

Author's name: Arvi Pakaslahti

Title of the paper: Act-utilitarianism, Relationships and Cheating

Institutional affiliation: University of Turku

E-mail address: arvi.pakaslahti@utu.fi

Telephone number: +358-2-333 6062

Mailing address: Sirkkalankatu 36 B 44, 20700 Turku, Finland

Act-utilitarianism, Relationships and Cheating

1 Introduction

According to one form of act-utilitarianism (call it AU1 and call act-utilitarianism in general AU), an act is morally right, if and only if its agent can justifiably believe (from an intellectual point of view) it to maximize expected utility (or well-being or welfare) in the world.¹ This means that when the moral agent does not act in accordance with this criterion of moral rightness of AU1, she does not comply with AU1.

AU1 can also be understood to have a non-moral criterion of goodness, which is that the more utility the world contains, the better the world is. This means that the best world in terms of AU1 is the world which contains most utility.

In this paper my main concern is with interaction problems regarding *perfectly complying AU1-believers*, by which I refer to moral agents who believe in AU1 (i.e. who think that it is the most plausible, best justified, right or true moral doctrine) and who never do anything which is morally wrong according to AU1. More specifically, I am concerned with the question whether it is reasonable to believe that an interaction of two perfectly complying AU1-believers, who are known to be perfectly complying AU1-believers (and who know that they are known to be perfectly complying AU1-believers), would be likely to cause such problems in terms of AU1 that it would be likely to be better in terms of AU1 if they were not known to be perfectly complying AU1-believers and if they were known to be believers and perfect compliers of a certain

¹ AU1 is one form of *subjective* AU. Of course, in some situations it is very questionable which act(s) its agent can justifiably believe to maximize expected utility, but in many cases it is not questionable at all. The reason why I am concerned with this kind of AU in this paper, and not with some other kind of AU, is simply that I believe it to be the most plausible form of AU.

deontologically-adjusted AU1, call it DAAU1. If it was reasonable to believe so, I would call AU1 *self-inactivating in a weak sense*. I would call AU1 self-inactivating in a weak sense, if it was reasonable to believe that in a society (or in a world) in which every moral agent would be aware of every moral agent's moral beliefs and how well each moral agent tends to comply with her own moral beliefs (and in which every moral agent would be aware that every moral agent is aware of every moral agent's moral beliefs and how well each moral agent tends to comply with her own moral beliefs), it would not be likely to be for the best in terms of AU1 if every moral agent was a perfectly complying AU-believer.²

One way to argue that AU1 is self-inactivating in a weak sense would be to argue that two perfectly complying AU1-believers, who are a couple and in a serious relationship,

² I would call AU1 *self-inactivating in a strong sense*, if it was reasonable to believe that the best society (or the best world) in terms of AU1 would not be likely to be such in which every moral agent would be a perfectly complying AU-believer. It might be so (although I doubt it) that AU1 is self-inactivating in a weak sense. But if this is the case, it does not have to mean that AU1 is self-inactivating in a strong sense, as the best society in terms of AU1 might be such in which every moral agent would be a perfectly complying AU-believer and in which some or all of them would be mistakenly perceived by some or all moral agents not to be perfectly complying AU-believers. My concept of self-inactivation can be seen as a modification of what Derek Parfit (1984, pp. 53-55) calls a directly collectively self-defeating theory.

It is important that my definitions end with "a perfectly complying AU-believer" rather than with "a perfectly complying AU1-believer", as otherwise AU1 would be self-inactivating simply because it is a form of subjective AU. This would make the concept of self-inactivation less fruitful. For example, one form of *objective* AU, call it AU2, says that an act is morally right, if and only if it in fact maximizes utility in the world. Surely in a society in which every moral agent would be aware of every moral agent's moral beliefs and how well each moral agent tends to comply with her own moral beliefs, it would be better in terms of AU1 if every moral agent was a perfectly complying AU2-believer than if all of them were perfectly complying AU1-believers. And surely the best society in terms of AU1 would be such in which every moral agent would be a perfectly complying AU2-believer, if the best society in terms of AU1 would be such in which every moral agent would be a perfectly complying AU-believer.

