

Review of Ingmar Persson's *Inclusive Ethics: Extending Beneficence and Egalitarian Justice*. Oxford: Oxford University Press, 2017, 288 pp.

TIMOTHY CAMPBELL

Institute for Futures Studies

Inclusive Ethics develops a moral theory governing any conduct that affects individuals that are, or could be, conscious. This is an ambitious task, and the book covers many topics—including the ethics of procreation, the disvalue of premature death, anthropocentric speciesism, the ethics of defensive killing, distributive justice, personal identity, strategies for increasing one's life-satisfaction, and the point of doing moral philosophy. *Inclusive Ethics* has something for anyone interested in normative or practical ethics.

The book is in three parts. Parts I and II focus on the two central components of Persson's moral theory: a principle of beneficence and a principle of equality. Part III focuses mainly on practical issues concerning the implementation of these principles.

Part I argues that there is moral reason to benefit individuals that are, or could be, conscious. Such beings include not only *persons*—those capable of rationality and self-awareness—but also barely sentient beings (such as lizards and frogs), non-conscious beings with the potential to become conscious (such as 3-month-old human fetuses), and even non-existent beings that would have the potential for consciousness if they were to exist (such as the fetus that would exist if some particular sperm and egg were to fuse). We benefit these individuals by promoting what is intrinsically good for them, and by reducing what is intrinsically bad for them.

Chapter 1 develops an account of what is intrinsically good (and bad) for individuals. On what Persson calls *the double-aspect account*, the only two things that are intrinsically good for one are (i) *well-being*, which is construed entirely in terms of experiences that are intrinsically good for one (one's pleasures), and (ii) the fulfillment of one's autonomous desires, those that are formed under certain ideal conditions that include one's being rational and well-informed. (i) and

AUTHOR'S NOTE: Funding from the Swedish Research Council is gratefully acknowledged.

(ii) jointly constitute one's *welfare*, which means that, for Persson, 'well-being' and 'welfare' are not synonymous.

Chapter 2 argues for an inclusive view of beneficence (IVB) according to which we can, and can have moral reason to, benefit individuals by creating them. On this view, our moral reasons to create individuals with positive welfare are *person-affecting*; these reasons derive not from the impersonal value of an individual's welfare, but from the fact that bringing this individual into existence would benefit *it*. Typically, those who defend IVB deny that the sense in which one benefits from being brought into existence entails that one is thereby made *better off* than one would have been had one not existed. For example, Jeff McMahan (2013, 6-7) argues that one's coming into existence could benefit one only in a *non-comparative* sense—that is, only in the sense that one's life contains more of what is intrinsically good than bad for one.

However, according to Persson, the fact that an act would benefit someone in this non-comparative sense is insufficient for there being any moral reason to perform that act. He points out that benefiting an individual in a non-comparative sense could also make that individual worse off than it would otherwise have been. For example, I could perform an act that results in your life containing more of what is intrinsically good than bad for you *and* prevents you from having *even more* of what is intrinsically good for you and no more of what is intrinsically bad for you. In this case, the fact that my act would benefit you in a non-comparative sense provides no reason for me to perform it. Rather, I have reason *not* to perform this act, since my not performing it would leave you better off. According to Persson, examples such as this strongly suggest that the benefits that we have moral reason to bestow are *comparative* in the sense of making individuals better off than they would otherwise have been.

But if this is right, then, as Persson acknowledges, IVB implies that causing someone to exist can make her better off than she would otherwise have been; and this seems to imply that if she had not existed *she* would have been *worse off*. While most find this implication absurd, Persson accepts it. On his view, merely possible individuals can be worse off than they would have been if they had existed. His argument for this view is based on two claims: (1) for one who has never existed nothing is either intrinsically good or bad, and (2) a state of affairs in which nothing is either intrinsically good or bad for one is worse for one than a state of affairs in which things are overall intrinsically good for one.

Claims (1) and (2) jointly imply that (3) it is worse for one not to have been brought into existence with a life that is overall good for one (61). A shortcoming of Persson's discussion of this argument is that he ignores the objection that a state of affairs in which nothing is either intrinsically good or bad for a being and a state of affairs in which things are overall intrinsically good for that being are *incomparable* with respect to that being's good. If this objection succeeds, then claim (2) is false and Persson's argument fails.

Persson argues in Chapter 2 that claims such as "There are (i.e. there *exist*) merely possible individuals that would have been better off had they begun to *exist*" are perfectly coherent and can be true (60). He considers the obvious objection that such claims are paradoxical because they seem to assert the *existence* of merely possible, and hence, *non-existent*, beings. His response to this objection is that such claims are paradoxical only if 'exist' has the same meaning in all of its occurrences, but that, in fact, different occurrences of 'exist' can have different meanings.

