Laitinen, Arto, and Otto Sahlgren. "Incommensurability, Incomparability, and Robotic Choices." In *Social Robots in Social Institutions: Proceedings of Robophilosophy*, pp. 389-398. IOS Press, 2022.

Incommensurability, Incomparability, and Robotic Choices

Arto LAITINEN[a,1] and Otto SAHLGREN[a]
[a] *Tampere University, Tampere, Finland*

Abstract. Some comparisons yield puzzling results. In the puzzling cases, neither item is determinately better than the other, but they are not exactly equal either, as improving one of them just slightly still does not make it determinately better than the other. What does this kind of incommensurability or incomparability mean for robots? We discuss especially Ruth Chang's views, arguing for four claims. First, we defend her view that despite appearances to the contrary, formal incomparability does not follow – comparison of "apples" and "oranges" is not impossible. Second, rational value-assessment turns out to be very complicated in virtue of the non-linear relations between descriptive and evaluative features. These complications pose considerable challenges to robots, whatever views about incommensurability are adopted. Ruth Chang's theory introduces a fourth value relation "being on a par", and we argue, thirdly, that (unlike its rivals) it will pose a considerable extra challenge for robots, as it is a non-transitive relation (unlike equality or betterness). Fourthly, we argue that exercise of normative powers - Chang's suggestion for hard choices in contexts of parity - is not available in the case of (fully autonomous) robots.

Keywords. Philosophy, incommensurability, incomparability, parity, value, robotic decision-making

1. Introduction

How robots are programmed to behave may affect humans and societies greatly. Their designs and functionings inevitably reflect certain values (at least implicitly), and the premise of ethical design is that they ought to be guided by what is desirable or valuable. But what if values are incommensurable and, by extension, the value-bearers between which they ought to choose incomparable, as in the proverbial comparison of "apples and oranges"? Indeed, comparisons can yield puzzling results. In puzzling cases, neither item is determinately better than the other, but they are not exactly equal either, as improving one of them just slightly still does not make it determinately better than the other [1, 16]. Think of ranking Mozart and Michelangelo in terms of creativity [5] or excellent poets and novelists in terms of literary merits [18]. Furthermore, existential choices often involve such comparisons. Consider the choice between a career as a lawyer or a career as a clarinetist [20], or the choice faced by Sartre's young man between joining the Resistance or staying with one's ailing mother, for example [21]. If faced with such "hard choices" [9], robots would have to be sensitive to the values at play, including their comparative relations. From delivery robots to recruitment chatbots and

1 Corresponding Author, Arto Laitinen, Philosophy, Faculty of Social Science, Pinni B 4148, 33014 Tampere University, Finland; E-mail: arto.laitinen@tuni.fi

possible autonomous weapon systems, any robots that make “independent” choices, and compare or rank items, may face these kinds of “puzzling situations” and “hard choices” that invoke questions of incommensurability or seeming incomparability.

In this paper, we look at both robotic evaluations and choices from the perspective of incommensurability and incomparability [1, 16]. Regarding evaluation, comparison, and ranking, on the one hand, we defend an optimistic view concerning comparability: often proverbial “apples” and “oranges” *can* be compared rationally. Drawing on the work of Ruth Chang [5-9], we propose that algorithmic comparisons are possible whenever the items can be formally comparable in light of a covering value. The comparison or ranking may need to allow for the special outcome called “parity” or “rough equality” of the formally comparable items. We focus on Chang’s [5-9] account, which differs from rival theories [2,3,15,18-20] in defending the need for a fourth comparative relation, “parity” We discuss what this implies for robots, the design of which is typically implicitly premised on the notion that tracking descriptive features suffices for tracking evaluative features, and by extension, value relations. We argue this is not the case, for many reasons.

In the context of robotic choice, on the other hand, we consider possible situations where two or more items between which a machine ought to choose are “on a par”. We discuss Chang’s proposal according to which normative commitments can create willbased reasons that function as tie-breakers for “hard choices”. We argue this proposal is not applicable in the context of robotic choices: robots lack normative powers necessary to create will-based reasons for hard choices, and in case the robots are functionally “autonomous”, it is uncertain whether designers and end-users exercising their powers will do the trick.

