

# Assessing Evidence for Purposes of Effective Altruism

KATHRYN E. JOYCE\*

## TABLE OF CONTENTS

I. EFFECTIVE ALTRUISM AND THE LIMITS OF THE INTERVENTION-CENTERED APPROACH . . . . .	759
A. <i>The Intervention-Centered Approach to Supporting Predictions</i> . . . . .	759
B. <i>Evidence from RCTs</i> . . . . .	760
C. <i>Aiming for Evidence of General Effectiveness</i> . . . . .	762
II. SHIFTING TO THE CONTEXT-CENTERED APPROACH . . . . .	764
III. ADDRESSING A COMMON CRITIQUE . . . . .	766
CONCLUSION . . . . .	767

Effective altruism (EA) embraces an evidence-based approach to charitable giving that resonates with the broader evidence-based policy (EBP) movement in medicine, education, and other areas of social policy. EBP holds that decisions about how to produce desired outcomes should be based on evidence of effectiveness from scientific research rather than conventions, intuitions, or personal impressions about what works or doesn't work. EA shares this commitment. At its core, "effective altruism is about using evidence and reason to figure out how to benefit others as much as possible and taking action on that basis."<sup>1</sup> Indeed, Peter Singer describes it as "a commitment not to a particular solution, but to following the evidence."<sup>2</sup>

Since its objective is to do as much good as possible for people in poverty, regardless of where they live, EA relies on research that aims to identify interventions that are effective across a wide range of contexts. In keeping with the dominant EBP approach, effective altruists view results from randomized controlled

---

\* Kathryn E. Joyce is a Postdoctoral Research Associate in Values and Public Policy at Princeton University's Center for Human Values. She will soon join the Ohio State University's Center for Ethics and Human Values as Director of the Civil Discourse Program. © 2023, Kathryn E. Joyce.

1. This definition comes from the Centre for Effective Altruism. See *CEA's Guiding Principles*, CEA, <https://www.centreforeffectivealtruism.org/ceas-guiding-principles> [<https://perma.cc/9982-JS23>].

2. PETER SINGER, *THE LIFE YOU CAN SAVE* 159 (2009).

trials (RCTs) as the gold-standard of evidence for these effectiveness claims. The general idea is that RCTs can be used to figure out whether an intervention ‘works’ or is replicable across a wide range of different settings. I call this the ‘intervention-centered approach’ to EBP because it focuses on the causal capacity of interventions.

Drawing on recent methodological critiques of EBP, I will argue that effective altruists should reject the intervention-centered approach in favor of a context-centered approach that centers the arrangement of causal factors present in target settings.<sup>3</sup> The key difference is that the intervention-centered approach aims to support claims about the general effectiveness of interventions, whereas the context-centered approach aims to support predictions about whether a candidate intervention is likely to be effective in a specific place. Since they aim to support different types of effectiveness claims, these approaches call for evidence from different types of research.

I begin by discussing methodological limitations of the intervention-centered approach that motivate the context-centered approach in general, arguing that they are especially salient given the aims of EA. Importantly, my methodological critique of the *intervention-centered approach* is distinct from recent methodological critiques of *effective altruism*.<sup>4</sup> Although both reference RCTs, the latter tend to concern deeper issues about the values underlying EA, the kinds of effects or interventions effective altruists should prioritize, engagement with target populations, and how decisions about poverty relief efforts should be made.<sup>5</sup> Indeed, some critics argue that EA involves ‘swooping in’ and imposing interventions on communities without considering local considerations and perspectives.<sup>6</sup>

To the extent that these criticisms challenge EA’s evidence-based methodology, they do so by confronting its fundamental commitments.<sup>7</sup> Rather than questioning EA’s commitments, in Section II, I argue that they speak in favor of shifting to the context-centered approach. In Section III, I suggest that doing so

3. Here, I’m employing Nancy Cartwright’s distinction between intervention-centered and context-centered approaches to EBP. See NANCY CARTWRIGHT, *NATURE, THE ARTFUL MODELER: LECTURES ON LAWS, SCIENCE, HOW NATURE ARRANGES THE WORLD AND HOW WE CAN ARRANGE IT BETTER* (2019) [hereinafter CARTWRIGHT, *NATURE*].

