

"Moral Weirdos": Effective Altruism and Empathy in Engineering Education

Dr. Richard A House, Rose-Hulman Institute of Technology

Richard A. House is Associate Dean for Professional Development and Professor of English at Rose-Hulman Institute of Technology. He received a B.A. from Illinois Wesleyan University and M.A. and Ph.D. from the University of California, Irvine. His interests include liberal education for engineers, engineering communication, and the rhetoric of science and technology. With Richard Layton, Jessica Livingston, and Sean Moseley, he is co-author of *The Engineering Communication Manual* (Oxford UP 2017).

“Moral Weirdos”: Effective Altruism and Empathy in Engineering Education

Abstract

Public attention to the effective altruism (EA) movement—in which utilitarian moral calculations are applied to career choices, seeking to maximize the good of an individual’s work—has exploded over the last year. Unfortunately, that attention is currently monopolized by the scandal around FTX founder Sam Bankman-Fried, EA’s highest-profile champion. At first glance, EA appears broadly consonant with broad trends in engineering education dedicated to applying engineering work toward global human problems. Indeed, “EA for Engineers” seems to attract humanitarian and civically engaged engineers and engineering students very similarly to initiatives like NAE’s Grand Challenges for Engineering and Engineers for a Sustainable World.

Yet EA also differs from those enterprises, most starkly in extolling the philanthropic “Money-Maker” over mere “Direct Benefitters.” Michael Lewis paraphrases MacAskill’s argument that a gifted young person should become a banker rather than a doctor [or, by implication, an engineer] serving the developing world: someone else does the investment banking, “very little of the replacement banker’s earnings would find its way to doctors in Africa. All those people you might have saved if you had become a banker and given away your money would die.” In this way, EA runs directly against most cases for “Engineering as an altruistic STEM career” [1]. EA doesn’t pose a higher good of civic involvement or service against the temptation to maximize elevated personal earnings. Most surprisingly, EA advocates often treat empathy, a central focus of much current engineering education research, as a mere shortcoming in decision-making, misdirecting altruistic impulses to causes that are emotionally engaging but logically subordinate to causes with more proven need or measurable benefit to be discovered by “running the numbers.” This paper looks at the prominence of effective altruism among STEM professionals, considers its tensions with practices in empathic and socially engaged engineering education, and treats the FTX collapse as a case study through which students can deliberate on these issues.

Introduction

When I started following the rapid collapse of the FTX cryptocurrency exchange in November 2022, I was already considering the possibility of a case study for my engineering students. Students at my institution (and, I suspect, elsewhere) had been enthusiastic about investments in cryptocurrency, even forming an official school club. I thought that studying a spectacular failure in crypto might leverage students’ interests in the manner of other case studies in engineering ethics and communication, such as the Challenger and Columbia space shuttle tragedies. I often develop lessons out of such news stories when they lead with an obvious ethical lapse, and I was particularly reminded of the 2018 fraud charges against the blood-testing startup Theranos. The FTX headlines weren’t too different: in essence, FTX deposits, converted into cryptocurrency tokens, had been misappropriated to fund risky investments by its sister company, Alameda

Research. FTX founder Sam Bankman-Fried (SBF), like Theranos's Elizabeth Holmes before him, gave a face to the scandal; even the romantic relationships among the companies' principals played similar roles in prompting tabloid-style coverage. (When I reached the point of discussions with my students some months later, I learned that several of them had followed the affair on YouTube.)

FTX can't be cleared of those lapses, and Sam Bankman-Fried has been convicted on seven counts of conspiracy and fraud, with additional charges pending. The company's and founder's relationship to professional ethics, though, appeared quite different once I learned that SBF was a leading figure in *effective altruism* (EA). The EA movement was started by Oxford philosophers to pursue utilitarian goals in philanthropy—for instance, identifying charitable causes that make the greatest measurable contributions to saving human lives for each dollar donated.