In Section 2 we defend – with Ruth Chang - the view that the hard choices do not imply formal incomparability (or noncomparability).

In Section 3, we draw attention to the more general phenomenon that comparing distances in descriptive features may not be a reliable way of assessing evaluative differences, which suggests that for robots or ML systems to realistically track value in their choices and recommendations, evaluative input is needed. This is so regardless of whether incommensurability poses an additional problem.

In Section 4, we zoom in on Ruth Chang’s understanding of the puzzle cases and her theoretical innovation of Parity. We discuss the radical departure that “parity” would require from existing ranking technology. In section 5 we argue that as robots do not have practical authority or possess normative powers, Chang’s suggested solution will not do (nor are any humans well-placed for the exercise of normative powers that would solve or pre-empt puzzle cases that robots may run into).

Section 6 is a Concluding Reflection. Before we proceed, a note on terminology: The terms “incommensurability” and “incomparability” are used with slightly different meanings in the general debate [1, 16]. The seeming failure of comparison to yield any positive result in terms of betterness, worseness, or equality is called “incommensurability” by some and “incomparability” by others (and some use these terms interchangeably). We will not follow either of these usages. We will use the term “incommensurability” to refer to two things lacking a common cardinal measure (i.e., cardinal measuring is not possible) and “formal incomparability” or “noncomparability” to refer to two items lacking a covering value with respect to which they can be compared (e.g., the number 54 cannot be compared to the color green in terms of tastiness). We call the cases where two items are formally comparable, but neither is better, and yet they are not exactly equal, “puzzle cases”.

2. Does formal incomparability or noncomparability follow?

Comparison of two items in terms of some property or feature F can take the form of the question: which is *F*er (i.e. which is more F)? [2]. When F is a feature that both items can have, the items are *formally comparable* in light of F, which is what Ruth Chang calls a “covering consideration”. In some cases, for example in comparing the weight of two items, things are determinate: either A is weightier, or B is weightier, or they are equally weighty.

The Trichotomy Thesis of value relations states that one of three relations – worseness, betterness, or equality – holds between any two items that are comparable. Note that robotics, machine learning (ML), and artificial intelligence (AI) more generally, are arguably premised on this view: Provided that the relevant concepts (*F*-ness and its components) are clearly operationalized in the language of mathematics and that the available data is accurate, there is no principled objection to computer systems being able to determine which of the three relations hold between items in terms of the relevant covering values.

Debates on (in)commensurability and (in)comparability revolve around cases where it seems that neither is better, but they do not seem to be exactly equal either, as minor changes in one do not seem to alter the situation. Paradigmatic examples include those considered above (e.g., comparing Mozart and Michelangelo, “hard” or existential choices between careers), although seemingly mundane choices (e.g., between different delivery routes and different films) can be construed similarly (with more realistic promise to be made by robots).

Note also that the covering consideration can be “multidimensional”: In recruitment, for example, there can be many different kinds of merits or good-making features that make someone a better choice for a given task, and the comparison must be able to combine these different dimensions. In easy cases, the merits in different dimensions may be summed up, but harder cases may lack a common cardinal measure in terms of which the aggregation is to be made. Sometimes the multidimensional value may *lack a conventional name*, but can clearly be defined with the help of the context, or via its constitutive elements (i.e., “hiring-worthiness” or “merits-in-X-and-Y”). Some comparisons may further be complicated in concerning items from *different* “*genres*” (e.g., comparing Poets and Novelists for a literature prize).

Because comparing excellence in different genres is difficult, and because it may require assessing many dimensions which lack a common cardinal measure, and because the value in terms of which the comparison is made may lack a conventional name, it is very tempting to conclude that any comparison is impossible or senseless – we shouldn’t try to compare apples and oranges. There is reason to think, however, that the comparison is “merely” difficult but not in principle impossible.