In one sense of 'exist', it is true that [there exist] merely possible beings because this follows from the clearly true claim that it is possible that some beings will begin to exist in the future. This is the sense in which I believe there to [exist] merely possible and, thus, non-existent beings. I cannot provide a philosophically adequate explication of this sense, but there are many commonsensical claims to which we can permissibly help ourselves, though we cannot accurately expound their sense philosophically (60-61).

However, this is puzzling because the true claim 'It is possible that there exists a being' does not obviously entail 'There exists a being'. Thus, the claim that there exist merely possible beings does not obviously follow from the clearly true claim that it is possible that some beings exist (or will exist). There are controversial views on which inferences of this kind are valid (see, among others, Lewis 1986, Linsky and Zalta 1994, and Williamson 2013). Unfortunately, Persson neither considers nor defends any such view.

As Persson acknowledges in chapter 3, IVB seems to imply that merely sentient beings (sentient beings that are not persons), and perhaps even persons, are *replaceable*. That an individual *X* is replaceable means that there is no more reason to benefit *X* by extending *X*'s life than there is reason to kill *X* and benefit some other

individual *Y* of the same kind as *X*, by bringing *Y* into existence with a good life, assuming that either alternative would produce the same total of benefits. In what follows, I will call the act of replacing an individual in the manner described above *replacement*, the alternative to it *non-replacement*.

Persson claims that in real-world conditions IVB does not imply that merely sentient beings are replaceable. Instead, he claims, IVB implies that replacement would typically be *worse* than non-replacement. His most intriguing argument for this claim appeals to a version of what Larry Temkin (2012) calls *the disperse additional great burdens view*, according to which it is better to inflict a very small harm on each of many individuals than to inflict a very serious harm on a single individual, even if the total amount of harm associated with the former option is *greater*.

According to Persson, the *disperse additional great burdens view* supports the claim that in most real-world conditions replacement would be worse than non-replacement. Whereas non-replacement bestows a significant benefit on a single individual, a merely sentient being, replacement is tantamount to giving each of many possible individuals a very small chance of a significant benefit. In any act of procreation there are many different sperm-egg combinations, each corresponding to a different possible individual. Hence, there are many possible individuals each with a *very small chance* of coming into existence. According to Persson, denying each of these possible individuals this very small chance of existence would inflict on each only a very small harm, whereas significantly shortening the life of a merely sentient being by killing it would inflict on it a very serious harm. Thus, given how procreation normally works, replacement would inflict a serious harm on one individual, whereas non-replacement would inflict only a very small harm on each of many individuals. And so according to *the disperse additional great burdens view*, replacement would be worse than non-replacement.

Persson neither gives an argument for the claim that a mere chance of a benefit *is* a benefit, nor for the claim that foregoing a mere chance of a benefit *is* a harm. But even if we accept these claims, Persson's argument is problematic. It ignores the fact that on IVB, replacing individual *X* with individual *Y* would not only give a very small chance of existence (and hence, let us suppose, a very small benefit) to each of

many possible individuals, but might also bestow a *very significant benefit* on *Y*.

Persson's rationale for ignoring this benefit to *Y* is as follows. Before one makes the choice between replacement and non-replacement, there is no possible individual of whom it is determinately true that it will receive the benefit of coming into existence if replacement is chosen. Therefore, on a *person-affecting* view of beneficence, there is no non-instrumental moral reason to produce this benefit. But this seems false. If replacement is *in fact* chosen, then there will *in fact* be an individual *Y* of whom it is determinately true that he is better off than he would have been if non-replacement had been chosen. One can therefore accept a person-affecting view and recognize the significance of the benefit that *Y* would have if *Y* were to exist.

Surprisingly, Persson claims that although merely sentient beings are not replaceable under real-world conditions, *persons* are replaceable in certain hypothetical situations. A person could rationally consent to being killed and replaced by an intrinsic duplicate of her, and perhaps even by someone with values sufficiently like hers. Persson's claim that persons are replaceable is inspired by Derek Parfit's well-known view that what matters in survival is not identity but psychological connectedness and/or continuity with any cause (Parfit 1984, 281-306). However, according to Persson, what matters in one's survival is not psychological continuity (with or without causation) but rather continuity of one's values or ideals.

Another challenge for IVB, discussed in Chapter 4, is that it seems to imply what Parfit called

The Repugnant Conclusion. Compared with the existence of very many people—say, ten billion—all of whom have a very high quality of life, there must be some much larger number of people whose existence, if other things are equal, would be *better*, even though these people would have lives that are barely worth living (Parfit 1986, 150).