The simplest account of FTX is surely that EA was a mere smokescreen to conceal financial crimes by Bankman-Fried and other executives. However, in studying interviews with SBF and the writings of the leading EA utilitarians—William MacAskill, Toby Ord, and Peter Singer—I have always sensed more authentic moral commitment than chicanery. Moreover, the strict metrics for accomplishing good made clear sense of MacAskill's claims that EA has always recruited most successfully in STEM fields [2], appealing to a small subset of philanthropic quants that he affectionately but self-consciously labels “moral weirdos” [3]. EA amounts to a cutting-edge version of utilitarianism's traditional appeal among those students. Nelson and Peterson's 1982 title might overgeneralize—“If You're an Engineer, You're Probably a Utilitarian”—but it's hard to argue with their assessment that the utilitarian's methodical “quantification is attractive to engineers.” [4]

For ASEE and LEES, that attraction is worth examining for several reasons. First, EA creates a direct appeal for engineers to contribute to a variety of public goods, and we should ask whether that appeal retains value despite its present reputational challenges. Second, in that same effort to connect engineering with world problems, EA conflicts with our present emphasis on *empathy*, especially in design education. For the strictest utilitarians, empathy is less a route to good outcomes than a source of moral error. According to psychologist Paul Bloom, author of *Against Empathy*, individual tragedies and local charities stir emotions, while rational calculation alone can direct our good intentions to their maximum impact. (According to EA philanthropic resources, that impact is presently greatest in a set of causes including malaria prevention and animal welfare.) Finally, the difference is even starker when it comes to individual career choices. Our attention to empathy in design aligns “doing good” with the practice of engineering itself; effective altruist Will MacAskill calls this a “direct benefit” strategy. He counters with the prospect of “earning to give,” in which the individual career is directed toward maximum financial return, betting on accumulated wealth that can outweigh a career of direct benefit once donated to charity according to EA principles.

I suspect that other socially engaged engineering educators might share my ambivalent initial reaction to this argument. On one hand, there is tactical appeal—especially for students who came into engineering motivated by the degree's financial return on investment, EA can make a convenient case that a career can aim at altruism without unreasonable self-abnegation. At the

same time, even a massive philanthropic payout may give us pause if our graduates are deliberately making the most lucrative decision at every turn. We can posit that an engineer who affiliates with EA takes seriously a calling to social welfare, and that the pursuit of wealth is merely a means to benevolent ends, and still worry—for instance, about ethical lapses made in the sincere belief that “the ends justify the means.”

SBF and FTX: The Road to Ruin

A basic description of the 2022 FTX collapse will suffice for this argument; this is the condensed treatment provided by *The New York Times*.

On Nov. 2, the crypto publication CoinDesk reported on a leaked document that appeared to show that Alameda Research, the hedge fund run by Mr. Bankman-Fried, held an unusually large amount of FTT tokens [the “native cryptocurrency token for FTX”]. Alameda’s need for funds to run its trading business was a big reason Mr. Bankman-Fried created FTX in 2019. But the way the two entities were set up meant that trouble in one unit shook up the other as crypto prices began to drop in the spring.

[Rival cryptocurrency exchange] Binance announced on Nov. 6 that it would sell its FTT tokens “due to recent revelations.” In response, FTT’s price plummeted and traders rushed to pull out of FTX, fearful that it would be yet another fallen crypto company.

FTX scrambled to process requests for withdrawals, which amounted to an estimated \$6 billion over three days. It seemed to enter a liquidity crunch, meaning it lacked the money to fulfill requests. [5]

The *New York Times* makes no mention of effective altruism in this story, but by all accounts, Bankman-Fried’s EA beliefs drove the career path that led to FTX and its failure. Here, I follow Michael Lewis’s account in *Going Infinite: The Rise and Fall of a New Tycoon* [3]. These are the broad contours:

1. Bankman-Fried’s utilitarian convictions originated with his parents, Barbara Fried and Joseph Bankman, professors at Stanford Law School. While both parents are professed utilitarians, Fried’s scholarship stands out here—her work routinely engages with major thinkers in both philosophical ethics and moral psychology, while Bankman focuses more narrowly on tax policy. Lewis carefully (and probably rightly) avoids any sense that the couple indoctrinated their son, describing the household more as an intellectual salon in which Sam was more generally shaped by deep conversations among the Bankman-Frieds and their guests. Nevertheless, Sam’s thinking came to resemble his parents’. In particular, Fried’s writing formulates moral problems—for instance, about the role of uncertain risks in utilitarian reasoning—to which Sam’s career can be seen as a response.
2. In 2012, Bankman-Fried was recruited to EA directly by Will MacAskill, a young Oxford philosopher who has written seminal EA texts and cofounded, with his colleague Toby Ord, most of the key EA nonprofit institutions. Lewis quotes a subsequent MacAskill

email to a colleague: “He was brought up as a utilitarian by his parents, two Stanford professors, is serious, dedicated, and committed to doing good, and seems really smart and sensible too (i.e. takes some of the weirder ideas seriously but isn’t fanatic about them). He’s thinking about earning to give or going into politics.” [3]