We maintain that cases of formal incomparability are rather rare. While the number 54 is noncomparable (or formally incomparable) with the color green in terms of tastiness, there typically is a covering consideration in the contexts of choosing delivery routes, recruiting people, grading essays, or choosing careers, and other paradigmatic cases discussed in the literature. The covering consideration is perhaps non-cardinal, nameless, or multi-dimensional, but nonetheless allows for formal comparability. Comparative judgments within a single "genre" can be easier (e.g., comparing sculptors to other sculptors). Yet while cross-genre comparisons are harder (e.g., comparing sculptors to composers), they are also possible when there is a relevant covering consideration.

Ruth Chang's [5-9] master argument to show that Mozart and Michelangelo are formally comparable in terms of creativity as an artist, appeals to an imagined very bad composer, Talentlessi. Talentlessi is not nearly as good a composer as Mozart. It seems clear that we can imagine a composer who is less creative than Michelangelo, and indeed we are invited to think that Talentlessi is one. So, despite appearances, we *can* compare Michelangelo with composers, and come to the conclusion that Michelangelo is more creative than many composers. The difficulty in telling whether Mozart or Michelangelo is more creative is precisely that, a difficulty. It is not a formally senseless question. We just do not know how to answer, as neither seems better. Chang continues her argument to suggest that they are "on a par" (see Section 4 below), but the argument so far has shown that Mozart and Michelangelo are at least comparable in terms of creativity.

John Broome [2] similarly approaches the question via a "standard configuration", where there is on the one hand a continuum (without gaps) from something very F to something not F at all, and then outside that continuum a point of comparison that Broome calls a "standard". For example, in comparing which is more impressive, Stonehenge or some church, and we can imagine a continuum of actual and possible churches ordered in terms of how impressive they are. Some are more impressive than Stonehenge, some are less impressive, and then there is a zone of indeterminacy where it is indeterminate which is more impressive, the church or Stonehenge.

The main thing so far is the good news: the difficult cases of comparing "apples and oranges" turn out to be comparable, at least in principle. There is no reason to expect a full paralysis when a sophisticated robot meets a hard choice, or a difficult comparison. But things will not be very straightforward either.

3. Evaluative Matters are Complicated Anyway

Machines perform comparisons based on data. The value of ML systems, for example, lies in part in their capacity to discover patterns in the available data, to detect subtle (dis)similarities between data items, and to help decision-makers make informed decisions, accordingly. In tracking which of two items is redder than the other, the differences concern descriptive, non-evaluative features. What about when the data is supposed to reveal something about evaluative features and value relations? Can a system that tracks descriptive features (in the input) manage to track evaluative features and value relations (in the output)? We suggest that, even if robotic comparisons and rankings are possible, comparing distances in descriptive features may not be a reliable way of assessing evaluative differences. This is so although evaluative features always supervene on descriptive features: making an exact copy of descriptive features of an item (and relevant features of the situation in cases in which they matter) produces an exact copy of the evaluative features of the item. But the distance in descriptive differences need not be matched by the distance in evaluative differences, and this is why descriptive differences represented in robots' input data as such are unreliable as guides to evaluative differences.

We will next illustrate some ways in which this is so, merely to pre-empt a wrong kind of response to the putative cases of incommensurability. The wrong kind of response is to think that "things would be neat and straightforward without incommensurability, so I better find a way to get rid of it". Evaluative matters, we suggest, are complicated anyway; incommensurability does not introduce messiness or unruliness to an otherwise neat picture of evaluative differences linearly matching descriptive difference – such a neat picture is illusory.

First, consider *diminishing marginal utility*, the familiar predicament in which some item's utility value does not increase in proportion to its quantity. One car maybe highly useful for a person in terms of allowing them to drive, but the tenth car will not make much of a difference. How steep the utility curve will be is a matter of value judgments, but in many cases it will not rise linearly. While robots may well be programmed to be sensitive to this phenomenon, the point is that descriptive difference is not directly proportional to evaluative difference.

