

## Utilitarianism and the Moral Significance of an Individual

James Cain

Oklahoma State University

A common criticism of classical hedonistic utilitarianism is that it does not fully capture the moral importance of the individual. Often this criticism is framed in terms of examples in which the utilitarian seems forced to endorse the violation of an individual's rights in order to bring about a greater quantity of good, e.g., an innocent person is intentionally executed in order to prevent rioting, murder and mayhem. A utilitarian is liable to respond to such objections in the following way:

Such things as justice, the individual, and rights are of secondary importance in ethics. Of fundamental importance is the utility of actions. Upholding justice, showing concern for the individual, and respecting rights *tends* to increase utility and *to that extent* such actions are morally justified. But justice, rights, and the individual play such a large and emotional role in our thinking about moral issues that we tend to overlook the fact that their moral importance is wholly grounded in considerations of utility. Thus we mistakenly think we are able to come up with counterexamples to utilitarianism.

Given the complexity of any philosophically adequate theory of justice, rights and the individual, it may be difficult to see how to resolve this impasse to the satisfaction of

those on both sides of the debate. I will approach the problem of whether utilitarianism properly respects the individual in a rather different way, one that appeals to considerations that are, so to speak, on the utilitarian's home turf. In doing so I hope to make some headway in bringing the utilitarian to see that the role of the individual in ethics cannot be *reduced* to mere consideration of quantities of utility.

I will develop my argument by way of a couple of examples. The examples are unusual in that the first involves a consideration of infinite utilities, and the second involves a consideration of finite utilities in a situation in which one must choose between actions that have infinite expected utilities.<sup>1</sup> In Section One, an example involving infinite utilities will be constructed in which (a) one action is clearly morally preferable to another, (b) classical utilitarianism can give no grounds for holding that action to be morally preferable, and (c) the comparative moral worth of the pair of actions can only be properly assessed by considering how utility is distributed among individuals. Such cases are highly problematic for classical utilitarianism as they show that utilitarianism fails to provide adequate criteria for evaluating actions. Yet a utilitarian may feel that arguments involving infinite utility are inconclusive or unreliable because our grasp of the infinite in ethics is not reliable. In Section Two I will address this concern by showing how cases involving finite utilities can be constructed which approximate the difficulties found in the infinite cases. I will present an example in which an action will generate less overall utility than its alternative and yet it is morally preferable to perform the action. Throughout the discussion the examples I use will be rather artificial so that the mathematics can be kept simple and taken in easily.<sup>2</sup>

## *Section One*

Consider the following scenario.<sup>3</sup> Imagine that Elm Street is infinitely long. For each positive integer  $n$ , there is a single person that lives at the address  $n$  Elm Street and that person, person  $n$ , is immortal. It has just been discovered that everyone on Elm Street has a rather peculiar disease. If untreated this disease manifests no symptoms for a certain period of time after which one suffers an unending headache. What is especially peculiar about the disease is the timetable for the appearance of the headache. At the beginning of year one, the first person on Elm Street, person 1, comes down with a headache. At the beginning of year two the next two people on Elm Street get a headache; the following year the next four people get headaches; then eight people; then sixteen, and so on as the headache progresses down the street. Each year the number coming down with headaches doubles. For our purposes the important thing to note and remember is that *each year the number of people who begin to get a headache that year is precisely one greater than the number who already have a headache.*<sup>4</sup> It will be convenient to suppose that the amount of disutility caused by a year of headache is a constant; we will refer to it a ‘headache-year’ of suffering.<sup>5</sup>

Before year one begins you have a choice to perform one of two mutually exclusive actions, each of which would prevent some of the suffering the people on Elm Street would otherwise have to endure. Action X delays the onset of each person’s headache by one year, but once the headache begins it still continues forever. Action Y brings it about that the headache lasts for only one year, but does not delay the onset of the headache.

Which action ought you to perform? Clearly the answer is action Y. You should perform the action that relieves each person from an eternity of pain rather than a single year of pain even though the relief comes a year later. Perhaps surprisingly if we take utilitarianism simply to be the doctrine that we ought to maximize utility, then utilitarianism does not give us a criterion which enables us to choose action Y over action X.

