

# Consequentialism, Utilitarianism, and Pareto Efficiency: Ethics and Welfare<sup>1</sup>

Julian Fink  
julian.fink@uni-bayreuth.de

(1) Consider a company that sells heart defibrillators, i.e. a device implanted into the heart to stimulate it with electric jolts when necessary. Defibrillators usually bring huge benefits to those who need them. However, if faulty, they also carry medical risks. Suppose a producing company discovers that a large batch of their defibrillators carries a mechanic defect, increasing the risk of device failure about 1-2%. Suppose that in extreme cases this defect can cause the death of a patient. What should the producing company do? Should they inform the public and their customers about the product failure? Or should they refrain from spreading this information?

In a real-world case (matching the described features), a producing company – Guidant – defended their position *not* to inform the public as follows: the defect only slightly increases the risk of device failure. However, if a patient discovers that her implanted device carries the mentioned failure, she is likely to ask for a replacement. Replacement requires surgery. This surgery carries a higher risk of damaging the patients' health than the faulty defibrillators. So, patients and the public should not be informed. Even if a patient does not choose to replace a faulty defibrillator, she would still suffer from unbeneficial anxiety and fear, *etc.* Hence, again, the patients and the public *should* not be informed.

The 'should' here seems to express an ethical prescription. But what is the basis of the reasoning behind this should? How can it be defended?

Guidant's ethical and normative defense seems to be based on consequentialism. Expressed simplified, the idea behind the defense seems this: if informed, *more* people would die, and *more* people would suffer. This cannot be right. So, patients and the public should not be informed. Hence, if consequentialism represents a correct ethical view, then Guidant's decision not to spread the information seems ethically justifiable.

## (2) Consequentialism

Unlike with a role-based approach, consequentialism does not refer to essential aims and norms that govern social roles in determining ethical obligations. Unlike a Kantian approach, consequentialism does not ask for the universalisation of the intention (maxim) behind the action. Instead, it simply asks what are the *actual* consequences of your action; and how do they compare to the consequences of the alternative actions available.

---

<sup>1</sup> This lecture is primarily based on chap. 5 of Sanbu's 'Just Business: Arguments in Business Ethics'.

Consequentialism describes an entire family of ethical theories. However, all consequentialist theories share the following premise: what makes a state of affairs ethically valuable comes prior to what is right. That is, what *makes it the case* that an act is right (or what *explains* an act's rightness) can be determined independently of what is ethically right.

Most consequentialist accounts are *maximizing* in nature. That is, roughly, they say that an act is morally right if and only if its consequences are *best* (i.e. have the highest ethically relevant value). We could also imagine a *non-maximizing* version of consequentialism. For example, we could say that an act is morally right if and only if its consequences reach a certain threshold level of morally relevant value.

Consequentialism says that only the morally relevant value of an act's outcome (i.e. the goodness of an outcome) and *nothing else* is relevant for the ethical rightness of an action. Schematically, this can be put as follows:

*Consequentialism schematically.* For all acts A: A is morally permissible if and only if A's consequences bring about at least the equal amount of morally relevant value as all available alternative acts.

However, there seems to be an immediate knockdown argument against consequentialism. Consequentialism may simply not be *practical*. We have no reliable way of telling which of the available acts will lead to at least as much morally relevant value as its alternatives?

But a solution seems available. Ethical theorists distinguish between *objective* and *subjective* rightness. Accordingly, an act A is *objectively* right if and only if A has *in fact* the best consequences. In contrast: A is *subjectively* right if and only if A has – *as far as we can know and expect* – the best consequences. What is objectively right may be of philosophical interest; however what we practically ought to do is determined by subjective rightness. Hence, subjective consequentialism may give rise to what we practically ought to do.

*Subjective consequentialism:* For all acts A: A is practically morally right if and only if – *given our rational beliefs and expectations* - A's consequences will bring about at least the equal amount of morally relevant value as all available alternative acts.

### (3) Utilitarianism

For consequentialism to be action guiding (i.e. to enable a person to make a practical decision), consequentialism needs to specify (at least) which properties of an outcome or an act's consequence are *ethically relevant values*. In sum, a consequentialist theory needs to state *what* we should maximize. Here is a list of possible contenders: (i) quantitative pleasure/happiness, (ii) qualitative pleasure/happiness, (iii) ideal goods (e.g. friendship, freedom, beauty, etc.), (iv) absence of pain, (v) preference satisfaction, *etc.*

In the following, we will concentrate on the most widely discussed version of consequentialism, i.e. *utilitarianism*. Utilitarianism can be characterized by pointing at two essential features of utilitarianism. Utilitarianism is a (i) *welfarist* type of consequentialism, and it (ii) *uses additive aggregation*.