3. At about the same time, Bankman-Fried tested the Earning to Give concept with an internship at the New York proprietary trading firm Jane Street Capital. Lewis notes that many MIT students pursue Wall Street careers; Bankman-Fried was notable in part for his philanthropic motives and in part for an almost preternatural skill navigating games of probability and uncertainty. Lewis links this ability to his favorite pastimes—the collectible card game *Magic: The Gathering* and then the computer game *Storybook Brawl*—but there is a connection to EA here as well. Like financial trading, the movement is predicated in part on sophisticated assessment of probabilities via Bayesian reasoning to chase *expected value*. EA leaders typically present the finance sector merely as an instrument, and sometimes acknowledge an anticipated distaste for Wall Street in their audience. When it comes to risk assessment, though, EA often shares the sophisticated trader’s decision processes and cognitive style.
4. Those decision processes informed Bankman-Fried’s decision to leave Jane Street to found Alameda Research and then FTX. That path began with subject-area interests—Jane Street prohibited its traders from any cryptocurrency transactions—but was navigated according to the EA/trader’s standard calculations: “I tried to estimate the expected value of each [alternative career path], and they were all very similar. A choice between Jane Street and any of them was a close call, but a choice between Jane Street and *all* of them is not. I was asking, *What is the likelihood that Jane Street is the best option?* Low. But it was pretty clear that I wasn’t going to figure it out at Jane Street.” [3]
5. Risk assessments in FTX and Alameda should seemingly have resembled those at Jane Street; the financial sector may change, but the Bayesian techniques should not. Obviously, though, the collapse revealed considerable (and now officially criminal) recklessness. John Ray III, the CEO appointed to oversee the companies’ bankruptcy, declared “Never in my career have I seen such a complete failure of corporate controls and such a complete absence of trustworthy financial information as occurred here.” [6] (This judgment was especially stunning given that Ray had played the same mess-clearing role for Enron after its implosion.)

The Causes of Ineffective Altruism

Between the unsound intercorporate relationship and the extreme volatility of cryptocurrencies, Alameda’s and FTX’s faulty business practices are relatively easy to diagnose. Accounts of the companies’ demise also essentially agree in assigning personal blame to Bankman-Fried (with somewhat more nuance in judging the relative culpability of other executives). Even his personal faults—inexperience, impatience, stubbornness—seem relatively clear. The role of EA, however, sparks considerable debate.

The simplest account of EA's role in FTX is that the philosophical position is a pretext for fraud, at least in Bankman-Fried's adoption of it and even, according to the most sensational YouTube critiques, by proponents like Ord and MacAskill. The most cynical view of the Oxford EAs is that their movement opened the wallets of tech billionaires for initiatives that might not always pass muster by EA's own criteria, such as the 2021 purchase of a 15th century Oxfordshire manor house, Wytham Abbey, to host EA events. A benefactor donated approximately £15 million to the Center for Effective Altruism for the purchase; its expected value was immediately debated in EA communities.

Another interpretation seeks to defend EA by asserting that FTX represents a misuse or unacceptable application of its ideas. MacAskill responded to the crisis on Twitter, citing the passages in his own books that recommend limiting EA actions with "common-sense ethical norms."