Second, *holism of value* refers to the idea that the value of a whole does not consist in the sum of the value of the elements. A machine requires all of its parts to function, and missing just one part might make the non-functioning assemblage worthless as a machine (even if each spare part retains their value individually). Some wholes may have positive value (appropriate punishment for a crime) even though its elements have negative value (the crime; restrictions on liberty of movement that count as punishment) [17].

Thirdly, and perhaps more controversially, some items may change their value depending on the context: lying can be good (of positive value) in the context of a game of Contraband, although elsewhere lying is bad (of negative value). [10].

A fourth way in which distances in value-properties may not follow the distances in descriptive properties relates to *"clumpiness" of range-properties*. A range property is something that you can have fully once you reach a threshold (say, the full right to selfdetermination by sufficiently capable adults), but before the threshold you either have it in degrees (say, the partial right to self-determination by the youth that depends on the level of capacity of self-determination) or lack altogether (say, the political rights of the underage citizens) [22]. Above the threshold one may become more and more capable, but no evaluative difference matches this descriptive difference. Similarly, one is equally and fully in Sweden whether located in Göteborg or Stockholm. This kind of "clumpiness" is at play, for example, in evaluating student essays as A's, B's, C's, and so on [15].

Finally, let us mention one more unruliness concerning descriptive and evaluative features, which derives from genre- or kind-dependence of value in some cases. It can be illustrated by the fact that a big flea is a small animal whereas a small elephant is a big animal, and a small elephant is bigger than a big flea. (For classic discussion of this, see [13], and in the context of incommensurability, see [4]).

The abovementioned phenomena suggest that the assumption that evaluative differences track descriptive ones in a linear fashion is an oversimplification. For robots or ML systems to realistically track values in their actions, choices, and recommendations, evaluative input is needed in the training data, accordingly, and evaluative understanding needs to be encoded in the programs or learnt models. This is true regardless of whether incommensurability poses an additional problem or not (and also illustrates one sense in which it matters *whose* evaluations robots are programmed to live by).

4. A Fourth Value Relation: Parity?

Let us now consider whether there is a further complication in comparative value relations, namely, whether the puzzle cases can or should be handled via theoretically introducing a fourth value relation. According to the standard view, "1) differences between cardinally ranked items can always be measured by a scale of units of the relevant value, and 2) all rankings proceed in terms of the *trichotomy* of 'better than', 'worse than', and 'equally good'" [8, see also 5,6,7,9]. Note that puzzling cases of comparison have in common their apparent tendency to challenge the Trichotomy Thesis of value relations, especially the second premise. Chang [6] provides a compelling argument for extending the traditional view of value relations. Consider again the puzzling case of comparing Michelangelo and Mozart in terms of creativity, for example. Chang asks us to imagine the following situation:

1) Michelangelo is neither better nor worse than Mozart in terms of creativity
2) A slightly "sweetened" version of Michelangelo, call them Michelangelo+, is better than Michelangelo in terms of creativity
3) Michelangelo+ is not better than Mozart in terms of creativity.

Chang argues that Michelangelo and Mozart are not related by any of the three relations included in the Trichotomy. A slightly improved version of Michelangelo is quite intuitively better than Michelangelo himself, but the slight improvement does not make him better than Mozart. The sculptors and the composer are "on a par" as Chang puts it. To explain what makes puzzling cases puzzling, Chang proposes "parity" as a fourth value relation in which certain items may stand. This proposal leaves the existing trichotomy as is, but merely extends the scope of value relations. The fourth relation is at stake, when neither is better, and a minor improvement does not suffice to make one better, but a major improvement would suffice.

Suppose Chang is right. What would "parity" mean for comparisons and choice situations in the context of robots or ML systems? One promise would be that once the robots or ML systems would be able to recognize cases of "parity", they would avoid the issue of illusory precision mentioned above: some differences in the "extra points" that a sweetened option might have would not make a difference.

