feel as though existential risk from out-of-control A.I. presented the same kind of "gut punch" as global poverty, but that they were generally ready to defer to the smart people who thought otherwise. Nishad Singh told me that he, like many longtermists, continues to donate to alleviate current misfortune: "I still do the neartermist thing, personally, to keep the fire in my belly."

One of the ironies of the longtermist correction was that, all of a sudden, politics was on the table in a new way. In 2020, Bankman-Fried donated more than five million dollars to Joe Biden's campaign, making him one of the top Democratic contributors. Given anxieties about the nuclear codes, his action wasn't hard to justify. But Bankman-Fried has his own interests—the only times he's been known to wear pants is in front of Congress, where he urges crypto deregulation—and electoral interventions are slippery. This year, Bankman-Fried's super PAC gave more than ten million dollars to support Carrick Flynn, the first explicitly E.A.-affiliated congressional candidate, in a crowded Democratic primary in a new Oregon district. Flynn ran on a longtermist message about pandemic preparedness. (His background was in A.I. safety, but this was clearly a non-starter.) He did little to tailor his platform to the particular needs of the local constituency, which has a substantial Latino population, and he lost by a large margin.

Part of the initial attraction of the movement, for a certain sort of person, was that E.A. existed in a realm outside the business of politics as usual. Now its biggest funder was doing something that looked a lot like an attempt to buy an open congressional seat. There wasn't necessarily anything wrong, from a means-end perspective. But it did seem as though, overnight, the ground had shifted underneath the movement's rank and file. From the perspective of the early days of hard benchmarks, the opportunity cost of ten million dollars spent on a long-shot primary was about twenty-five hundred lives.

One of the reasons MacAskill is so venerated by his followers is that, despite his rarefied intellect, he seems to experience the tensions of the movement on a somatic level. In late March, he gave a talk at EAGxOxford, a conference of some six hundred and fifty young E.A.s. MacAskill had celebrated his thirty-fifth birthday the night before—a small group of largely E.A.-unaffiliated friends had gone out into the fields in pagan costumes to participate in a Lithuanian rite of spring. MacAskill told me that he'd never been happier with his life, and he definitely looked a little worse for wear. At the conference, he was introduced, to rapturous applause, under a portrait of George III wearing a gold damask suit. The room featured a series of ornately carved wooden clocks, all of which displayed contrary times; an apologetic sign read "Clocks undergoing maintenance," but it was an odd portent for a talk about the future. Afterward, MacAskill had a difficult time negotiating his exit from the marbled hall—he was constantly being stopped for selfies, or interrupted to talk about some neglected nuclear risk by a guy dressed like Mad Max, or detained by a teen-ager who



"Remember, he created us in his image a really long time ago."



• •

wanted to know how he felt about the structural disadvantages that kept poor countries poor.

One young woman, two months shy of her high-school graduation, told him that she had stayed up all night fretting—she felt bad that she had paid for private lodging for the weekend, and wanted to know how to harmonize her own appetites with the needs of others. When MacAskill speaks, he often makes a gesture that resembles the stringing of gossamer in midair, as if threading narrow bridges across pitfalls in understanding. He told the young woman that he tried to cultivate his own disposition so that the contradictions disappeared: "E.A. has motivated me to do stuff that's hard or intimidating or makes me feel scared, but our preferences are malleable, and these activities become rewarding." He warned her, however, that it was "pretty easy to justify anything on altruistic grounds if your reasoning is skewed enough. Should I have a less nice apartment? Should I not have Bluetooth headphones?" He sighed and fluttered his eyelids, unable to provide the answers she sought. "After all this time, I guess I don't have a better suggestion for what to do than to give ten per cent. It's a costly signal of your moral commitment." Beyond that, he continued, "try to do the best you can and not constantly think of the suffering of the world."