Ad (i). A utilitarian theory is consequentialistic: it only considers the *consequences* of an act when it comes to evaluating if an act is right or wrong. The outcome matters, not the act itself. In particular, what matters for a utilitarian is how much *welfare* or *well-being* an act brings about (compared, of course, with the welfare and well-being the available alternative acts will lead to). In its classical formulation, utilitarians see an act as right if it ensures 'the greatest happiness for the greatest number'.

Put simply, utilitarianism is a well-being calculation. Suppose you have to choose between two possible actions,  $A_1$  and  $A_2$ . You want to find it if it is  $A_1$ , or  $A_2$  or both are morally permissible. What you need to establish is whether  $A_1$  leads to an higher expected well-being than  $A_2$ .

*Welfareism* (i.e. the view that we should maximize well-being or welfare) may seem obvious to many. However, one may think of other things than well-being as being morally relevant. One thing that may be relevant are *lives*. Suppose you believe that *avoiding* deaths is morally valuable even if the lives saved do not come with positive well-being (i.e. 'happy' lives). Or consider values that involve a way of *self-realization*. Suppose an artist aims at producing art at the cost of sacrificing his happiness and well-being. We may think of this as adding to the moral value of an outcome, even though it produces negative well-being.

Utilitarianism is logically distinct from consequentialism. One can be a consequentialist without being a utilitarian. That is why even by accepting consequentialism, the utilitarian position still needs to be justified against possible alternatives.

Ad (ii). Whose welfare are utilitarians concerned with? The idea is simple: *everyone's* well-being matters, and it matters *equally*. As a guide to our actions, utilitarianism thus tells us to take into account how the available actions will affect the well-being of *each* affected person (or, perhaps even more broadly, every 'sentient being'). In fact, ethical actions aim choosing the action that brings about the greatest total net gain of well-being when all affected beings are considered.

For example: take a manager who faces the choice between using profits to raise salaries for the workers or, alternatively, to pay higher dividends to shareholders. The two actions clearly affect the two groups' well-being differently. Utilitarianism thus says that the company should resolve this trade-off by finding the wage/dividends distribution that *maximizes* the sum of total individual well-being of both workers and shareholders. Consequently, utilitarianism is premised on *additive aggregation* of individual welfare.

Compare also *act vs. rule utilitarianism* (Sanbu, pp. 74-6).

Two distinct problems of utilitarianism: (i) equality and (ii) basic rights.

#### (4) 'Efficiency' and the Pareto Criterion

As mentioned, on a theoretical level, utilitarianism demands individual well-being to be added up across individuals. On a practical level, it requires that we act 'efficiently' in the sense that we should not let any potential aggregate happiness go to waste. This presupposes that we can say – in principle – that an action increases some people's well-being more than it reduces someone else's well-being.

But can such a judgment ever be justified? In extreme cases, the answer may be 'yes': consider the addition of well-being when €20 are used to cure someone's blindness compared to the effect €20 have when added to the wealth of a millionaire. But in less stark cases, such judgments lack even pragmatic justification.

Additive aggregation thus makes many ethical theorists suspicious of utilitarianism. Some consequentialists thus turn to alternative theories that make no use of additive aggregation. One interesting alternative can be based upon the *Pareto criterion*:

*Pareto superiority.* A state of affairs X is better than another state of affairs Y if and only if nobody is worse off (has less well-being) in X than in Y, and at least one person is better off (has more well-being) in X than in Y.

Utilitarianism requires one's acts to maximize the total sum of well-being. It does so even if thereby *some* people are made much worse off. In contrast, if using *Pareto superiority* as a moral criterion, the *Pareto superiority* only requires actions that make everyone better off (or that make some people better off without making anyone worse off). So, if an act results in a state of affairs where at least one person is better off without anyone else being worse off, this act leads to a *Pareto improvement*.

The *Pareto criterion* also encapsulates a version of efficiency. But this differs from the kind of efficiency utilitarianism demands. Utilitarianism says: do not let any well-being go to waste. An act is ethically impermissible if there is an alternative act that could bring about more total well-being. Pareto superiority, in contrast, says that an act is ethically impermissible if every single individual could be made better off (without making anyone worse off). Defined formally, this sense of efficiency can be defined as follows:

*Pareto efficiency.* A state of affairs is Pareto efficient if and only if there is no other available state of affairs that is Pareto superior.




FIGURE 1

Suppose in figure 1, the two axis represent the well-being of a person. Each point (A to E) thus represents a distribution of well-being between two people. The axes measure the amount of well-being each person has in a state. If we assume that we can add up the person well-being, what would a utilitarian and a Pareto-criterion based ethics recommend?