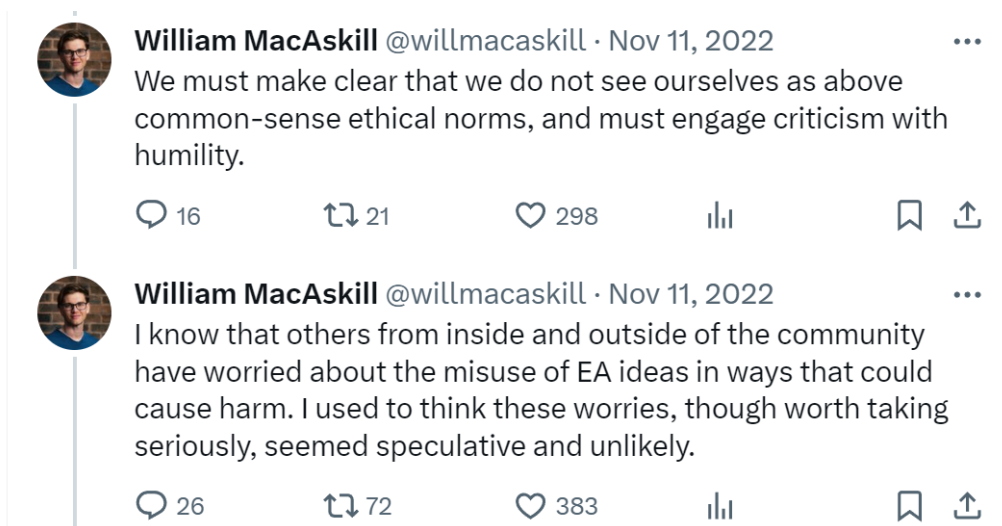


Figure 1. MacAskill's public response to the FTX scandal [7]

One reply, although intensely supportive of EA, calls this thread out as "disingenuous nonsense ...because the whole idea of EA is to step above common-sense ethical norms" to produce better consequences. And MacAskill does seem to invoke weak moral pluralism: if he concedes that expected value calculations aren't always the right way to make moral decisions, he certainly doesn't give us any firm principled limits or criteria beyond intuitive or *ad hoc* judgments. This anonymous poster also correctly finds the same issues in MacAskill's book *What We Owe the Future*:

He tells us that "naive calculations" that the good will outweigh the harm are "almost never correct", without talking about more sophisticated calculations that are correct; that violating rights is "almost never" the best way to bring about good results (almost?); and that the ends "do not always" justify the means (not always - but sometimes? 5% of the time? 49%?). That kind of hedging ... doesn't support the absolute prohibitions declaimed in high dudgeon in MacAskill's twitter thread. [8]

Finally, from a social justice standpoint, EA is criticized not for its individual moral judgments but for the very goal of pursuing the good by aggregating individual acts of altruism. One of the YouTube videos is representative in dismissing it as “just a reheated form of liberal utilitarianism, used to justify the conditions of global exploitation which produces the harms which it apparently seeks to ameliorate.” [9] However, this implication of ideological dishonesty is not a necessary part of this viewpoint. Rather than a broadly underhanded scam, EA may merely present a system with guidance for individual acts but not for collective action or systemic transformation.

STEM, Altruism, Rigor, and Empathy

Effective altruism might be less important for engineering educators if our programs weren’t expressly targeted by the movement’s recruiting pitches, like the one that Lewis describes as prompting Bankman-Fried’s conversion:

Even before [MacAskill] was done, he knew the sort of person who’d be coming up to speak with him: the sort of person who scored an 800 on their math SAT, and understood that the test was too crude to capture their full aptitude.... Roughly three in four of the people who approached MacAskill after one of his talks were young men with a background in math or science. ‘The demographics of who this appeals to are the demographics of a physics PhD program,’ he said. ‘The levels of autism ten times the average. Lots of people on the spectrum.’ [3]

Caroline Ellison, CEO of Alameda Research and Bankman-Fried’s sometime romantic partner, agrees: “I was attracted to people thinking about what to do in a quantitative, rigorous way. Before that, I don’t think I had much of an impulse to do good in the world.” [3]

Engineering educators will recognize the implication that formalized calculations make EA more credible than other, “softer” forms of moral reasoning, which might be tainted by emotion, bias, or other subjective sources of error. This case for EA is made most assertively by the psychologist Paul Bloom:

Empathy is biased; we are more prone to feel empathy for attractive people and for those who look like us or share our ethnic or national background. And empathy is narrow; it connects us to particular individuals, real or imagined, but is insensitive to numerical differences and statistical data... Laboratory studies find that we really do care more about the one than about the mass, so long as we have personal information about the one. In light of these features, our public decisions will be fairer and more moral once we put empathy aside. [10]

The *Boston Review* hosted an interdisciplinary forum in which well-known public intellectuals responded to these ideas; interestingly, one of the most nuanced responses comes from Barbara Fried, who endorses Bloom’s call for more rational/utilitarian public policy but insists that “this is not an argument against empathy; it is an argument against taking it to extremes or in pathological directions.” [11]