MacAskill, who still does his own laundry, was deeply ambivalent about the deterioration of frugality norms in the community. The Centre for Effective Altruism's first office had been in an overcrowded firetrap of a basement beneath an estate agent's office. "I get a lot of joy thinking about the early stages—every day for lunch we had Sainsbury's baguettes with hummus, and it felt morally appropriate," MacAskill told me. "Now we have this nice office with catered vegan lunches. We could hire a hedge-fund guy at market rates, and that makes sense! But there's an aesthetic part of me that feels really sad about these compromises with the world."

I asked about the slippage, in his response, from moral to aesthetic propriety. He said, "Imagine you're travelling through a foreign country. During a long bus ride, there's an explosion and the bus overturns. When you come to, you find yourself in a conflict zone. Your travel companion is trapped under the bus, looking into your eyes and begging for help. A few metres away, a bloody child screams in pain. At the same time, you hear the ticking of another explosive. In the distance, gunshots fire. That *is* the state of the world. We have just a horrific set of choices in front of us, so it *feels* virtuous, and morally appropriate, to vomit, or scream, or cry." MacAskill replenishes his own moral and aesthetic commitment through his personal giving, even if he can now fund-raise more in an hour than he could donate in a year.

In "What We Owe the Future," he is careful to sidestep the notion that efforts on behalf of trillions of theoretical future humans might be fundamentally irreconcilable with the neartermist world-on-fire agenda. During a break in the conference, he whisked me to a footpath called Addison's Walk, pointing out the fritillaries, and a muntjac deer in the undergrowth. "We need to stay away from totalizing thinking," he said. "These thought experiments about suffering now versus suffering in the future—once you start actually doing the work, you're obviously informed by common sense. For almost any path, there's almost always a way to do things in a coöperative, nonfanatical way." Pandemic preparedness, for example, is equally important in the near term, and some people think that A.I. alignment will be relevant in our lifetimes.

Members of the mutinous cohort told me that the movement's leaders were not to be taken at their word—that they would say anything in public to maximize impact. Some of the paranoia—rumor-mill references to secret Google docs and ruthless clandestine councils—seemed overstated, but there was a core cadre that exercised control over public messaging; its members debated, for example, how to formulate their position that climate change was probably not as important as runaway A.I. without sounding like denialists or jerks. When I told the disaffected E.A. that MacAskill seemed of two minds about longtermism as an absolute priority, she was less convinced of his sincerity: "I think Will does lean more toward the fanatical side of things, but I think he has the awareness—off the merit of his own social skills or feedback—of the way the more fanatical versions sound to people, and how those might affect the appeal and credibility of the movement. He has toned it down in his communications and has also encouraged other E.A. orgs to do the same." In a private working document about how to pitch longtermism, extensive editing has reduced the message to three concise and palatable takeaways.

The disaffected E.A. warned me to

be wary whenever MacAskill spoke slowly: these were the moments, she said, when he was triaging his commitment to honesty and the objectives of optimized P.R. With so many future lives at stake, the question of honor in the present could be an open one. Was MacAskill's gambit with me—the wild swimming in the frigid lake—merely a calculation that it was best to start things off with a showy abdication of the calculus?

But, during my week in Oxford, it was hard to shake my impression of him as heartrendingly genuine—a sweaty young postulant who had looked into the abyss and was narrating in real time as he constructed a frail bridge to the far side. I asked him what made him most apprehensive, and he thought for a moment. "My No. 1 worry is: what if we're focussed on entirely the wrong things?" he said. "What if we're just wrong? What if A.I. is just a distraction? Like, look at the Greens and nuclear power." Panic about meltdowns appears, in retrospect, to have driven disastrously short-term bets. MacAskill paused for a long time. "It's very, very easy to be totally mistaken."

We returned to the conference courtyard for lunch, where an eclectic vegan buffet had been set up. The line was long, and MacAskill had only five minutes free. He tried to gauge the longest amount of time he could spend queuing, and in the end we contritely cut in at about the halfway point. The buffet table had two stacks of plates, and a fly alighted briefly on one of them. In MacAskill's presence, it's difficult not to feel as though everything is an occasion for moral distinction. I felt that I had no choice but to take the plate the fly had landed on. MacAskill nodded approvingly. "That was altruistic of you," he said.