4. Some criticisms of effective altruists’ reliance on RCTs also arise in the literature on EBP.

5. For example, see Emily Clough, *Effective Altruism’s Political Blind Spot*, BOSTON REV. (Jul. 14, 2015), <https://www.bostonreview.net/articles/emily-clough-effective-altruism-ngos/> [https://perma.cc/X9Q5-RBFH]; Amia Srinivasan, *Stop the Robot Apocalypse*, 37 LONDON REV. BOOKS (2015). See also Iason Gabriel, *Effective Altruism and its Critics*, 34 J. APPLIED PHIL. 457, 457–73 (2016) (discussing several criticisms).

6. For example, see Kate Manne, *Against Swooping in*, MORE TO HATE (Jun. 30, 2022), <https://katemanne.substack.com/p/against-swooping-in> [https://perma.cc/Y8FM-CP5T]; Jennifer Rubenstein, *Response to Effective Altruism*, BOSTON REV. (Jul. 1, 2015) [hereinafter Rubenstein, *Response*], [https://www.bostonreview.net/forum\\_response/response-jennifer-rubenstein/](https://www.bostonreview.net/forum_response/response-jennifer-rubenstein/) [https://perma.cc/QNY2-9T7U]; Jennifer Rubenstein, *The Lessons of Effective Altruism*, 30 ETHICS & INT’L AFFS. 511 (2016) [hereinafter Rubenstein, *Lessons*].

7. For instance, some condemn EA’s use of RCTs because they think efforts to relieve poverty should be developed locally or because they favor strategies that cannot be tested in RCTs.

would help assuage some common concerns about neglecting local considerations and perspectives in target sites.

### I. EFFECTIVE ALTRUISM AND THE LIMITS OF THE INTERVENTION-CENTERED APPROACH

Across domains, policymakers must decide how best to achieve their goals. EBP involves using the best available evidence to justify policy predictions—*ex ante* causal claims about what will happen if an intervention is used in a target setting. EBP consists of a research side that produces reliable knowledge and a practice side that uses that knowledge as evidence to support predictions.<sup>8</sup> Of course, even warranted predictions do not independently support decisions about policy or practice. Among other things, such decisions require information about cost effectiveness, the relative value of expected effects, and feasibility. The predictions EBP supports can be seen as one premise in an argument that justifies policy decisions.<sup>9</sup> As effective altruists recognize, however, it is a very important premise. Decisions about policies or interventions cannot be justified without good reason to think they will produce effects that contribute to the desired outcomes, whether they be domestic social welfare programs, mathematics curriculum, or humanitarian aid.

#### A. *The Intervention-Centered Approach to Supporting Predictions*

The dominant EBP strategy is intervention-centered in that it aims to support effectiveness predictions about particular target sites indirectly by identifying interventions that are generally effective across settings of interest. RCT results, along with meta-analyses and systematic reviews of them, are widely considered “gold-standard” evidence for these effectiveness claims. For example, the Abdul Latif Jameel Poverty Action Lab (J-PAL) conducts RCTs around the world and summarizes results, highlighting general policy insights in an attempt to reduce poverty by “ensuring that policy is informed by scientific evidence.”<sup>10</sup> The general idea is that RCTs can be used to determine whether an intervention ‘works’ or is replicable across a wide range of different settings. If so, we are justified in predicting that it will work in a particular setting of interest (unless it is highly atypical), provided it is implemented with fidelity (i.e., just as it was implemented

---

8. Likewise, according to [effectivealtruism.org](https://effectivealtruism.org), EA consists of “a research field that aims to identify the most effective ways of helping others, and a practical community of people who aim to use the results of that research to make the world better.” *What Is Effective Altruism?*, EFFECTIVEALTRUISM.ORG, <https://www.effectivealtruism.org/articles/introduction-to-effective-altruism#faq> [<https://perma.cc/2DW9-BXK7>] (last visited March 26, 2023) (under “What is the definition of effective altruism?” in FAQs section).