According to Parfit, one's quality of life or life-quality is determined by how much of "the best things in life" (such as the most profound human relationships, or the most exquisite aesthetic experiences) one's life contains (1986, 150).

One argument for *The Repugnant Conclusion* involves a sequence of outcomes in which the number of individuals gets progressively larger while average life-quality gets progressively lower. Suppose outcome *A* contains ten billion individuals all with very high life-quality, while outcome *B* contains five times as many individuals (50 billion) each with a life-quality just minimally lower than that of each individual in *A*. *B* seems better than *A*. Moreover, Persson claims that his view of beneficence supports the view that *B* is better than *A* (104). Not only do the individuals in *B* vastly outnumber those in *A*, but the individuals in *A* and *B* have *benefited* from coming into existence, and the benefit that each individual in *B* enjoys is only slightly smaller than the benefit that each individual in *A* enjoys. For this reason, it seems, the total amount of benefit in *B* is greater than in *A*.

Next consider *C*, “which is related to *B* as *B* is to *A*” (104). *C* seems better than *B* for the same reason that *B* seems better than *A*. And again, Persson claims, his view of beneficence supports this result. We can continue in this way until we reach an outcome *Z* which contains an enormous population (many billions of people) in which each person’s life-quality is so low that her life is just barely worth living. Yet intuitively, and on Persson’s view, *Z* is better than *Y*, which is better than *X*, ..., which is better than *B*, which is better than *A*. If the relation *better than* is transitive, *Z* is better than *A*. But that is *The Repugnant Conclusion*!

Persson argues, against Parfit, that what makes *The Repugnant Conclusion* repugnant isn’t that the larger population with lives barely worth living is devoid of the best things in life. Instead, he argues, the root of its repugnance is embodied in the following claim:

Root: For every intensity of (positive) well-being, however low, there is an instance (or set of instances) of it of sufficient (aggregate) duration for this instance (set of instances) to be better in respect of well-being than a (minimal) instance of a higher intensity of well-being, however high it may be (105).

Persson rejects the *Repugnant Conclusion* and *Root* by denying that *better than* is transitive. His argument against transitivity is novel in that it does not appeal to what Temkin (2012) calls *the essentially comparative view of outcome goodness*—according to which an outcome’s goodness can depend not only on its internal features but

also on what other outcomes it is compared with. On Temkin's view, transitivity fails only if the essentially comparative view of outcome goodness is true. However, Persson argues that the failure of transitivity follows from the claim that value is a *supervenient property*. On Persson's view, the supervenience of value entails the non-transitivity of perfect similarity with respect to value, which entails the non-transitivity of *better than*. On Persson's understanding of the notion of supervenience, 'S is a supervenient property of X' means there are more basic or subvenient properties B of X in virtue of which X has S. This is stronger than standard definitions of 'supervenience' according to which the supervenience of A-properties on B-properties means that two worlds differ in their A-properties only if they differ in their B-properties.

Persson's argument from the supervenience of value to the non-transitivity of perfect similarity in value proceeds as follows (107-110). Suppose X, Y, and Z are three different painful (bad) experiences. Suppose S is the felt badness of an experience and B the basic physical properties in virtue of which an experience has S. There can be differences in respect of B between X and Y, and between Y and Z, which are not sufficient for differences in respect of S between X and Y, or between Y and Z, but that *are* sufficient for differences in respect of S between X and Z. As a result, it is logically possible that X is perfectly similar to Y, and Y to Z, with respect to S, and yet, X and Z differ with respect to S.

For example, the difference between two patterns of the firing of C-fibers in a subject's brain may be so small that the subject cannot feel any difference between the corresponding pains that these two patterns produce; yet if the difference between two such patterns were sufficiently large, the subject *would* feel that one was more painful than the other. We can therefore imagine that while the differences in such patterns between X and Y, and between Y and Z are not large enough to produce a difference in felt painfulness, the difference between X and Z *is* large enough. Thus, it is possible that X and Y, as well as Y and Z, are perfectly similar with respect to the felt badness of pain, while X and Z are *not*.

Persson shows how his putative counterexample to the transitivity of perfect similarity in the (dis)value of pain can be extended to produce putative counterexamples to the transitivity of *better than*. Suppose Y's *duration* is noticeably slightly shorter than X's, while there is no

difference in felt *intensity* between *X* and *Y*. Assuming intensity and duration are the only relevant bad-making features for pain, it seems *Y* is *better than X*. Next, suppose *Z* is noticeably slightly shorter than *Y* but there is no difference in felt intensity between *Y* and *Z*. Again, it seems *Z* is *better than Y*. Still, *Z* may *not* be better than *X* “because *Z* is felt to be more intense than *X*, and this difference is judged to outweigh the longer duration of *X*” (114).