Whether such sensitivity to relations of parity could be practically achieved is another question, however, due to the possibly grave issue that "parity" is not a transitive relation. Michelangelo may be on a par with Mozart, who may be on a par with Michelangelo+, but Michelangelo+ is better than Michelangelo (and not on a par). For a ML system to learn to predict who is (according to the best human judges) on a par with whom, for example, it would have to be trained to be sensitive to the possibility of pairs who are "on a par". This would require that it be possible to program the system to (learn to) recognize what features would increase or decrease the value of items within one "genre" (e.g., what descriptive differences would explain the difference between Michelangelo+ and Michelangelo) and within another "genre" of items, respectively (e.g., Mozart and Mozart+). Moreover, to compute parity-sensitive rankings, it would need to be programmed to (learn to) recognize what kind of differences lead (best) human evaluators to judge that two items from different "genres" are on a par. Importantly, this need not be a straightforward result of the descriptive differences, thanks to the complicated relationship between descriptive and evaluative features (but a result of the descriptive differences nonetheless).

To our understanding, this would call for novel kinds of datasets about pairs of items that relate to each other in the four different ways that Chang envisages. Exact equality and betterness would result from individual evaluative features (in transitive ways) just like "being-taller-than" results from the tallness of individuals. Intransitive parity, however, would have to be presented in the training data as a feature of *pairs of items* making the dataset hugely complex (just like variables "A loves B", which does not result from features of A and B but is irreducibly relational). So even if evaluative matters are complicated anyway, introducing non-transitive parity to evaluative comparative relations, would add to the complexity considerably, and would be more than a drop in the ocean. Of course, Chang may turn out to be right, but that would be bad news for robots and ML systems aiming to track evaluative comparative relations.

Further, if parity would obtain between two items in a set of ranked items, this could suggest that the generated ranking should be partially ordered as opposed totally ordered [12]. However, due to intransitivity, a partially ordered ranking might obfuscate the minor differences between items *within* a given rank: one might have exactly equal items and items that are "on a par" within a given rank.

5. **Robots do not have "normative powers"**

Let us comment one more aspect of Chang's theory. Ruth Chang [9] suggests that in cases of "hard choices" between incommensurable options it is rational to commit oneself, to create a reason by an act of will. For Chang, "hard choices" are cases where given reasons (i.e., reasons grounded in normative facts) do not suffice to break the tie between options that are "on a par", but normative commitments or decisions can create "will-based reasons" that do so. This is analogous to how a meeting can collectively make a binding, unique decision, or how autonomous persons can make promises, contracts, and so on. What matters is that a normative move is made by people having the suitable authority, and the output consists of extra reasons or duties.

Goodman [14] has applied this view to algorithmic fairness, where he notes that the question regarding what notions of fairness algorithms should reflect is not a decision to be left to machines, but one resolved through human agency alone. This is indeed a relevant consideration: A normal adult person can make promises, contracts, adopt autonomous policies, or commitments, and so create duties to oneself. This involves "bootstrapping" an extra reason into existence (most easily seen in the case of promises and contracts), and after the extra reason is taken into account, the normative tie is resolved. Robots themselves do not have such normative powers or personal authority. They may end up doing one thing rather than another, but they do not have a will to exercise or personal autonomy of the sort that should be respected. Robot choices do not generate reasons in the way that human promises, contracts, or commitments generate reasons.

Indeed, it seems that to the extent that the solution to robot choices should come from exercise of normative powers, it should be from exercise of human normative powers. But which humans? Perhaps the designers or users of robots?

Start by considering the designers and suppose they *can* build a robot that reflects *their* commitments. We suggest the key issue remains unresolved, even if the robot is in this way designed to address "hard choices" between items P and R (which are on a par) by always choosing P over R. This would imply, first, that the robot does not actually recognize the choice as “hard”. If P is better than R in light of the programmed settings, there is no “hard choice” to be made between P and R.

Second, and more importantly, the robot might yet face other “hard choices” precisely due to the “programmed commitment” creating new indeterminate cases (e.g., the robot might need to choose between P and a "sweetened" R+). In other words, prior commitments precisely fix the “standards” and the corresponding “zones of indeterminacy” within which “hard choices” arise. It is characteristic to “hard choices” that they are *token* choices that arise in the flux of normative life.