9. Here I’m invoking Nancy Cartwright’s argument theory of evidence for use. See Nancy Cartwright, *Evidence, Argument and Prediction*, in EPSA11 PERSPECTIVES AND FOUNDATIONAL PROBLEMS IN PHILOSOPHY OF SCIENCE 3, 3–17 (V. Karakostas ed., 2013); NANCY CARTWRIGHT & JEREMY HARDIE, *EVIDENCE-BASED POLICY: A PRACTICAL GUIDE TO DOING IT BETTER* (2012); Julian Reiss, *What’s Wrong With Our Theories of Evidence?* 29 THEORIA 283, 283–306 (2014).

10. See *Policy Insights*, J-PAL, <https://www.povertyactionlab.org/policy-insights> [<https://perma.cc/5MNJ-46S8>] (last visited Dec. 1, 2022).

in the studies).<sup>11</sup> Put roughly, if an intervention—like microfinancing—alleviates poverty in many different places, then we can expect it to alleviate poverty to a similar degree in a new target setting.

In keeping with the intervention-centered approach, EA relies on RCTs to produce evidence about what ‘works.’ Indeed, GiveWell explicitly considers RCTs gold standard evidence of effectiveness; it identifies top charities in large part by considering whether they are using programs that are “highly evidence-based and cost-effective” according to research from organizations like J-PAL for anti-poverty programs and the Cochrane Library for health programs.<sup>12</sup> As Singer explains: “[W]e can learn what works and what does not work, but to do so we must use the most objective possible means of assessing programs. Ideally that will mean randomized controlled trials.”<sup>13</sup>

RCTs are favored for evaluating the impact of interventions because, when internally valid, they justify ascribing observed effects to the intervention. The basic idea is simple: RCTs attempt to create groups that are identical except that one receives the intervention being tested (treatment group) and the other (control group) does not. Participants are randomly assigned to treatment and control groups because random assignment is supposed to help balance known and unknown causal factors between groups. If the groups are similar enough in terms of factors that could affect the outcomes, then differences in outcomes can be ascribed to the intervention. So, RCT results are *causal ascriptions* of the form: “X caused Y in the study.”

Conducting RCTs in real-world environments is challenging, and they are sometimes compromised.<sup>14</sup> However, since I am interested in the inferences drawn from RCT results to other settings for policy purposes, my discussion will take internal validity for granted. The question is whether well-conducted, internally valid RCTs can play the role that the intervention-centered approach assigns to them.

### B. Evidence from RCTs

Causal ascriptions cannot directly support general effectiveness claims or predictions. The fact that an intervention *worked* somewhere cannot show that the

11. This fidelity condition is important because RCTs test a particular version of an intervention and so only provide evidence about the effects of that version. The claim that RCT results can be replicated—or that an RCT-tested intervention will produce similar results somewhere else—assumes that the intervention will be implemented with fidelity.

12. *Research on programs*, GIVEWELL, <https://www.givewell.org/research> [<https://perma.cc/U5QZ-58MJ>] (last visited Dec. 1, 2022).

13. SINGER, *supra* note 2, at 161. MacAskill clarifies that RCTs are not sufficient for decisions about what to do because they don’t answer all relevant questions, like questions about human extinction or ethics. Still, he seems to think they can do the work that EBP assigns to them.

14. See Thomas Cook, *Twenty-six assumptions that have to be met if single random assignment experiments are to warrant “gold standard” status: A commentary on Deaton and Cartwright*, 210 SOC. SCI. & MED. 37, 37–40 (2018); see also Angus Deaton & Nancy Cartwright, *Understanding and misunderstanding randomized controlled trials*, 210 SOC. SCI. & MED. 2 (2018).

intervention *works* or that it *will work* in some other target.<sup>15</sup> These inferences are only justified if the study participants and settings *represent* the entire population of interest or a particular target site. The RCT design does not guarantee representation. Random assignment only balances treatment and control groups. Random sampling (i.e., conducting an RCT on a population that is randomly selected from the whole population of interest) *could* produce a representative sample in theory, if it were large enough, but researchers do not typically draw random samples from the entire population of interest to policymakers.<sup>16</sup>

For effective altruists, the population of interest is quite large—it includes many people in many different places. So, we cannot simply take representation for granted. We need evidence showing that the population or target site is *like* the study site(s) in ways that are *causally relevant* to the intervention. To know if the study sample and the target are similar in the relevant ways, we need information about how the intervention works and information about the contextual factors that affect its effectiveness. RCTs do not provide that. Statistical approaches are often used to make trial populations more closely resemble the target population or to make predictions about a target population using RCT results.<sup>17</sup> However, these will only work if the causally relevant factors are known, but they usually are not—that is why random assignment is taken to be essential for reliable causal ascriptions.