and who are known to be perfectly complying AU1-believers, would be likely to have a better relationship in terms of AU1 (i.e. a relationship which would produce more utility), if they were known to be *perfectly complying deontologists about cheating* (i.e. perfectly complying deontology-believers about infidelity)³ *and perfectly complying AU1-believers about all other kinds of acts* (call them perfectly complying DAAU1-believers). This claim is based on the argument that the possibility of being cheated by a partner, who is a perfectly complying AU1-believer, is bigger, than the possibility of being cheated by a partner, who is a perfectly complying DAAU1-believer, and that because of this some counterproductive feelings (in terms of AU1) that would arise in both parties of the relationship, if both parties were known to be perfectly complying AU1-believers, would not arise in either party, if both of them were known to be perfectly complying DAAU1-believers.

In this paper my main argument is that this line of thought does not work. I leave it open whether there might be some way of arguing plausibly that AU1 is self-inactivating in a weak sense.⁴ If it could be argued plausibly that AU1 is self-inactivating in a weak sense and if AU1-believers were aware that AU1 is self-inactivating in a weak sense, it would probably make most AU1-believers feel at least a little bit uncomfort-

³ It is, of course, far from obvious how exactly cheating (or infidelity) should be understood. In this paper I understand cheating in the following way. A person cheats on her partner (with whom she forms a couple and with whom she is in an intimate relationship), if she is sexually or romantically involved with someone else, if she has not got her partner's approval for that, and if the relationship which she is in has not been agreed by her and her partner to be an open relationship. This definition of cheating can, of course, be questioned, but I think it is good enough for the purposes of this paper. I am also aware that "being sexually involved with" and "being romantically involved with" can be understood in many different ways, but I do not think that the subject area of this paper requires an analysis of how exactly "being sexually involved with" and "being romantically involved with" should be understood.

⁴ It seems to me that arguments based on lying and promise-breaking would not work either. See Singer (1972).

able even if they did not think that it would make AU1 implausible.⁵ Many opponents of AU, on the other hand, would probably try to use it against AU1. I do not discuss whether it could be used plausibly against AU1. It is a question that AU1-believers do not really need to worry about as long as it has not been argued plausibly that AU1 is self-inactivating in a weak sense, although it would surely put AU1 on a firmer basis, if someone argued plausibly that even if AU1 was self-inactivating in a weak sense, it would not make AU1 implausible.⁶

2 Perfectly Complying AU1-believers and Cheating

It may be argued: “If both parties of a couple, who are in a serious relationship, were known to be perfectly complying AU1-believers, both parties would be distrustful and suspicious, as both of them would know that the other party of the relationship would cheat on her (or him) *every time* when she (or he) should do that from the point of view of AU1.⁷ Since both parties know that the other party is so little committed to her, it

⁵ See Singer (1972, p. 94).

⁶ I think that there are good reasons for believing that even if AU1 was self-inactivating in a weak sense (or even in a strong sense), it would not make it an implausible moral doctrine. See e.g. Norcross (1997, pp. 387-393); Railton (1984, pp. 116-117).

⁷ Certainly there can be situations in which a person can justifiably believe to maximize expected utility by cheating on her partner even if they had a good relationship. Norcross (1997, pp. 382-383) writes: “John is married to Jane. He has always professed love and loyalty to Jane, and she has to him. She has been an excellent spouse and given him no cause for complaint...John is at a party also attended by Jane and Mary...John goes home with Mary... [W]hat if Jane is unaware of John’s behavior, and so doesn’t feel any pain at all? Wouldn’t the utilitarian have to say that there is nothing at all wrong with John’s behavior in this case? In fact, it might even be admirable behavior, if he and Mary enjoy it enough.” Surely, from the point of view of AU1, John should also take into account (before cheating) how big a