Bloom opens his argument, in both essay and book-length versions, by acknowledging its provocative quality: “I often say I am writing a book about empathy. People tend to smile and nod, and then I add, ‘I’m against it.’ This usually gets an uncomfortable laugh... I’ve come to realize that taking a position against empathy is like announcing that you hate kittens—a statement so outlandish it can only be a joke.” [10]

I imagine a similar reaction if a title like *Against Empathy* were on the program at a conference like ASEE. In his genealogy of the concept of empathy in engineering education, Xiaofeng Tang found that in 2017, 78 ASEE conference papers contained the word “empathy”—exponential growth since the late 1990s, when it might appear in a single paper [12]. The trend continued into the pandemic years, reaching 131 “empathy” papers in 2022 (before the count unexpectedly cratered to 10 last year). Bloom would no doubt see the field as caught up in the same enthusiasm, seen across academic disciplines and popular culture alike, that he seeks to temper and correct. Indeed, he precedes his critique by making some concessions to the work of C. Daniel Batson, originator of the “empathy-altruism hypothesis”—a common foundation for research on the development of empathy in engineering students. [13] [14]

Such research often illuminates the question that animates both Bloom’s work and the responses to it. Empathy is traditionally understood primarily as a *feeling*—but models like Batson’s encompass a range of cognitive as well as affective experiences, from “theory of mind” to “empathic distress.” [13] [14] As a result, contested definitions and distinctions create complex lines of disagreement about which types of other-oriented perceptions or sensations count as “empathy” and which have the most beneficial effects.

In the engineering education literature, a convenient starting point is the “operational definition” given by Justin Hess and Nicholas Fila: “empathy includes both affective experiences and cognitive processes that may be primed automatically or within the subconscious, and that may operate in isolation or concurrently, but which tend to have a cyclical relationship.” [14] Hess and Fila organize the cycle as follows, with each quadrant representing one variety of empathy in Batson’s scheme:

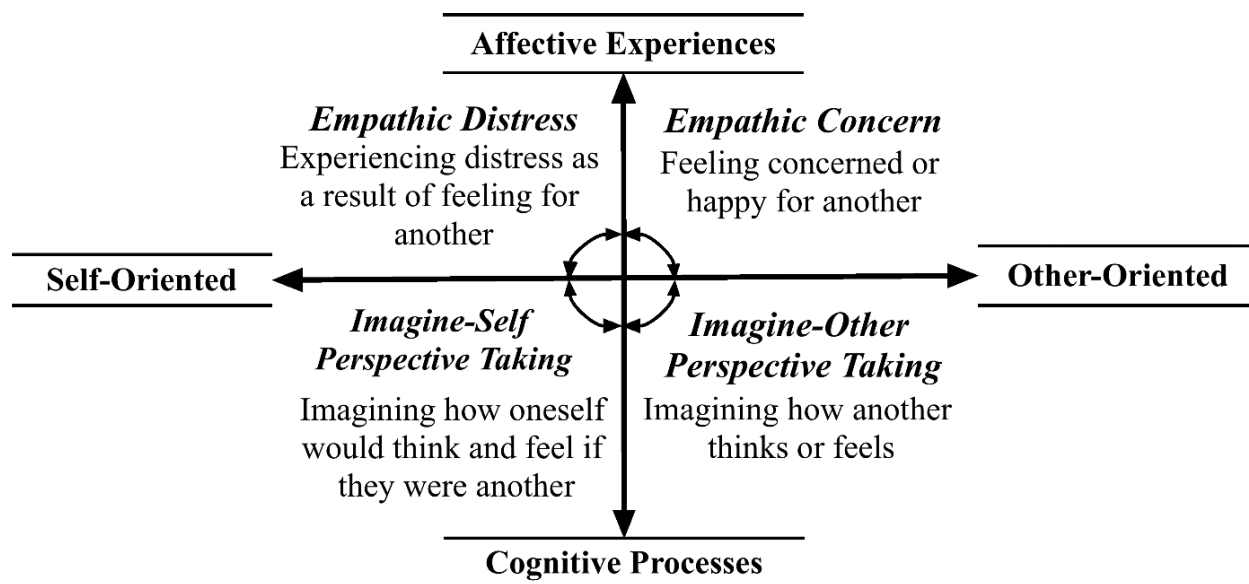


Figure 2. Hess and Fila’s model of “interrelationship between empathy types” [14]

The cyclical model, of course, suggests that growth or development in any or all of the quadrants constitutes desirable personal development. It may also support critiques of *Against Empathy*; Fried doesn’t use such a visualization but notes that Bloom may sometimes conflate the *affective/cognitive* distinction and the *self-/other-oriented* distinction.