We see this as follows. Both X and Y will result in the prevention an infinite number of headache-years of pain.<sup>6</sup> If we look at it on a year by year basis, then each year action X prevents a finite positive number of headache-years of pain.<sup>7</sup> It does this for an infinite number of years and thus has infinite utility. The same reasoning holds for action Y, except that action Y does not begin to prevent pain until the second year<sup>8</sup>; nonetheless in each year after the first, action Y prevents a finite positive number of headache-years of pain and thus over the course of all time Y will prevent an infinite number of headache-years of pain.

During any given year action X prevents more pain than action Y; in fact each year X prevents one more headache-year of pain than does Y. This follows from the fact we saw earlier that if neither X nor Y were performed then from year one onward the number of people experiencing their first year of headache would always exceed by one the number who have had a headache for over a year. Thus after the passage of  $n$  years, action X would have prevented  $n$  more headache-years of suffering than action Y would prevent. This might make it look like action X results in greater total utility than action Y. On the other hand the fact that for everyone on Elm Street action X only prevents one headache-year of pain and action Y prevents an infinite number of headache-years of pain

may seem to show that action Y results in a higher total utility. The fact is that if we are merely comparing the number of headache years prevented by actions X and Y, there is no absolute basis for saying that one prevents more headache-years of pain than the other. Both prevent  $\aleph_0$  (aleph null) headache-years of pain.

I can now formulate the 'infinite utility' objection. Classical utilitarianism does not give an adequate account of the moral importance of the individual since:

(1) Mere consideration of the quantity of utility generated by X and by Y is insufficient to determine whether action X or action Y is preferable.

(2) Action Y is in fact preferable to action X, and this follows from a consideration of both the utility of actions X and Y and features of the way the utility is distributed among individuals.

Thus, (3) the role of the individual in determining the morality of actions cannot be reduced to considerations of total utility of actions.

Thus we see that classical utilitarianism by itself cannot fully account for the moral importance of the individual.

### *Section Two*

Given that the moral importance of the individual is not exhausted by its role in increasing total utility, the question arises whether features of the individual can be

brought to bear in such a way that an action can be seen to be preferable though it has *less* utility than an alternative. I believe that my example can be modified in such a way that considerations similar to those we just looked at can be used to give us a case in which the preferred action will definitely result in less total utility than the available alternative action. Furthermore, it can be done in a way that only involves a consideration of finite utilities.

In the last example we imagined that the only treatment for the disease stems from the performance of either action X or action Y. Imagine the situation is just as before except that an additional treatment will be given for the disease, and this treatment will eventually eradicate the disease though an indeterminate number of years will pass before the treatment has its effect and until that time it is only our action X or Y that can help relieve the suffering caused by the disease. How this additional treatment works is a bit mysterious, but the part that concerns us is the following. At the beginning of year one a special container will be opened and a single atom of Rubidium<sup>87</sup>, which has a half-life of 47 billion years, will be dropped in. At the end of the year in which this atom finally decays the disease will be cured and any headache resulting from the disease will end. We will assume that our choice between X and Y has no bearing on when the atom finally decays. As before, we have the question, which is preferable: action X or action Y?

The utilitarian calculation tells us that action X is preferable to action Y. So long as the disease is not yet cured, each year action X will prevent one more headache-year of suffering than Y prevents. So if  $k$  years pass before the disease is cured then action X will prevent  $k$  more headache-years of suffering than action Y. Since the disease will

only last for some finite, though indeterminate, number of years actions X and Y will each bring about a finite amount of utility and action X's utility will exceed action Y's utility.<sup>9</sup>

