

Consequences, Numbers, and Reasons

Berke Çelik | Boğaziçi University

Introduction

Consequentialism is usually criticized because many consequentialist theories oblige, or permit, killing one person to save five people, in popular thought experiments, such as Trolley and Transplant. In this paper, I argue that these thought experiments that are usually employed to provide counter-arguments against consequentialism are too charitable towards consequentialism. Most consequentialist theories that permit killing one to save the five, also provide justifications for killing one thousand to save one thousand and one. After exploring consequentialist reasoning in cases where there are trivial stakes in agent-neutral axiological terms, I show why it's quite hard for consequentialists to avoid counter-intuitive verdicts in trivial value cases.

The core issue I identify is consequentialism's lack of stake sensitivity - its inability to appropriately weigh the magnitude of outcomes against other moral considerations. I'll offer a quasi-consequentialist account that avoids counter-intuitive verdicts in trivial value cases that's essentially symmetrical to threshold deontology: A consequentialist or quasi-consequentialist hybrid theory should be sensitive to the question of "how much is at stake?" After exploring Setiya's and Scheffler's attempts to improve consequentialism, I'll argue that these accounts also fail to give intuitive verdicts in trivial value cases. I then briefly explore how recent work on justifying and requiring normative reasons may provide the consequentialist with a way to develop a hybrid theory that can capture the weight of non-consequentialist reasons in trivial cases in a way that's more aligned with our intuitions.

Consequentialism

There are various versions of consequentialism, and although I will tend to specify which variant of consequentialism, I believe that most consequentialist theories are vulnerable to the critique from trivial value increase arguments. As the counter-intuitive verdicts in trivial cases stem from the heart of consequentialism, which could be defined in the following way:

The Heart of Consequentialism: The goodness of outcomes is the only relevant moral criterion to evaluate objects of moral judgment.

Maximizing act consequentialists should kill one person to save five people, and should kill five to save six. Agent-neutral value calculus dominates everything, and the

evaluative focal point is acts. However, some may think for consequentialists whose evaluative focal point isn't acts, killing five to save six is better than not killing may not be obligatory. Although act consequentialists may be more vulnerable to my upcoming counter-arguments, the vulnerability to counter-arguments from other cases persists for any moral theory that adheres to "the heart of consequentialism". For clarity and length-related reasons, I will focus on maximizing act consequentialism, only briefly discussing other variants of consequentialism, as the problem I identify persists for global consequentialists and non-maximizing consequentialist theories.

Trivial cases

Transplant Case: *There are five patients who are about to die. A health patient checks into the hospital, whose organs can save five patients. Should we kill or refrain?*

Act consequentialists must kill. And now consider the following case:

Strictly Utilitarian Surgeon: *Dr. Jeremy is an omniscient hospital administrator and, in his hospital, there are one thousand and one patients who will soon die unless they receive a transplanted organ. The region is hit by an earthquake and a thousand patients with very healthy organs check in to the hospital. Should Jeremy kill 1000 patients to save 1001 patients?*

The outcome with 1001 lives is superior to 1000 lives if all the other morally relevant, agent-neutral properties are the same according to consequentialism (not only maximizing act consequentialists), and although some consequentialist theories, such as Hooker's rule consequentialism (Hooker 2000), may claim that Jeremy isn't permitted to kill a thousand people, from an axiological standpoint the universe in which Jeremy kills is better than the universe in which Jeremy doesn't kill. Consequentialists may be permitted to not kill, but they must, axiologically, admit that the outcome in which Jeremy kills is better than the other outcome (Howard 2021).

This highly idealized and odd thought experiment doesn't provide us a reason to reject consequentialism wholesale, however, it does point towards something deeply counter-intuitive about consequentialism. According to consequentialism we are not only justified to kill the few to save the many, we are justified to kill n number of people to save $n+1$ number of people or create n units of suffering to increase $n+1$ units of happiness. The issue at stake here is the lack of the following moral property:

Sensitivity to Trivial Increases (STI): *The strength of consequentialist reasons for action weakens in proportion to how trivial the potential value gain becomes.*