would bring distrust and suspicion into their relationship, which would naturally not be ideal in terms of AU1. Because of this, it would be better in terms of AU1, if both parties of the relationship were known to be perfectly complying DAAU1-believers. Perfectly complying DAAU1-believers would always comply with the principle “Do not cheat on your partner!”. Apart from this, they would always comply with AU1. The fact that they would always comply with the principle “Do not cheat on your partner!”, would certainly be bad in terms of AU1 in a sense that each party of the relationship would fail to comply with AU1 every time when she should, from the point of view of AU1, cheat on her partner. However, situations in which a person should, from the point of view of AU1, cheat on her partner, are so rare that in the calculus of AU1 the problem that perfectly complying DAAU1-believers would never cheat on their partners in situations in which that is what they should do from the point of view of AU1, would be outweighed by the lack of distrust and suspicion in their relationship.”

I believe that this argument does not hold. I do believe that situations, in which a person should cheat on her partner from the point of view of AU1, are rare in relationships in general and would also be rare in relationships formed by perfectly complying AU1-believers. But this is the reason why neither of the perfectly complying AU1-believers would be distrustful or suspicious. Since each party of the relationship knows that there is only a very small possibility that the other party will ever cheat on her, the argument of the previous chapter does not give enough reason to believe that there is a real problem in terms of AU1. The best description of the situation is not that each party

risk there is that Jane will find out about the cheating, how sad and indignant she would become if she found out about the cheating, how likely it would be that John would feel somehow uncomfortable after the cheating, and so on. Nevertheless, the point is that we could give such a description of this case so that it would be clear that John can justifiably believe to maximize expected utility by cheating on Jane with Mary. In that case, from the point of view of AU1, it would be John’s moral obligation (besides or instead of being morally admirable) to cheat on Jane with Mary, if he can do it.

would cheat on her partner *every time* when that is what she should do from the point of view of AU1. The best description is that each party would cheat on her partner *only if* that is what she should do from the point of view of AU1. (It should be noticed that in many real life relationships which involve cheating, this is clearly not the case.)⁸

It may be responded: “You are right, neither of the perfectly complying AU1-believers would be distrustful or suspicious, or if they were, their distrust and suspicion would be so mild that it would be a smaller problem in terms of AU1 than the problem that they would never, if they were perfectly complying DAAU1-believers, cheat on each other in situations in which they should cheat from the point of view of AU1. But both of them would nevertheless be *happier*, if they knew that the other party of the relationship is a perfectly complying DAAU1-believer and thus knew that the other party would be *so committed* to her that she (the other party) would definitely never cheat on her. This is enough to outweigh in the calculus of AU1 the problem that each party would not comply with AU1 in those rare situations in which she should cheat on her partner from the point of view of AU1. This means that it would be better in terms of AU1, if both parties were known to be perfectly complying DAAU1-believers than if they were known to be perfectly complying AU1-believers.”

It can be argued, however, that if both parties of the relationship were perfectly complying AU1-believers, they would not mind that the other party is not committed to her in an absolute, deontological, sense, because both of them would know that if her partner would ever cheat on her, it would happen on the grounds that she (by “she” I am referring to “the knower”, not to the partner of the knower) believes to make cheating morally justifiable and even morally obligatory. Certainly each party *hopes* that her partner will never cheat on her. Each party hopes that she is able to sexually satisfy her

⁸ It seems to me that in the world which we are living, an overwhelming majority of the acts of cheating that actually occur, are morally wrong from the point of view of AU1.

partner so completely, is able to show so much love towards her partner, and is so much loved by her partner, that her partner will never cheat on her. But on the other hand, they both morally accept the possibility of being cheated. Since each of them knows that her partner would cheat on her only on the grounds that she (the knower) considers morally justified, each party knows that she will never be treated by her loved one in a way that she (the knower) considers morally wrong. What is important for a perfectly complying AU1-believer is not that she knows that her partner will definitely never cheat on her. What is important for her is that she is able to sexually satisfy her partner so completely, is able to show so much love towards her partner, and is so much loved by her partner, that the possibility of being cheated by her partner would be as small as possible.⁹ A perfectly complying AU1-believer would be likely to feel sad (although not morally indignant), if she found out that her partner has cheated on her and that she was morally obliged, from the point of view of AU1, to cheat on her. But this does not mean that she would prefer that her partner would be a perfectly complying DAAU1-believer rather than a perfectly complying AU1-believer.