Moreover, the cyclical model provides a nuanced way of considering variations in individual development or temperament—“that individuals vary in their propensity for specific empathy types” [14]—rather than dividing cleanly, as the EA movement tends to do, between rational altruism and emotionally driven moral errors. MacAskill, for instance, cautions, “When a disaster strikes, the emotional centers of our brain flare up: we think—*emergency!* ... Because disasters are new and dramatic events, they inspire deeper and more urgent emotions, causing our subconscious to mistakenly assess them as more important or worthy of attention.” [15] Peter Singer’s assessment is comparable: “My admittedly impressionistic observation is that effective altruists are not especially empathetic... Many of them have backgrounds in math or computing. They are moved by the fact that a relatively modest level of donation can help so many people. Rarely do they mention being emotionally moved.” [16]

Accounts like these exceed cautionary notes about emotional excesses or distortions; at least by emphasis, they suggest instead that any positive role for emotion is negligible compared with its potential to undermine rational optimization.

Solving for SBF

MacAskill and Singer clearly appreciate the benefits for moral calculation of intensive education in STEM subjects—and looking at the curriculum taken by my students, I have to imagine that more extensive and rigorous instruction in statistics and probability might be their prescription for creating more effective altruists. It’s less certain that they would want such courses prioritized

over service-learning or human-centered design, but it's hard to imagine them being motivated, as are many engineering education researchers, by affective growth. Where Hess and Fila express interest in "alleviating absolutist/positivistic biases" through participatory design alongside users [14], for instance, it's easy to imagine EA arguments that positivism in such settings corrects bias rather than introducing it.

In recent weeks, while my students work through the later stages of a community-based service-learning project for a local nonprofit, they have also been reading excerpts from Lewis's *Going Infinite* and evaluating effective altruism through Bankman-Fried's story. We have confronted the question of whether EA would treat their design project as a suboptimal use of their time on a cause of merely local significance. We have also, ironically enough, read Bankman-Fried's expressions of loathing for books as an object of study. Lewis quotes a blog written during his sophomore year at MIT:

I could go on and on about the failings of Shakespeare...but really I shouldn't need to: the Bayesian priors are pretty damning... When Shakespeare wrote almost all Europeans were busy farming, and very few people attended university; few people were even literate—probably as low as ten million people. By contrast there are now upwards of a billion literate people in the Western sphere. What are the odds that the greatest writer would have been born in 1564? The Bayesian priors aren't very favorable. [3]

Bankman-Fried's chafing at humanities requirements is hardly unusual among MIT students. This argument, though, already shows the cognitive style that would bring him to EA, to Jane Street, and to FTX—the reduction of assigning value to a calculation of probability—judging aesthetic merit by dismissing any attention to aesthetic objects themselves.

Lewis's book has been pilloried for "credulous" [17] and "generous" [18] coverage of Bankman-Fried. Critics are surely right to see another version of the math-whiz misfit type that Lewis has covered affectionately since *Moneyball*—but Lewis has also written the passages that they cite as evidence that Bankman-Fried is "monstrous." [18] One anecdote, about a clever bet on the results of a series of wagers among Jane Street interns [19], seems expressly intended by Lewis to illustrate Bankman-Fried's cruelty:

Everyone but Sam was unsettled by Asher's humiliation. And yet Sam was ever so slightly taken aback, a few weeks later, when his superiors expressed their dismay at what he'd done... He wasn't surprised to learn that Asher Mellman felt wounded. What surprised him was that his Jane Street bosses thought that he somehow missed the effect he was having on other people. He'd known exactly what he was doing. [3]

One conclusion—along with the Jane Street executives—might be that Bankman-Fried's MIT education could have used more empathy development after all. Another is that the incident actually shows the presence of interpersonal feelings—albeit sadistic ones—that Bankman-Fried denies elsewhere in the book. It's harder to conclude that readers are to leave the book with our bets placed on effective altruism as a reliable path to just decisions.