I'd argue that most people would agree that we are not justified in killing a great number of people - let that number be x - to save $x+1$ people. We may have strong reasons to kill one to save a thousand people, but weak reasons to kill 100 people to save 101 people. A traditional consequentialist might not inherently reject the stake-sensitivity point and may agree that the strength of consequentialist reasons to "kill" depends on the context. However, consequentialism lacks a robust framework for weighing these differences intuitively. Specifically, it lacks a systematic approach to consider how the strength of consequentialist reasons may vary based on stakes and how these reasons should be balanced against competing non-consequentialist considerations. Moreover, while one might argue that the counter-intuitiveness of killing 100 to save 101 stems from our difficulty in internalizing large numbers, the issue is more fundamental: it concerns consequentialism's inability to intrinsically value non-consequentialist factors and competing considerations, such as the inherent badness of certain act-types (like killing).

This stake-sensitivity is somewhat symmetrical to threshold deontology (Alexander 2000), which argues that we should be deontologists with stake-sensitivity: Murder may be wrong even when it will prevent two murders, but murdering someone to prevent the murder of a million people may be permissible, or obligatory. The weight of deontological reasons varies in relation to how much is at stake. Later in the paper, I'll argue that consequentialists should incorporate a similar, yet asymmetrical form of stake-sensitivity.

Agent-relative reasons and personal projects

The intuitions I'm appealing to in trivial value cases could be formulated in agent-relative terms, but also in agent-neutral terms (for example by an appeal to value conservatism (Nebel 2022)). Our strongest intuition in the trivial cases above may be considered as agent-relative or agent-centered, in that we may feel like we have agent-relative reasons to refrain from killing to produce optimal outcomes. Before delving deeper into agent-relativity, I would like to discuss why our intuitions about agent-relative reasons to refrain from doing some harm to produce an optimal outcome may be unreliable.

If we have the considered intuition that we ought not to kill for trivial gains, then a moral account should entail or provide normative reasons that would render killing for trivial gains impermissible not because of our aversion to "getting our hands

dirty” (Coady 2024). Most people, including myself, don’t like the idea of killing, even when the act of killing means saving 5 or 50 people. Although this dislike towards killing and the demandingness of this idea may push us to reject consequentialist reasoning, I would argue otherwise. I believe that one of the best features of consequentialism is its “passive demands” (Sobel 2008): Wouldn’t a moral theory that rejects pulling a lever that will result in one killing to save the five be too demanding on behalf of the five people who lie on the tracks? I will not discuss reasons why consequentialism is *worth saving*, but merely discuss how it could be modified to give more intuitive verdicts in trivial value cases.

The main aim of this paper isn’t to point towards the counter-intuitive nature of consequentialism to reject it, but to discuss how it could be rendered more intuitive with revisions that are not too *theoretically* costly. I would argue that the consequentialist obligation to prioritize aggregated value across the population or producing the best outcome over our personal projects should remain but with emphasis on *priority*. The specific feature of mainstream consequentialism I’m criticizing is its inability to prioritize and compare various normative reasons, as it only accepts consequentialist reasons. Contrary to virtue ethics or certain deontological theories that may value beneficence intrinsically (Formosa & Sticker 2019), consequentialist theories are unable to entail or respond to strong non-consequentialist reasons in cases with trivial value gains.

I would like to reiterate that the focal point of this paper is the axiology of consequentialism rather than its decision procedures, which is the reason why moving away from “naive” act consequentialism to sophisticated global consequentialism does not solve the problem, as sophisticated global consequentialism captures non-consequentialist normative reasons only indirectly, and does not capture our intuitions in trivial value cases. Although moving away from act consequentialism to more sophisticated and seemingly expressive versions of consequentialism may recommend more intuitive types of moral reasoning and verdicts in certain cases or real-world contexts, counter-intuitiveness in trivial cases stems from the “Heart” of consequentialism, and as a result, changing the evaluative focal point or decision procedures is not sufficient.¹ Consider the following case:

Squirrel Makes the Difference:

Let’s say I can either save a or b:

- a) 100 people, 10 of whom I have great and deep relationships with*
- b) 100 people I’m not even acquainted with, plus a squirrel*

¹ For the axiological identity of global and act consequentialism, see Thornley (2022).