⁹ Surely, from the point of view of AU1, other things being equal, stronger love (both giving and getting it) is always better than weaker love, and more sexual satisfaction is always better than less sexual satisfaction, as love (both giving and getting it) and sexual satisfaction are very pleasurable feelings. The more a person gets these things from one partner and the more she loves one partner, the less likely it is that she should ever cheat on her partner from the point of view of AU1. This is so because the more love and sexual satisfaction a person gets from her partner, and the more she loves her partner, the less she can benefit from cheating on her and the less likely she is to benefit from cheating on her. It can also be argued that the better her partner satisfies her sexually, the more she is loved by her partner, and the more she loves her partner, the more she would lose if her partner found out that she has cheated on her. This means that if a person was sexually very satisfied with her partner, was very much loved by her partner, and loved her partner very much, she should not, from the point of view of AU1, cheat on her partner in some situation even if the risk of getting caught for that was extremely small and even if she was likely to benefit quite a lot from cheating in that situation.

3 AU1-believers and Cheating

I am unable to come up with an argument that would give enough reason to believe that the cheating issue is so problematic in terms of AU1 that it shows AU1 to be self-inactivating in a weak sense.¹⁰ It should be noticed, however, that two people, who are a couple and in a serious relationship, and who are known to be AU1-believers (and who know that they are known to be AU1-believers), might be likely to have a better relationship in terms of AU1 if they were known to be DAAU1-believers. This means that AU1 is *self-effacing in a weak sense*. I call AU1 self-effacing in a weak sense, because it is reasonable to believe that in a society (or in a world) in which every moral agent would be aware of every moral agent's moral beliefs (and in which every moral agent would be aware that every moral agent is aware of every moral agent's moral beliefs), it would not be likely to be for the best in terms of AU1 if every moral agent was an AU1-believer.¹¹ This is so ultimately because it is reasonable to believe that some of them would not be very good compliers of AU1.

¹⁰ I have, however, ignored one important argument. It can be argued that someone who would be prepared to cheat, and who would cheat, on her partner, if that is what she should do from the point of view of AU1, would not be capable to genuinely love her partner and that this makes AU1 self-inactivating, as being capable of loving genuinely is very valuable in terms of AU1. I believe that this a very complex argument and that it would require such an extensive moral psychological discussion that I cannot discuss it here. See Mason (1998); Norcross (1997); Railton (1984) on issues which are related to this argument. It should be noticed that this is not only an argument that AU1 is self-inactivating in a weak sense. It is also an argument that AU1 is self-inactivating in a strong sense, as this argument does not depend on whether both parties of a relationship know that the other party is a perfectly complying AU1-believer.

¹¹ I would call AU1 *self-effacing in a strong sense*, if it was reasonable to believe that the best society (or the best world) in terms of AU1 would not be such in which every moral agent would be an AU-believer. My concept of self-effacement is a modification of Parfit's (1984, pp. 40-43) concept of self-effacement.

It seems to me that some people would be likely to cheat on their partners in some situations in which it is morally wrong from the point of view of AU1, if they were AU1-believers, but would be unlikely to ever cheat on their partners, if they were DAAU1-believers. This means that a relationship of two AU1-believers might be such in which neither party could trust that the other party would cheat on her *only* in those rare situations (which would probably never even occur during their relationship) in which it would be justified from the point of view of AU1.