By rejecting transitivity, Persson blocks sequential arguments for *The Repugnant Conclusion* and provides a defense of the *disperse additional great burdens view* to which he appeals in Chapter 3. As Persson acknowledges, the rejection of transitivity implies that value cannot be quantified.

Part II of *Inclusive Ethics* develops and defends the second main part of Persson’s inclusive moral philosophy—an egalitarian principle of justice. Chapter 7 defends *Extreme Egalitarianism* (EE) according to which justice requires that everyone capable of being well or badly off is *equally* well off unless some autonomously choose (or would choose) to be worse off (163). According to Persson, if some autonomously chooses to be worse off, the resulting inequalities in welfare are neither just nor unjust. Justice simply fails to apply to them. EE is extreme in that it rules out the possibility of *just* inequality. In defense of this, Persson argues that there is nothing that can *make* inequalities in welfare just.

His argument is as follows:

- (1) Nothing can make it just that some are better off than others, unless this is something for which they are responsible.
- (2) There is nothing for which anyone is *ultimately responsible*.

From (2), we get

- (3) There is nothing for which anyone is ultimately responsible that can make it just that some are better off than others.
- (4) If there is nothing for which anyone is ultimately responsible that can make it just that some are better off than others, then there is nothing for which anyone is *responsible* that can make it just that some are better off than others.

From (3) and (4), we get

- (5) There is nothing for which anyone is responsible that can make it just that some are better off than others.

Finally, from (1) and (5), we get

- (6) There is nothing that can make it just that some are better off than others.

Persson calls this the *Demolition Argument* against just inequality (158) since it demolishes what he takes to be the only plausible candidate explanations of how inequality can be just. One such candidate is that some people deserve to be better off (or worse off) than others. Another is that some people have a right to be better off (or worse off) than others. According to Persson, such claims are true only if those who are better off (or worse off) are ultimately responsible for certain of their actions. But according to the *Demolition Argument's* premise (2), there is *nothing* for which anyone is ultimately responsible.

Persson's argument for premise (2), an argument against the possibility of ultimate responsibility, resembles familiar arguments in the literature on free will (see, for example, Strawson 1994). One's being ultimately responsible for a certain fact *F* requires that the fact that one is responsible for *F* obtains only in virtue of facts for which one is also responsible. But as Persson convincingly argues, when one is responsible for a certain fact *F*, the facts in virtue of which one is responsible for *F* always seem to include facts for which one is *not* responsible, for example the fact that one was born with certain genes, or that the atoms in one's brain have a certain configuration.

Persson's intuition is that one's lack of ultimate responsibility undercuts the significance of one's responsibility as it pertains to matters of justice (160). I think that many (including me) share Persson's intuition, but I also suspect that many would have the opposite intuition, namely that one's being responsible for *F* can be relevant to matters of justice even if one is not *ultimately* responsible for *F*. Folk intuitions appear to be split on the question of whether physical determinism would undermine free will (see, for instance, Nahmias et al. 2006). Since this question seems closely related to the question of whether lack of ultimate responsibility (in Persson's sense) would undermine responsibility *tout court*, one might expect folk intuitions to be similarly split on the latter question.

Part III discusses the implementation of Persson's moral theory—the inclusive principle of beneficence and the egalitarian principle of justice. Chapter 11 argues that given the current state of the world, the morality that follows from Persson's theory is of little or no practical use, and that therefore morality (currently) has no point. Persson assumes that for morality to have a point it must be implementable. One putative necessary condition for a principle's being implementable is that it must

be practically possible for people in general to follow it (234). Persson thinks his two principles probably do not satisfy this condition. These principles conflict starkly with common-sense morality (CSM), and so wouldn't be generally accepted. Persson concedes that his morality is very demanding especially for well-off members of affluent nations. His morality includes nothing that permit these well-off people to retain their enormous wealth rather than use it to help the world's poor, which, if done effectively, would not only greatly increase the sum of individuals' benefits but also greatly reduce unjust inequality. Persson concludes that unless the current political and economic climate, or the moral psychology of human beings changes dramatically, morality will have no point.