A maximizing act consequentialist who holds “sentient lives” as the ultimate good would argue that we are obligated to opt for b. This seems highly implausible. The issue is less acute for scalar (and arguably satisficing) consequentialists, as they don’t use the language of “obligation,” but according to scalar consequentialism b is still more preferable than a, so it’d be better if I killed, but I am not obligated to. The issue persists for global consequentialists who haven’t revised their axiology but merely use a (much) more sophisticated decision procedure. Although it may be argued that global consequentialism would not have permitted the usage of a naive decision procedure in cases such as this, according to the axiology, the consequentialists must admit that in this isolated instance, the world in which I save b is better than the world in which I choose a.

We may prefer a lower number of deaths, and I am not arguing against the idea that we should aim to minimize the number of deaths. But a moral account that provides no space to considerations other than minimizing badness and maximizing goodness seems incomplete: Obligation to allow the death of my close friends to refrain from choosing an option that is only trivially better in agent-neutral terms. Rather than provide an argument against agent-neutrality, my aim with this thought experiment is to show that consequentialist reasons for action seem weaker as the stakes get lower. In this case, the majority of the people would have a strong non-consequentialist intuition that would push them to choose a, rather than b, due to special relations. Would a person be obligated to kill a friend to prevent a nuclear war? I’d argue yes. But the strength of normative reason that obliges or permits to killing of a friend gets weaker as the stakes get more trivial, as we see in *Squirrel Makes the Difference*.

Hybrid moral theories

Agent-neutral consequentialism purports that we ought to be indifferent toward the interests of our children versus other children (Lazar 2018). This seems counter-intuitive and wrong to many people. There has been a small industry of philosophers that have attempted to introduce agent-relative considerations into consequentialism. Before delving deeper into these attempts, I will briefly discuss and define agent-centered restrictions.

Agent-centered restrictions

Agent-centered restrictions correspond to the following idea: We are morally prohibited from committing (certain) morally objectionable actions, even if doing so would reduce the overall occurrence of similar objectionable actions (Alexander 1987). For example, let’s say killing is bad and should refrain from that act-type regardless of the

outcome it will produce. But this implies that murdering in order to prevent five murders from happening is bad. This is in tension with the maximizing conception of rationality: If something is bad, shouldn't we want less of it? If violating a moral principle is bad, shouldn't we aim to minimize the number of violations of that moral principle?

The answer "No" to the questions above has been thought to be in tension with rationality, hence perceived as a paradox by some philosophers, such as Scheffler, who rejects agent-centered restrictions partially because of this paradox and its tension with rationality (Scheffler 1985). I agree with Scheffler, especially in cases where we can minimize a great number of morally bad acts or increase total value greatly. I would argue that a moral account should give the verdict that we are required to kill in cases with astronomical numbers. Surely, if I can prevent a nuclear war with a murder, I should murder. Should I kill someone to prevent 1000 murders? 500? 100? My answer to all three of these question marks would be yes. Symmetrical to SAI, now consider:

Sensitivity to Astronomical Increases (SAI): *As the magnitude of normatively significant stakes becomes astronomically high, the requiring strength of consequentialist reasons for action significantly increases in proportion, dominating competing non-consequentialist reasons.*

I will explain why I used the word *requiring* in the final section. So according to SAI, in scenarios where committing a morally objectionable act will prevent an astronomical number of similar violations or result in a greatly better outcome, consequentialist reasons dominate competing normative reasons, hence requiring us to commit that morally objectionable act.

I would argue that a theory that would reject SAI would be too passively demanding on people who are on the receiving, rather than acting, end of normatively significant actions. I am talking about axiology here, rather than a decision procedure: We have very strong reasons to refrain from killing even when we believe that the stakes are high. But vaguely, I believe that a moral account should *oblige* us to inflict harm if there are astronomical stakes. If we face the option of keeping our hands morally clean or preventing millions of deaths, we ought to prioritize the claims of people who will die, over our urge and normative reasons to refrain from committing morally objectionable acts.

So, in the remainder of the paper, I'll review various strategies employed by consequentialists to render (agent-neutral) consequentialism more intuitive. My goal will be the following: Evaluate whether these strategies enable consequentialist agents to not violate both STI and SAI. I will mainly focus on two strategies: Setiya's (2018) and Scheffler's (1994).