If someone thinks that there is something intrinsically morally bad or wrong about cheating, she might feel more antipathy towards cheating (and feel more remorse after having cheated) than she, if she was an AU1-believer, would feel antipathy towards that kind of cheating which she knows to be morally wrong from the point of view of AU1 (and feel remorse after having cheated in a way which she knows to be morally wrong from the point of view of AU1).¹² So if she evaluates the moral badness and wrongness of cheating only on the grounds of the maximization of subjective expected utility, she may not take the ethics of cheating as seriously as she would take it if she believed that cheating is intrinsically morally bad or wrong, especially if she thought that cheating is intrinsically morally *very* bad or wrong. This means that she might comply with AU1 worse if she was an AU1-believer than if she was a DAAU1-believer. Surely, if she was a DAAU1-believer, she would be more likely not to cheat on her partner in those rare situations in which she should, from the point of view of AU1, cheat on her partner, than if she was an AU1-believer. However, in the calculus of AU1 this problem is outweighed by the fact that she would be more likely, if she was a DAAU1-believer, not

¹² See also Parfit (1984, p. 40).

to cheat on her partner in those *common* situations in which cheating would be morally wrong from the point of view of AU1, than if she was an AU1-believer.¹³

4 Conclusion

It may be argued that my discussion about perfectly complying AU1-believers does not make much sense or at least is not very fruitful, as no perfectly complying AU1-believer exists, has ever existed or will ever exist. It could be argued that at most there can be moral agents who are AU1-believers and who comply with it very well, in other words, who rarely do acts which, according to AU1, are morally wrong.

It is probably true that no perfectly complying AU1-believer exists, has ever existed or will ever exist. People are often tired, busy or influenced by some strong emotions, and these things may often cause AU1-believers to make bad judgments from the point of view of AU1. So it seems to me that even the most unselfish AU1-believers who would always be very motivated to comply with AU1, would sometimes do acts which are morally wrong from the point of view of AU1. It also seems to me that there could not be a person who would *always* be motivated enough to pay sufficient attention to

¹³ Another reason why some people might comply with AU1 better if they believed in DAAU1 than if they believed in AU1 is that some people, if they believed in AU1, might be prone to deceive themselves into believing that cheating is morally right from the point of view of AU1 in many situations in which it is not, whereas they could never deceive themselves into believing that cheating is morally right from the point of view of DAAU1 in some situations, as according to DAAU1, cheating is never morally acceptable. This complication could, however, be avoided if these people as AU1-believers would always use a deontological *decision procedure* regarding cheating. In other words, these people should understand that if they always complied with the principle “Do not cheat on your partner!”, they would be more likely to comply with AU1 better in the long run than if they did not always comply with the principle “Do not cheat on your partner!”.

AU1. It seems to me that it is inevitable that even the best compliers of AU1 would sometimes do acts that are morally wrong from the point of view of AU1.

I believe, however, that my discussion about perfectly complying AU1-believers makes sense and is fruitful, as it helps us to understand AU1 on a very theoretical level. I believe that act-utilitarian interaction (or co-operation) problems can be understood best as a whole, if we first look at them on a very theoretical level, as that can help us to understand them better on a more practical level which is closer to the real world.

Bibliography

Mason, Elinor (1998): "Can an Indirect Consequentialist Be a Real Friend?", *Ethics*, Vol. 108, 386-393.

Norcross, Alastair (1997): "Consequentialism and Commitment", *Pacific Philosophical Quarterly*, Vol. 78, 380-403.

Parfit, Derek (1984): *Reasons and Persons*. Oxford University Press, Oxford.

Railton, Peter (1984): "Alienation, Consequentialism, and the Demands of Morality", *Philosophy and Public Affairs*, Vol. 13. Reprinted in Samuel Scheffler (ed.) *Consequentialism and Its Critics* (1988). Oxford University Press, Oxford, 93-133.

Singer, Peter (1972): "Is Act-Utilitarianism Self-Defeating?", *The Philosophical Review*, Vol. 81, 94-104.