Combining RCT results to create a larger sample, as organizations like J-PAL do, makes representation for the whole population more likely. If successful, it gives us an accurate estimate of the average treatment effect (ATE) for a broad population, but that is not very useful for making predictions. A positive ATE does not show that the intervention generally produces positive effects across many different sites and thus can be expected to do so in new places—it tells us what the average effect would be for the general population if each individual received the intervention.<sup>18</sup> But that is not the kind of general effectiveness claim that supports predictions.

Even if accurate, the ATE for the entire population of interest does not directly support predictions about a subset of that population unless we know that the subset is properly represented by the broader population. In other words, the target would have to be a miniature of the broader population. That would mean, for

---

15. See CARTWRIGHT, NATURE, *supra* note 3; see CARTWRIGHT & HARDIE, *supra* note 9; see also Kathryn Joyce & Nancy Cartwright, *Bridging the Gap Between Research and Practice: Predicting what will work Locally*, 57 AM. EDUC. RSCH. J. 1045 (2020).

16. But, as Cartwright and Deaton point out, representation is not guaranteed and changes that occur post-selection can undermine it. See Deaton & Cartwright, *supra* note 14.

17. See Kathryn Joyce, *The key role of representativeness in evidence-based education*, 25 EDUC. RSCH. & EVALUATION 43, 43–62 (2019) [hereinafter Joyce, *Representativeness*]; see also Deaton & Cartwright, *supra* note 14.

18. As Jack Martin observes in the context of education, the term ‘general’ is often used in different ways when discussing empirical research. Essentially, ‘general’ can mean ‘on average’ for the general population or it can mean ‘common to all’ in the general population. See Jack Martin, *In Defense of Robin Barrow’s Concerns about Empirical Research in Education*, 26 PHIL. INQUIRY IN EDUC. 137–45 (2019).

example, that students in a school district in Kansas have the same distribution of relevant causal factors (e.g., background knowledge, resources, English language skills) as the entire population of students in the U.S. Again, we cannot simply assume this is the case and assessing it requires causal knowledge about how the intervention operates and the contextual factors that affect its effectiveness.

Further, the ATE we get from combining RCT results could be misleading if effects varied significantly across study sites. An intervention that is very effective in some places but not in others can have a very small ATE when RCT results are combined. Although this does not mean that it is not very effective, nor that the trials where it shows large positive effects were mistaken, that is what it might look like. For example, consider Head Start, a free pre-school program for low-income children across the U.S. The ATE for all students in the U.S. who attend Head Start is positive, but modest. Given its low ATE, Head Start seems ineffective, especially relative to its cost. However, researchers examining impact variation across sites found that Head Start is highly effective in some places—the national ATE is low because there were also sites where the effects were negligible.<sup>19</sup> As this example shows, a program that would do a lot of good in some places can appear to have an insignificant impact if we look at the ATE for the larger population. To capitalize on the effects of programs like Head Start, we need to understand why the impact varies.

### *C. Aiming for Evidence of General Effectiveness*

I have argued that RCTs, alone or combined, provide insufficient support for the sort of general effectiveness claims that would warrant predictions about particular cases. This is a limitation of the intervention-centered approach that stems from its reliance on RCTs. But, aside from that, should EA adopt an intervention-centered approach that aims to identify interventions that will work reliably across the wide range of settings with which they are concerned in the first place?

Supporting predictions by establishing general effectiveness claims is common in medicine because interventions often have stable causal capacities. For example, acetaminophen has a stable causal capacity to reduce fevers—it reliably does so across a broad, diverse range of individuals and circumstances. The fact that it is generally effective in this sense justifies the expectation that it will reduce fevers in new cases. The intervention-centered approach assumes that interventions have stable causal capacities. Therefore, adopting it only makes sense for EA if certain interventions (e.g., anti-poverty programs, medication to prevent HIV transmission) can be expected to produce positive effects across many different contexts (e.g., countries with high rates of poverty or of HIV/AIDS) when implemented with fidelity.