Agent-centered prerogatives without restrictions

Scheffler argues that agent-centered restrictions lack a robust foundation and should be rejected, however, he argues that we can accept agent-centered prerogatives while rejecting agent-centered rejections (Scheffler 1994). Maximizing consequentialists argue that we are obligated to produce the best outcome available to us, and it's impermissible for agents to produce a suboptimal outcome (MacAskill et. al 2023). Scheffler believes that the introduction of agent-centered prerogatives to consequentialism would permit agents to produce a suboptimal outcome due to agent-relative considerations. Agent-centered restrictions may be paradoxical, but prerogatives that permit agents to produce a suboptimal outcome for agent-relative reasons are not paradoxical, as they are not in tension with maximizing property of rationality (Alexander 1987).

Schefflerian prerogatives enable consequentialist theories to permit agents to save the group that includes their friends, rather than saving the random group of people that includes a squirrel. An agent-centered prerogative that permits us to give greater weight to our commitments and (personal) projects, hence enabling agents to respond to and act in accordance with agent-relative reasons that stem from special obligations or other things. The resulting theory is more robust compared to traditional variants of consequentialism especially when it comes to Williams "utilitarian inability to commit to any personal projects and commitments" critique (Williams 1973).

Scheffler's account doesn't violate STI in cases where we have agent-relative reasons that permit us to do *less than best*. But prerogatives could be used to violate SAI, by permitting agents to produce a suboptimal outcome even when the stakes are really high. Also, in the revised transplant, Schefflerian agents should still hold that the outcome in which we kill in the Revised Transplant Case is better and that killing is still permissible, just not required. The problem is, that our intuition in the revised transplant case is that killing is impermissible, and that outcome in which the surgeon kills (not us) may not actually be better. Although the Schefflerian move does a good job of capturing agent-relative intuitions, it fails to capture the cases in which i) stakes are astronomical and not producing the suboptimal outcome is impermissible and ii) value gain is so trivial that producing the best outcome is impermissible due to weightier non-consequentialists reasons for opting for a less than best outcome.

Setiya's account

Another account of consequentialism that permits agents to refrain from killing one person to save the five is put forward by Kieran Setiya. Setiya's consequentialism (Setiya 2018), consists of the following claims:

Action-Preference Nexus: Among the actions available to you, you should perform one of those whose consequences you should prefer to all the rest.

Agent-Neutrality: Which consequences you should prefer is fixed by descriptions of consequences that make no indexical reference to you.

Setiya then considers the following popular case

Footbridge Thought Experiment: *A trolley is about to kill five people, and you can prevent this by pushing an innocent person onto the tracks.*

Setiya argues that pushing an innocent person in front of the tracks isn't permitted even when that act would save five lives. Relatedly, his theory recommends that we should not kill in the Revised Transplant Case, not due to some agent-centered restrictions, but because it would result in a worse outcome, as he holds that killing in order to prevent deaths is worse than allowing deaths, appealing to intuitions and commonsense. So, we can say that Setiya's account does not violate STI, and admirably provides agent-neutral reasons for not violating STI. The problem is, that Setiya's account also deems killing impermissible in non-trivial cases, and arguably in cases where killing would result in an astronomical increase of value.

Agent-relative arguments are argued to be motivated by a shaky argument called "Keeping Your Hands Clean" (Howard 2021), and one way in which Setiya's account is better than other moral accounts is it avoids this problem, by reformulating this intuition in an agent-neutral way. It'd be better to kill to save the five, but becoming a murderer seems bad, and I am permitted to avoid becoming a murderer. Setiya is very successful in capturing this intuition in an agent-neutral way, by arguing that we should defend that others too should not kill in scenarios where we also should not kill.

However, I contend that this argument encounters a counter-argument that is parallel to the "Keeping Your Hands Clean" argument. Our inherent aversion to being labeled as murderers significantly shapes our reasoning in such moral dilemmas. While the removal of self-reference might reinforce the argument against committing murder, it doesn't alter the fundamental unreliability of this underlying intuition. Consider a scenario with exponentially higher stakes: is it justifiable to kill 100 people to save 500? I posit that it is, given the magnitude of the situation (400 lives at stake). Although the act of killing a hundred individuals is undeniably grave, a moral

framework that precludes the action necessary to save 900 lives does not adequately account for scenarios involving such astronomical numbers, hence violating SAI.