Third, to program a commitment surely constitutes an exercise of power similar to when any relevant collective sets a policy. It is not necessarily an exercise of *normative* power in the technical sense, however. Specifically, it is unclear whether the designers are normatively positioned so that they *have* the relevant “normative powers” or can exercise them in advance – even in principle. Compare this to being a chairperson in a meeting: only the chairperson has the relevant ability to run the meeting, others simply lack the ability. Similarly, no-one can in advance *make* the unique decisions made collectively in the meeting; that is part of what normative powers are like. Surely, robots can be programmed to not recognize the puzzle cases of incommensurability, or to “value” some choices over others. In these cases, the designers use some kind of power but not exactly their “normative powers” in the sense that promises, contracts *etc* are exercises of normative powers. One can design faulty voting machines, for instance, and that is an exercise of power that pre-empts the citizens’ use of their normative powers, but it is not an exercise of “normative power”.

By contrast, users may meaningfully exercise *their* normative powers by technological means, but again robots do not seem to fit the bill. To the extent that robots are “autonomous” (in the engineers’ sense of not being remote controlled but selfdirected) in relation to the users, and make choices, then what is at stake are not choices of the users. So, while humans can exercise their normative powers in a technologically mediated way with the help of say mobile phones and computers, “autonomous robots” seem independent of their users. There are different ways in which robots could be programmed to align with one’s values as a user in advance (e.g., through customizable settings), but even those values could then face situations of incommensurability. (To the extent that the robots would affect the lives of others, such as in the case of autonomous vehicles, perhaps they should have a say as well – in general, and not only in the puzzle cases.) For a token hard choice in a token puzzle case to be made by a user would require real-time remote-control of the robot (e.g., a delivery robot could be paralyzed and phone the user to make the decision). This is of course one option (of the humans-in-the-loop kind), but restricts the scope in which we might use "fully autonomous" robots.

We remain skeptical about the applicability of Chang’s view in the context of robotic choices. There are, moreover, reasons to criticize the way Chang links normative powers and parity. First, Chang considers normative powers only in cases where different options are on a par, but presumably normative powers can be exercised also in other situations: meetings can make collective decisions in clear-cut cases, individuals may promise to do things they would in any case have sufficient reason to do, or promise to do things they otherwise would have sufficient reason not to do, etc. Second, exercises of normative powers may create new puzzling situations downstream: what if I need to choose between two promises, that I cannot fulfil, and neither of the promise is weightier than the other, but they are not equally weighty either. Or supposing that I have made a commitment in a situation where options P and R were on a par, deciding to favour P in situations like this, also in the future. This may create a new indeterminate case, where I have to choose between P and R+.

6.Concluding Reflections

The paper has first of all presented some good news about comparability in general, suggesting that at least the puzzle cases are in principle formally comparable.

Secondly, it has suggested that the relations between descriptive and evaluative features are complicated even if incommensurability can somehow be explained away. Robots or AI systems aiming to track evaluative features will have to be pretty sophisticated “anyway”. So would accepting Ruth Chang’s proposal of fourth positive relation, “parity” be more than just a drop in the ocean? We have suggested that *if* the alternative theories [e.g. 2-3, 15, 19] retain the *transitivity* of the comparative value relations, they are technologically much easier for robot operations. Each of the rival theories may of course have their downsides that we haven’t addressed here, but at least they may avoid that specific challenge faced by Chang. Yet, they may be able to share the insight that excessive precision in evaluative comparisons is often imaginary – in fact it may be indeterminate which of two options is better, so robots or AI systems guided by very precise calculations may be committing an error.

Thirdly, the paper has argued against Ruth Chang’s specific suggestion that exercises of normative powers and creation of will-based reasons would provide a way out of the puzzle cases. More specifically, that avenue does not seem open for *robots* who do not possess normative powers, and for different reasons, not for designers or users either. But more generally, creating more reasons may just work to displace the puzzle cases – new puzzle cases may be found downstream from the creation of willbased reasons.