Nonetheless it seems to me that Y is by far the preferable course of action. This seems so obvious as to not need an argument, but it may be useful to lay out a couple of considerations. First note that the *expected utility* of each individual on Elm Street with respect to his or her own headache will be many billion times greater with action Y than with X. We may see this as follows. Consider person 1. Action X gives person 1 an expected utility of one since without action X he will have a headache throughout year one regardless of whether the rubidium atom decays during the first year, and action X saves him from this year of headache and only this year of headache. Action Y on the other hand has an expected utility of many billion headache years since there is a 50% probability that his disease will last for at least another 47 billion years. So person 1 has a higher expected utility given action Y. Now consider any other person on Elm Street, say person  $n$ . Given this person's address on Elm Street there will be a particular year, say year  $i$ , in which she is destined to begin having a headache if neither X nor Y is performed and the disease has not yet been cured. Furthermore, since actions X and Y have no bearing on the time at which the disease is cured there will be a fixed probability  $p$  that the disease will still be uncured at the beginning of year  $i$ . If year  $i$  arrives and the disease hasn't yet been cured, then *at the beginning of year  $i$*  this person will have the same expected utilities as person 1 did *at the beginning of year one*. So to obtain person  $n$ 's expected utility *at the beginning of year one* with respect to actions X and Y we simply multiply person 1's expected utilities at the beginning of year one by the

probability  $p$ . Since for person 1 the expected utility of Y is billions of times greater than for X, it follows that the same will hold for person  $n$ .

Second, let's consider the way the headaches are potentially to be distributed among the individuals if the disease continues for a long duration. Most likely the disease will continue for many billion years. To get something of a feel for the distribution of pain among the residents of Elm Street, let's arbitrarily suppose that the disease lasts for just a billion years and consider the consequences of actions X and Y. Whether X or Y is performed a huge number of people will be affected, just one short of  $2^{1,000,000,000}$  (two to the billionth power). I'll refer to this group as *the population*. If action Y is performed everyone in the population will suffer one headache-year of pain. If action X is performed then over the course of the billion years there will have been a billion fewer total headache-years of suffering. About one half of the population will not suffer any headache.<sup>10</sup> About one quarter of the population suffers one year of headache<sup>11</sup>; about one eighth suffers two years of headache<sup>12</sup>; about one sixteenth suffers three headache years<sup>13</sup>; about a thirty-second suffers four years of headache<sup>14</sup>, and so on. Huge numbers of people will suffer for extraordinarily long periods. Person 1 suffers one year less than a billion years of headache. Two people, persons 2 and 3, each fall two years short of having a billion years headache. Four people, persons 4 through 7, each fall three years short of having a billion headache years of pain. Over a billion people fall thirty-one years short of suffering a billion year headache.<sup>15</sup> Over a trillion people fall forty-one years short of having a billion year headache, etc.<sup>16</sup> Over a googol people each fall 335 years short of suffering a billion year headache.<sup>17</sup> The picture that emerges is that to prevent a fairly modest amount pain (one headache-year of pain per year) a substantial



number of individuals would have to bear an enormous amount of suffering. Clearly in this case the potential distributions of suffering among individuals affected by actions X and Y shows us that Y is preferable to X despite the fact that X will have a greater total utility than Y. The problem now confronting utilitarianism is a form of the distribution problem, but it brings out the same issues raised by the case involving infinite utilities.

By way of summary and conclusion, there has been a long-standing objection to classical utilitarianism that it fails to give a proper account of the moral significance of the individual. I concur with this objection. I have attempted to bring forward arguments for this view that I hope will be of some weight for a person who is inclined to accept classical utilitarianism. I have tried to show by my first example that the importance of the individual cannot be *reduced* to considerations that merely treat of amounts of utility. The second example was formulated to bring out the fact that the role of the individual in our moral considerations may lead to an action's being preferable even though it will result in less utility. Because utilitarianism does not capture the moral significance of the individual it does not have adequate resources to give the correct moral assessment of the first example and it gives the wrong moral assessment of the second.