Threshold consequentialism? Numbers and strength of reasons

Well-being may be morally significant, or it may be the most morally significant thing, but we need not to assert that dignity or autonomy does not matter, or reject the importance of non-welfarist normative reasons (Chappell 2022). Chappell argues that utilitarianism is justified in focusing on well-being, but it may be wrong to not focus on anything besides well-being and coins "beneficentrism" which entails the claim that well-being is the most important thing, but not the only important thing.

A similar point could be made about consequentialism: Promoting good may be the most important source of normative reasons, but it need not be the only. Promoting good, or bringing out better outcomes or world-states may be the most important moral goal or source of moral reasons; however, this does not imply that non-consequentialist considerations do not matter at all, or matter only indirectly. We may come up with a theory that is mostly consequentialist but accepts that there are, at least in some cases, weighty moral reasons for not acting like a naive or sophisticated consequentialist.

The issue I am discussing here is very relevant to the recent literature on normative uncertainty (Macaskill et. al 2020). I won't give a detailed account of how recent work on normative uncertainty may contribute to the goal of building an expressive quasi-consequentialist theory, but I will merely show how a richer conception of normative reasons may enrich our understanding of conflicting normative reasons and hybrid moral theories.

Requiring and justifying strength of reasons: A promising avenue

We have strong consequentialist reasons for action in most cases. However, this strength isn't absolute; the strength of consequentialist reasons is influenced by the value at stake. Consider the act of killing a thousand people: while the inherent badness of this action remains constant, its weight in our moral calculation can vary significantly depending on the outcomes. For instance, our consequentialist reason to kill a thousand people to save ten thousand might be stronger than our reason to kill a thousand to save only a thousand and one. This illustrates how consequentialism correctly recognizes that the value of an act depends on the stakes involved.

However, what consequentialism fails to admit is that while consequentialist reasons for acting may get stronger depending on how much value we may bring out and overall stakes, it gives no room for understanding how non-consequentialist

reasons may also get stronger depending on “stakes.” For example, our consequentialist reasons are stronger in the original Transplant scenario (killing one to save five) compared to a Revised Transplant scenario where we might kill a thousand to save a thousand and one. Yet, consequentialism doesn’t account for how non-consequentialist considerations (like the prohibition against killing) might also strengthen as the number of lives taken increases.

This is quite similar to the difference between deontology and threshold deontology: We should (usually) be deontologists, but when there are astronomically high stakes, we should act like a consequentialist. To reformulate in this section’s language: Consequentialist reasons trump non-consequentialist agent-neutral, or agent-relative reasons when the stakes are too high.

However, we can go even further by utilizing Gert’s work on normative reasons (Gert 2003). Rather than viewing reasons as having one dimension of value (strength or weight), Gert argues that we should view reasons as having two dimensions of value: Requiring and justifying. Requiring strength of reasons determines whether an action is obligatory while justifying strength (of a reason) determines the permissibility of an action. This conception is especially good for accommodating agent-relative intuitions. For example, in Peter Singer’s Shallow Pond case (Singer 1972), it may be impermissible to not save a child because our suit will get wet (low cost), requiring strength of our consequentialist reason is not balanced by neither the requiring or justifying strength of competing reasons. But, if an agent faces a decision in which she will lose two arms to save a child, justifying the strength of her non-consequentialist or agent-relative reason is strong enough to override the consequentialist reason for action, hence permitting not losing two arms and not saving the child (Pummer & Munoz 2022).

Requiring strength of our agent-neutral reasons may dominate other reasons when facing a decision to kill one to save five, but when we face a decision to either save ten friends or save eleven random people, the strength of justifying agent-relative reason may be much greater than the requiring strength of our agent-neutral reasons for killing our friends to opt for a trivially better outcome. The same conceptual machinery -requiring and justifying normative reasons- also enables moral theories to capture SAI: When the stakes are astronomically high, requiring the strength of consequentialist reasons dominates anything else. A symmetrical claim can be utilized to express STI. Requiring strength of consequentialist reasons diminishes as the value increase in question gets lower and lower.

I won’t give a detailed account of how to weigh consequentialist reasons against non-consequentialist reasons, or when we are permitted to act based on agent-relative

reasons. What I'm mainly pointing towards is, that requiring and justifying the conception of normative reasons is a good tool that could be utilized to build hybrid theories that don't violate STI and SAN. Additionally, they are also good concepts to use for the introduction and formalization of agent-relative reasons (and certain non-consequentialist reasons that are not agent-relative) in general. For example, W.D. Ross' account of competing prima facie duties could also be consequentialized using this machinery (Philips 2019).