Could the solution be simpler then – that of simply treating all of the options in a hard choice as eligible, as “good enough” choices? The choice between them would not create new reasons, but it would solve the dilemma or threatening paralysis by simply choosing one of the options and marching on. If so, then the practical importance of recognizing zones of indeterminacy would be that one would avoid excessive precision, recognize that the options that are roughly equal as roughly equal (instead of one being slightly better than the other) – and that the robot can “flip a coin” and move on?

That option is worth a closer look. It is good to remember though, that there is hope for a robot or AI system to track value reliably only once it manages to take the sorts of complications mentioned in Section 3 into account. Only then is there hope that the puzzle cases can be recognized as puzzle cases. So even if the puzzles cases would turn out to be happy choices between eligible options, there is a lot of work to do to enable robots to be sensitive to the very existence of puzzle cases in the first place.

References

[1] Andersson H, Herlitz A, editors. Value Incommensurability. Ethics, Risk, and Decision-Making. New York: Routledge; 2022.

[2] Broome J. Is Incommensurability Vagueness?. In: Chang R, editor. Incommensurability, Incomparability, and Practical Reason. Harvard University Press; 1997.
[3] Broome J. Are Intentions Reasons? And How Should We Cope with Incommensurable Values?. In: Morris C, Ripstein A, editors. Practical Rationality and Preference: Essays for David Gauthier. Cambridge University Press; 2001.
[4] Bykvist K. Cross-Categorical Value Comparisons. In Andersson H, Herlitz A, editors. Value Incommensurability. Ethics, Risk, and Decision-Making. New York: Routledge; 2022:162-181.
[5] Chang R. Introduction. Incommensurability, Incomparability, and Practical Reason. In: Chang R, editor. Cambridge: Harvard University Press; 1997.
[6] Chang R. The Possibility of Parity. Ethics. 2002, Jul;112(4):659–688.
[7] Chang R. Parity, Interval Value, and Choice. Ethics. 2005, Jan;115(2):331–350.
[8] Chang R. Parity, Imprecise Comparability and the Repugnant Conclusion. Theoria. 2016, May;82(2):182214.
[9] Chang R. Hard choices. Journal of the American Philosophical Association. 2017, May;3(1):1-21.
[10] Dancy J. Moral Reasons. Oxford: Blackwell; 1993.
[11] Dobbe R., Gilbert TK, Mintz Y. Hard choices in artificial intelligence. Artificial Intelligence. 2021;300.
[12] Fleisher W. 2021. What's Fair about Individual Fairness?. In: AIES '21: Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society; 2021 May 19-21; USA: ACM; p. 480–490.
[13] Geach PT. Good and Evil. Analysis. 1956, Dec;17(2):33-42.
[14] Goodman B. Hard Choices and Hard Limits for Artificial Intelligence. arXiv:2105.07852 [Preprint]. 2021. [cited 2016 Jun 1].
[15] Hsieh NH. Equality, clumpiness and incomparability. Utilitas. 2005, Jun;17(2):180-204.
[16] Hsieh NH, Andersson H. Incommensurable Values. In: Zalta EN, editor. The Stanford Encyclopedia of Philosophy [Internet]. 2021 Jul [cited 2022 May 10]. Available from: https://plato.stanford.edu/archives/fall2021/entries/value-incommensurable/.
[17] Moore GE. Principia Ethica. Cambridge: Cambridge University Press; 1903.
[18] Parfit D. Reasons and Persons. Oxford: Oxford University Press; 1987.
[19] Regan D. Value, Comparability, and Choice. In: Chang R, editor. Incommensurability, Incomparability and Practical Reason. Cambridge: Harvard University Press; 1997 [20] Raz J. Morality of Freedom. Oxford: Oxford University Press; 1986.
[21] Sartre JP. Existentialism Is a Humanism. World Publishing Company; 1956.
[22] Rawls J. A Theory of Justice. Cambridge, MA: Harvard University Press; 1971.