According to Persson, if morality has no point, *moral philosophy* could still have a point by engendering rational consensus regarding the truth of a moral theory. Chapter 12 argues that the prospects for such a consensus are dim. According to Persson, moral philosophy is likely to remain inconclusive because philosophical "disagreements—such as between consequentialism and deontology—are too deep to be settled by arguments acceptable to all parties" (234). Contemplating the history of analytic moral philosophy, Persson "cannot think of a single instance of [its] having definitely resolved a pre-theoretical disagreement about what is morally right or wrong" (236).

But perhaps *a* point of moral philosophy is to steer morally motivated individuals in a better direction, allowing them to have a more positive impact. Moral philosophy *seems* to have already accomplished this to some extent. For example, the growing movement known as *effective altruism* has been influenced by philosophical pioneers like Peter Singer. Key figures in this movement, such as William MacAskill, report being inspired by the philosophical arguments of Peter Singer (Rogan 2017). This point of moral philosophy is admittedly less ambitious than what the method of philosophical argument seems designed to do. Nevertheless, it seems extremely important.

Chapter 13 presents *a philosophy of life*—a set of recommendations for coping with general features of life such as the dominant role of luck in shaping one's prospects and the fact of one's inevitable death. A philosophy of life focuses mainly on the *intrapersonal* dimension of morality, which is concerned only with how one's acts affect one. Persson's philosophy of life emphasizes *autarchy*—"the attitude of striving to change our own affective and conative attitudes so as to

make them conform to the world around us” (250). Autarchy is contrasted with what Persson calls *heterarchy*—the attitude of striving to make one’s environment conform to one’s desires.

Persson’s autarchic philosophy is inspired partly by the writings of the Stoic Epictetus (2004), although Persson’s philosophy is much more reasonable than Epictetus’s. Persson rejects Epictetus’s absurd advice that we stop trying to avoid disease and famine because they are beyond our control. He also offers interesting advice that one won’t find in Epictetus. For example, Persson suggests increasing one’s level of contentment with one’s life by contemplating how unlikely it was that one came to exist in the first place. He also suggests contemplating the view, considered in Chapter 3, that what matters in one’s survival is not one’s identity but that one’s values are carried forward.

Persson emphasizes a connection between his philosophy of life and the problems concerning the implementation of his morality (260-261). A possible result of becoming more autarchic is that one finds morality less demanding. Much of morality’s demandingness comes from a conflict between it and desires to accumulate wealth and social status. By becoming more autarchic, one’s desire for these things diminishes, making it easier to act morally.

Inclusive Ethics is incredibly ambitious. It seeks to develop not only a general framework for evaluating behavior that impacts individuals that are or could become conscious, but also a general philosophy of life, an account of the purpose of morality, and an account of the purpose of doing moral philosophy. The book is filled with novel and intriguing arguments, many of which have not been touched upon here. However, some of the book’s central claims, such as the claim that non-existent individuals can be worse (or better) off than if they had existed, are vulnerable to crucial philosophical objections that Persson does not consider. Though less carefully argued than one might hope, the book is interesting, remarkably innovative, and overall well worth reading.

REFERENCES

- Epictetus. 2004. *Enchiridion*. Translated by George Long. Mineola: Dover Publications.
- Lewis, David. 1986. *On the Plurality of Worlds*. New York: Oxford University Press
- Linsky, Bernard and Edward N. Zalta. 1994. “In Defense of the Simplest Quantified Modal Logic.” *Philosophical Perspectives* 8: 431-458.
- McMahan, Jeff. 2013. “Causing People to Exist and Saving People’s Lives.” *Journal of Ethics* 17: 5-35.

- Nahmias, Eddy, Stephen G. Morris, Thomas Nadelhoffer, and Jason Turner. 2006. "Is Incompatibilism Intuitive?" *Philosophy and Phenomenological Research* 73: 28-53.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, Derek. 1986. "Overpopulation and the Quality of Life." In *Applied Ethics*, edited by Peter Singer, 145-164. Oxford: Oxford University Press.
- Rogan, Joe. 2017. "The Joe Rogan Experience Podcast 930: Interview with William MacAskill." Podcast audio. *The Joe Rogan Experience*. Accessed November 24, 2017. <https://www.youtube.com/watch?v=buyBzK5yM-s>.
- Strawson, Galen. 1994. "The Impossibility of Moral Responsibility." *Philosophical Studies* 75: 5-24.
- Temkin, Larry. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. New York: Oxford University Press.
- Williamson, Timothy. 2013. *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Tim Campbell is a postdoctoral researcher at the Institute for Futures Studies, Stockholm. He works on population ethics, the metaphysics of personal identity, and the implications of the latter for the former. Contact e-mail: <timothy.campbell@iffs.se>