Finally, justifying and requiring strength of reasons is also relevant for moral reasoning under (moral) uncertainty. Let's consider an agent that has a high credence in the idea of agent-neutral consequentialism, but at the same believes that Kill seems egregious in the revised transplant case. Rather than domesticating this theory with Schefflerian prerogatives, it may seem more preferable for him to accept that the requiring strength of agent-relative reasons or dignity-related reasons are greater in certain cases compared to competing consequentialist reasons in the Revised Transplant. As the Schefflerian revision would still require him to admit that it's permissible for him to kill his friends for the greater good and that the outcome in which he does that is better. Accepting that moral betterness relationships are not only governed by agent-neutral axiology, similar to the idea that all normative reasons for actions don't have to be consequentialist in nature seems like a good accommodation strategy for consequentialists who would like to build a more commonsensical theory or take normative uncertainty seriously.

Concluding remarks

Justifying and requiring reasons may increase the expressiveness of consequentialist theories without adding additional problems. Although I have been vague with regard to points at which consequentialist reasons are dominated by non-consequentialist reasons, I have provided a potential avenue for consequentialist theories that may render them more intuitive and expressive.

I believe that moving towards a richer conception of normative reasons provides the most promising avenue for the consequentialist. Introducing agent-centered restrictions implies a hard-to-avoid paradox, and Setiya's account captures the intuition behind agent-centered restrictions in an agent-neutral way, and while really innovative, it neither accepts agent-relative reasons nor entails SAI. While Scheffler's account accepts the existence of agent-relative reasons, it fails to capture impartial considerations implying STI. Accepting (usually weak) non-consequentialist reasons rather than rejecting them and distinguishing two dimensions of normative reasons may provide foundations for a more expressive hybrid consequentialist theory.

References

- Alexander, L. (2000). *Deontology at the Threshold*
- Alexander, L. (1987). "Scheffler on the independence of agent-centered prerogatives from agent-centered restrictions," *The Journal of Philosophy*, 84(5), 277-283
- Chappell, R.Y. (2023). "Beneficentrism," *Good Thoughts*
- Coady, C.A.J (2024). "The problem of dirty hands," *Stanford Encyclopedia of Philosophy*
- Formosa, P. and Sticker, M (2019). "Kant and the demandingness of the virtue of beneficence," *European Journal of Philosophy*, 27(3), 625-642.
- Gert, J. (2003). "Requiring and justifying: Two dimensions of normative strength," *Erkenntnis*, 59(1), 5-36
- Hooker, B. (2000). *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford University Press
- Howard, C. (2021). "Consequentialists must kill," *Ethics*, 131(4), 727-753
- Lazar, S. (2018). "Deontological decision theory and agent-centered options"
- MacAskill, W., Bykvist, K., and Ord, T (2020). *Moral Uncertainty*. Oxford University Press
- Muñoz, D. and Pummer, T (2022). "Supererogation and conditional obligation," *Philosophical Studies*, 179(5), 1429-1443
- Nebel, J.M (2022). "Conservatism about the valuable," *Australasian Journal of Philosophy*, 100(1), 180-194
- Ord, T. (2008). "How to be a consequentialist about everything"
- Portmore, D.W. (forthcoming). "Consequentializing," *Stanford Encyclopedia of Philosophy*
- Portmore, D.W. (2023). "Consequentializing agent-centered restrictions: A Kantse-quentialist approach," *Analytic Philosophy*, 64(4), 443-467
- Scheffler, S. (1985). "Agent-centred restrictions, rationality, and the virtues"
- Scheffler, S. (1994). *The Rejection of Consequentialism*, Oxford University Press
- Setiya, K (2018). "Must consequentialists kill?" *Journal of Philosophy*, 115(2), 92-105.
- Singer, P (1972). *Famine, Affluence, and Morality*. Oxford University Press.
- Smart, J.J.C. and Williams, B. (1973). *Utilitarianism: For and Against*, Cambridge University Press.
- Thornley, E (2022). "Is global consequentialism more expressive than act consequen-tialism?" *Analysis*, 82(1), 75-84.
- Utilitarianism.net (2023). "Elements and types of utilitarianism"