- [1] J. Lakin, D. Marghitu, V. Davis, and E. Davis, "Introducing Engineering as an Altruistic STEM Career," *The Science Teacher*, pp. 29-35, March/April 2021.
- [2] M. Lewis, *Going Infinite: The Rise and Fall of a New Tycoon*. New York: W. W. Norton, 2023.
- [3] "Ezra Klein Interviews Will MacAskill," *The New York Times*, Aug. 9, 2022. [Online] Available: <https://www.nytimes.com/2022/08/09/podcasts/transcript-ezra-klein-interviews-will-macaskill.html>
- [4] C. Nelson and S. Peterson, "If You're an Engineer, You're Probably a Utilitarian," *Issues in Engineering: Journal of Professional Activities*, vol. 108, no. 1, pp. 13-18.
- [5] K. Huang, "Why Did FTX Collapse? Here's What to Know," *The New York Times*, Nov. 10, 2022. [Online] Available: <https://www.nytimes.com/2022/11/10/technology/ftx-binance-crypto-explained.html>
- [6] E. Brown, "New CEO Says FTX Suffered 'Complete Failure of Corporate Controls,'" *The Wall Street Journal*, Nov. 17, 2022. [Online] Available: <https://www.wsj.com/livecoverage/stock-market-news-today-11-17-2022/card/new-ceo-says-ftx-suffered-complete-failure-of-corporate-controls--anHR8VGla37oobtZDlnh>
- [7] W. MacAskill, "We must make clear that we do not see ourselves as above common-sense ethical norms, and must engage criticism with humility." *Twitter*. [Online] Available: <https://twitter.com/willmacaskill/status/1591218036781895680>
- [8] "Ours is a High and Lonely Destiny: overcoming disingenuous cant from high priest of Effective Altruism," *Further or Alternatively*, Nov. 14, 2022. [Online] Available: <https://furtheroralternatively.blogspot.com/2022/11/ours-is-high-and-lonely-destiny.html>
- [9] J. Duncan, "Effective Altruism: Bourgeois Morality for the 21st Century," *YouTube*, Jun. 1, 2023. [Online] Available: https://www.youtube.com/watch?v=jvpGdoKTK_I
- [10] P. Bloom, "Against Empathy," *Boston Review*, Aug. 20, 2014. [Online] Available: <https://www.bostonreview.net/forum/paul-bloom-against-empathy/>
- [11] B. H. Fried, "Emotional Empathy Is Not the Culprit," *Boston Review*, Aug. 20, 2014. [Online] Available: (https://www.bostonreview.net/forum_response/barbara-h-fried-response-against-empathy-fried/)
- [12] X. Tang, "From 'Empathic Design' to 'Empathic Engineering': Toward a Genealogy of Empathy in Engineering Education," *ASEE Annual Conference Proceedings*, 2018.
- [13] C. D. Batson, "These Things Called Empathy: Eight Related but Distinct Phenomena," in *The Social Neuroscience of Empathy*, J. Decety and W. Ickes, eds. Cambridge: MIT Press, 2009, pp. 16-28.
- [14] J. Hess and N. Fina, "The Development and Growth of Empathy Among Engineering Students," *ASEE Annual Conference Proceedings*, 2016.

- [15] W. MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference*. New York: Avery, 2015.
- [16] P. Singer, “Doing the Most Good,” *Boston Review*, Aug. 20, 2014. [Online] Available: https://www.bostonreview.net/forum_response/peter-singer-response-against-empathy-peter-singer/
- [17] J. Szalai, “Even Michael Lewis Can’t Make a Hero Out of Sam Bankman-Fried,” *The New York Times*, Oct. 2, 2023. [Online] Available: <https://www.nytimes.com/2023/10/02/books/review/going-infinite-michael-lewis.html>
- [18] J. Szalai, “Going Infinite by Michael Lewis review – falling for the antihero,” *The Guardian*, Oct. 25, 2023. [Online] Available: <https://www.theguardian.com/books/2023/oct/25/going-infinite-by-michael-lewis-review-falling-for-the-antihero>
- [19] R. Hecklen, “What to expect when you’re expecting expected value,” *Bayes Shammai*, Dec. 14, 2023. [Online] Available: <https://bayesshammai.substack.com/p/what-to-expect-when-youre-expecting>