Not all ‘causes’ are stable—social interventions and other causes operating in social domains often are *not* stable, or if they do have a stable capacity, the range

---

19. Pamela Morris et al., *New Findings on Impact Variation from the Head Start Impact Study: Informing the Scale-Up of Early Childhood Programs*, 4 AERA OPEN 1 (2018).

of stability is narrow and subject to change as contexts shift and evolve.<sup>20</sup> To say that an intervention has a stable causal capacity across a range of settings is to say that the factors necessary to support it will be present and an appropriate causal pathway will be available in all of them. This is the case for acetaminophen across human beings because it reduces fevers through a biological causal pathway that is available across the vast majority of normally functioning human bodies despite expected differences (e.g., age, height, blood type, resting heart rate, geographical location).

It seems unlikely that this will be true of anti-poverty interventions across most settings of interest to effective altruists given their significant cultural differences, the complex set of factors that cause poverty, and the fact that social settings change in ways that might be relevant to EA programs. Indeed, in its evaluation of global education programs, GiveWell acknowledges the relevance of contextual variation: “We are cautious about drawing general conclusions from individual education evaluations because we think that education interventions, and the mechanisms that they work through, are particularly complex, and there are large differences between different settings where they are implemented.”<sup>21</sup> Since the same concerns exist about many anti-poverty interventions and their target settings, a similar degree of caution is warranted.

If we reject the idea that (most) of the interventions in question will have stable causal capacities, then we should not be trying to support predictions about what *will be effective in particular places* with claims about what *is effective in general*. Instead, we should be considering why, when, and where an intervention works, so we can identify appropriate target sites. That requires studying interventions’ causal mechanisms and considering contextual factors in the target site that bear on them.

EBP proponents agree that context needs to be considered at the implementation stage, but only to determine how the intervention can be implemented with fidelity in the target site, as opposed to whether an intervention *can* work there if it’s implemented with fidelity. As Hakan Seckinelgin observes, the intervention-centered approach

considers the context as something of an ‘add-on’ consideration. The implicit claim here is that the integrity of the research presented stands alone on its scientific grounds and that context becomes only an issue once we move to consider policy implementation, focusing on particular target/risk groups and the way we can deliver to them what we know works.<sup>22</sup>

---

20. CARTWRIGHT & HARDIE, *supra* note 9.

21. *Education in Developing Countries*, GIVEWELL (April 2018), <https://www.givewell.org/international/technical/programs/education> [<https://perma.cc/K46Q-QXR9>] (last visited March 26, 2023).

22. HAKAN SECKINELGIN, *THE POLITICS OF GLOBAL AIDS: INSTITUTIONALIZATION OF SOLIDARITY, EXCLUSION OF CONTEXT* 132 (2017).

Contextual factors cannot be deemed relevant only to implementation because, as Seckinelgin's discussion of Global AIDS policy demonstrates, implementing interventions with fidelity does not ensure their success in a new setting.

The upshot is that predictions about particular target sites require more information than we can get from RCTs or other studies that measure the short or long-term effects of an intervention on the sample population. Additionally, since effective altruists want to produce the best possible outcomes in each target site, they should not favor implementing interventions with fidelity unless they have reason to think that the intervention would do the most good there as a result. However, they would still need more information to figure out how to adjust interventions to maximize their positive effects.

## II. SHIFTING TO THE CONTEXT-CENTERED APPROACH

These limitations of the intervention-centered approach motivate the context-centered approach. Since the latter does not assume that interventions have stable causal capacities, the context-centered approach aims to support predictions about particular targets directly rather than supporting them via general effectiveness claims. For the context-centered approach, the research question is: "what will work here?" rather than "what works?" As Nancy Cartwright explains:

The problem with intervention-centering is that it is too limited in the range of knowledge it employs. It looks chiefly at whether the policy has worked elsewhere; then it uses loosely warranted markers to settle whether the new site is sufficiently similar to those where the policy has been successful to support the same causal pathways. The context-centered approach takes on the difficult job that intervention-centering ducks: understanding the details of the new context well enough to figure out what causal pathways it can and cannot afford. Context-centering can have real advantages. It should not only provide more reliable predictions about the effectiveness of the proposed policy in the new site; it can also ground new proposals for bespoke policies geared to the causal pathway available there.<sup>23</sup>

If Cartwright is right, EA's commitment to using evidence-based thinking favors the context-centered approach. I will briefly describe this approach before considering how adopting it helps EA avoid some common criticisms.