---

<sup>1</sup> For discussions of the infinite utility problem see: Krister Segerberg, "A Neglected Family of Aggregation Problems in Ethics," *Nous* 10 (1976): 221-44; Mark Nelson, "Utilitarian Eschatology," *American Philosophical Quarterly* 28 (1991): 339-47; Peter Vallentyne, "Utilitarianism and Infinite Utility," *Australasian Journal of Philosophy* 71 (1993): 212-17; James Cain, 'Infinite Utility,' *Australasian Journal of Philosophy* 73 (1995) pp. 401-404; Peter Vallentyne, "Infinite Utility: Anonymity and Person-Centeredness," *Australasian Journal of Philosophy* 73 (1995): 413-20; and Peter Vallentyne and Shelly Kagan, "Infinite Value and Finitely Additive Value Theory," *The Journal of Philosophy* XCIV (1997): 5-26. For more recent discussions, see the following papers and the references contained in them: Donniell Fishkind, Joel David Hamkins and Barbara Montero, "New Inconsistencies in Infinite Utilitarianism: Is Every World Good, Bad or Neutral?," *Australasian Journal of Philosophy* 80 (2002): 178-90; and Luc Lauwers and Peter Vallentyne, "Infinite Utilitarianism: More is Always Better," presented at the Econometrics Workgroup associated with the Department of Economics of the Katholieke Universiteit

---

Leuven (K.U.Leuven) – the paper is available at  
[http://www.econ.kuleuven.ac.be/ew/academic/econmetr/agenda/a0203\\_1.HTM](http://www.econ.kuleuven.ac.be/ew/academic/econmetr/agenda/a0203_1.HTM).

<sup>2</sup> For those interested in seeing a bit more of the mathematics I will provide footnotes. One does *not* need to read the footnotes to take in the paper.

<sup>3</sup> This example is a modification of some examples developed in my 'Infinite Utility,' *Australasian Journal of Philosophy* 73 (1995) pp. 401-404.

<sup>4</sup> In year one this is trivially true. For  $n \geq 1$ , in year  $n + 1$  there will be  $2^n$  people who begin to get a headache and there will be  $2^0 + 2^1 + \dots + 2^{n-1}$  people who already have a headache.  $2^n$  is one greater than  $2^0 + 2^1 + \dots + 2^{n-1}$ .

<sup>5</sup> Of course there may be forms of suffering other than headache pain that are relevant here. However, to keep the discussion simple I will only deal with headache pain. If we want to include other forms of suffering (e.g., the suffering involved in anticipating future pain) we can vary the example in the following way: let the intensity of the headache diminish to the extent that some other form of suffering is present so that the total suffering associated with a year of headache is constant.

<sup>6</sup> To be exact,  $\aleph_0$  (aleph null) headache-years of pain.

<sup>7</sup> In the  $n$ th year, action X prevents  $2^{n-1}$  headache years of pain during that year, and the cumulative total for all years up through the  $n$ th year on action X will be  $2^n - 1$  headache years of pain prevention.

<sup>8</sup> In the  $n$ th year, action Y prevents  $2^{n-1} - 1$  headache years of pain during that year, and the cumulative total for all years up through the  $n$ th year on action Y will be  $2^n - (n + 1)$  headache years of pain prevention.

<sup>9</sup> Here I am ignoring the worry that there is an infinitesimally small chance that the rubidium atom will never decay and thus X and Y may each have infinite utility. This problem can be taken care of in the following way. At the end of year one if the rubidium atom has not yet decayed the disease will go into temporary remission. If the atom eventually decays, after, say,  $k$  more years, then the headache will progress down Elm Street for the next  $k$  years as set out in the example. On the other hand, if the atom never decays then the headache never resumes—and so action X prevents one headache-year of pain (since it prevents a headache during year one) and action Y prevents zero headache-years (since it didn't prevent a headache the first year). Thus even if the atom never decays action X will have a greater finite utility than action Y. Using this modification all my other arguments will still hold.

<sup>10</sup> To be precise,  $2^{999,999,999}$  people.

<sup>11</sup>  $2^{999,999,998}$  people.

<sup>12</sup>  $2^{999,999,997}$  people.

<sup>13</sup>  $2^{999,999,996}$  people.

<sup>14</sup>  $2^{999,999,995}$  people.

<sup>15</sup> At the end of year 31,  $2^{30}$  people (over a billion people) begin to have a headache that lasts to year one billion.

<sup>16</sup> At the end of year 41,  $2^{40}$  people (over a trillion people) begin to have a headache that lasts to year one billion.

---

<sup>17</sup> At the end of year 335,  $2^{334}$  people (well over a googol, i.e.,  $10^{100}$ ) people begin their headache.