The Future Fund has offices on a high floor of a building in downtown Berkeley, with panoramic views of the hills. The décor is of the equations-on-a-whiteboard variety, and MacAskill told me that the water-cooler talk runs the gamut from "What are your timelines?" to "What's your p(doom)?"—when will we achieve artificial general intelligence, and what's your probability of cataclysm?

When I visited recently, Nick Beckstead, the C.E.O., assembled the team for a morning standup, and began by complimenting Ketan Ramakrishnan, a young philosopher, on his dancing at an E.A. wedding they'd all attended. The wedding had been for Matt Wage, the early earning-to-give convert. The employees had planned to go to Napa for the weekend, but they were completing their first open call for funding, and there was never a moment to spare. First, some had skipped Friday's rehearsal dinner. Then they figured that they wouldn't be missed at the Sunday brunch. In the end, they'd left the reception early, too. Wage understood. The opportunity cost of their time was high. The Future Fund agreed to finance sixty-nine projects, for a total of about twenty-seven million dollars. The most heavily awarded category was biorisk, followed by A.I.-alignment research and various forecasting projects; the team had funded, among other things, the mushroom caterers of the coming nuclear winter.

Beckstead's new role, and accumulated life experience, seemed to have mellowed his more scholarly inclinations. "I personally find it tough to be all in on a philosophical framework in the fanatical sense," he said. "Longtermism has been my main focus, and will be my main focus. But I also feel like there's some value in doing some stuff that does deliver more concrete wins," and which shows that "we're morally serious people who are not just doing vanity projects about fancy technology." It remains plausible that the best longtermist strategy is more mundanely custodial. In 1955,



the computer scientist John von Neumann, a hero of E.A.s, concluded, "What safeguard remains? Apparently only day-to-day—or perhaps year-to-year—opportunistic measures, a long sequence of small, correct decisions." MacAskill had worried that one of the best new initiatives he'd heard about—the Lead Exposure Elimination Project, which was working to rid the world of lead poisoning—might be a hard sell, but everyone had readily agreed to fund it.

From the outside, E.A. could look like a chipper doomsday cult intent on imposing its narrow vision on the world. From the inside, its adherents feel as though they are just trying to figure out how to allocate limited resources—a task that most charities and governments undertake with perhaps one thought too few. "A.I. safety is such an unusual and uncertain area that it's tempting to simply hope the risks aren't real," Ramakrishnan said. "One thing I like about the E.A. community is that it's willing to deal with the knottiness, and just try to reason as carefully as possible." There were also signs that E.A.s were, despite the hazard of fanaticism, increasingly prone to pluralism themselves. Open Philanthropy has embraced an ethic of "worldview diversification," whereby we might give up on perfect commensurability and acknowledge it is O.K. that some money be reserved to address the suffering of chickens, some for the suffering of the poor, and some for a computational eschatology. After almost a decade of first-principles reasoning, E.A.s had effectively reinvented the mixed-portfolio model of many philanthropic foundations.

One sweltering afternoon, MacAskill and I went for a walk in the Berkeley hills. What had begun as a set of techniques and approaches had become an identity, and what was once a diffuse coalition had hardened into a powerful but fractious constituency; the burden of leadership fell heavily on his shoulders. "One of the things I liked about early E.A. was the unapologetic nature of it," he said. "Some charities are better by a lot! There was this commitment to truth as a primary goal. Now I constantly think of how people will respond to things—that people might be unhappy." He strung another invisible thread in the air. "Am I an academic who says what he thinks? Or am I representative of this movement? And, if so, what responsibilities do I have? There are some things that become compromised there." We passed People's Park, which had become a tent city, but his eyes flicked toward the horizon. "Sometimes, as I think about what I'm going to do after the book comes out, I think, I have a job as the intellectual face of this broader movement, and sometimes I just want to be an independent pair of eyes on the world." ♦