The context-centered approach calls for evidence about what causal pathways to desired outcomes are available in the specific context where interventions are to be implemented.<sup>24</sup> Here, 'causal pathway' refers to the arrangement of causal factors present in the target setting. Such factors include available resources, cultural norms, and infrastructure. The arrangement of causal factors in the target influences what will happen if an intervention is implemented there. Information

---

23. CARTWRIGHT, NATURE, *supra* note 3, at 73.

24. See CARTWRIGHT & HARDIE, *supra* note 9; see also Joyce, *Representativeness*, *supra* note 17.



about them is therefore necessary to support the predictions central to EA. A causal pathway to a desired outcome is present if the arrangement of causal factors supports the intervention's causal process. Consequently, it must include the support factors necessary for the intervention to operate.<sup>25</sup> This is because a 'cause' (e.g., an anti-poverty program) is usually a set of causal factors that, together, contribute to the occurrence of an effect.<sup>26</sup> But in addition to the set of causal factors that comprise it, interventions rely on support factors from the setting where they are implemented. Therefore, the context in which a causal mechanism (i.e., intervention) produces effects is itself part of that mechanism.<sup>27</sup>

Assuming the necessary support factors are present, the intervention's causal process must be able to operate uninterrupted. Thus, causal processes occurring in the target must be compatible with the intervention's causal process. On the other hand, factors in the target setting could undermine, counteract, or dilute an intervention's effects, even if the setting allows its causal process to unfold. For instance, programs that are redundant are unlikely to produce significant positive effects.<sup>28</sup> A causal pathway will only be available, then, if no factors counteract, undermine, or dilute the intervention's effects.<sup>29</sup>

Warranted predictions about how an intervention will perform in a local setting require evidence that a causal pathway will be available during the time of implementation. Predictions will usually be uncertain (i.e., "it will probably work here"), then, because this information is difficult to obtain in advance due to the complex, contingent nature of social settings. Still, as Joyce and Cartwright explain, theory and research that uses a mix of methods, can inform predictions by providing information that is relevant to assessing causal pathways—such as identifying factors that would support or derail the intervention.<sup>30</sup>

Rather than relying solely on RCTs, the context-centered approach calls for an array of theory and research that provides this kind of information to support predictions about particular target settings. As this brief discussion suggests, predictions that include premises about support factors and causal pathways also require local knowledge and judgment about what is possible or desirable in the target setting. This information often comes from local sources. However, RCT

---

25. CARTWRIGHT & HARDIE, *supra* note 9; Joyce & Cartwright, *supra* note 15; *see also* Kathryn Joyce & Nancy Cartwright, *Meeting our standards for educational justice: Doing our best with the evidence*, 16 THEORY & RSCH. EDUC. 3 (2018); TONE KVERNBEKK, EVIDENCE-BASED PRACTICE IN EDUCATION: FUNCTIONS OF EVIDENCE AND CAUSAL PRESUPPOSITIONS (2016).

26. J.L. Mackie, *Causes and Conditions*, 2 AM. PHIL. Q. 4 (Oct. 1965). In technical terms, some causes are INUS conditions for their effects: a set of factors where each factor is *Insufficient* but *Necessary* for producing an effect and the set of factors itself is *Unnecessary* but *Sufficient* for producing the effect.

27. Joseph A. Maxwell, *Causal Explanation, Qualitative Research, and Scientific Inquiry in Education*, 33 EDUC. RESEARCHER 2 (2004).

28. Kathryn Joyce & Nancy Cartwright, *How Should Evidence Inform Education Policy?*, in HANDBOOK OF PHILOSOPHY OF EDUCATION (Randall Curren ed., 2022).

29. CARTWRIGHT & HARDIE, *supra* note 9. *See also* Joyce & Cartwright, *supra* note 15.

30. *See* CARTWRIGHT & HARDIE, *supra* note 9; SECKINELGIN, *supra* note 22.

results (and other trustworthy casual ascriptions) can help support predictions when buttressed with information about the intervention's causal mechanism and the contextual factors that influenced it in the study setting. After all, RCTs show that an intervention can work under some circumstances—those in the study. We can therefore learn about some of the relevant causal factors by looking at that case.

### III. ADDRESSING A COMMON CRITIQUE

EA has been subject to many criticisms. One theme is that focusing on evidence of effectiveness leads effective altruists to neglect local perspectives and considerations. Jennifer Rubenstein's critique represents this view. Rubenstein claims that EA has a "hidden curriculum" that conveys the message that "doing the most good does not require listening to those affected by the issues one is trying to address."<sup>31</sup> Further, Rubenstein says that "[b]y excluding poor people and encouraging a savior complex and insularity among its members, the effective altruism movement fails to meet normative criteria of democracy and equality."<sup>32</sup> She urges effective altruists to adopt a 'pluralist' approach that "includes poor people as partners or follows their lead, even if this means less certainty about doing the *most* good."<sup>33</sup> This involves disavowing "the belief that technically savvy, well-meaning outsiders can kick down doors and save the day."<sup>34</sup>

Effective altruists express a desire to avoid a 'top-down' approach, in part by working with local populations, but they also identify some challenges to doing so. For one, locally designed interventions are not always effective. We cannot just assume that, because they are developed locally, they will have a positive impact, and it is difficult to find evidence that supports predictions about them (even if we do not insist on RCTs). It is also difficult to know how to involve local perspectives in useful ways.

The context-centered approach addresses some of these problems by identifying a role for local perspectives that actually contributes to EA's main goal of using evidence and reason to do good. Because the context-centered approach bases predictions on evidence about causal pathways available in target settings, it requires information about how local norms and culture will affect the behavior of potential aid recipients. This information can come from researchers, but obtaining it would involve engagement with the local community. As the previous section indicates, predictions—not just implementation planning—require information about relevant contextual factors (those that bear on the effectiveness of the intervention), including the social, political, and economic structure. But, importantly, the context-centered approach does not recommend institutional interventions, as its focus is still on evidence of effectiveness.

---

31. Rubenstein, *Lessons*, *supra* note 6, at 519.

32. Rubenstein, *Response*, *supra* note 6.

33. *Id.*

34. *Id.*

Still, the context-centered approach would adjust effective altruists' relationship to the communities they are trying to help. Seckinelgin, describing the intervention-centered approach to Global AIDS policy, finds that

by presenting context as an externality to research, policy thinking is oriented towards controlling [context] so that the research outcomes observed can be made replicable within a generalized policy framework. In other words, the concern here about context is about how to control it, rather than asking what it means for the research output.<sup>35</sup>

The intervention-centered approach transports interventions that have worked somewhere to new settings. This could feel like an imposition, especially if it is insensitive or unresponsive to community features that matter to participants. By contrast, the context-centered approach starts by learning about the community and considering what would work best there. Moreover, since the intervention-centered approach aims to replicate results by implementing interventions with fidelity, it could involve adjusting the context to support implementation in ways that are unwelcome. Instead of focusing on replication and fidelity, the context-centered approach provides resources for adjusting interventions to fit the context so they can produce the best possible results in that setting.

For these reasons, adopting the context-centered approach (or supporting charities that do) rebuffs some criticisms about neglecting local considerations while preserving EA's commitment to using evidence of effectiveness to produce the best results. However, it does not assuage this type of critique entirely. Since EA still favors effectiveness over local involvement, it would not address critics who think values like equality and democracy require local involvement (or a 'bottom-up' grassroots approach) at the cost of effectiveness. Still, it seems as though this is the best EA can do without abandoning its commitment to EBP.

#### CONCLUSION

Effective altruism requires evidence to make predictions about particular places. I argue that effective altruists should shift from an intervention-centered approach that aims to identify interventions that will work across contexts using RCTs to a context-centered approach that requires more information than we get from RCTs and other research that measures the (short or long term) effects of interventions on sample populations. Adopting the context-centered approach should improve predictions about target settings, which should help effective altruists produce better outcomes. Additionally, the context-centered approach involves more engagement with local communities, which means attending to local considerations and giving recipients a more active role.

---

35. SECKINELGIN, *supra* note 22, at 132.