

Ethics Offsets

Posted on January 4, 2015 by Scott Alexander



I

Some people buy voluntary carbon offsets. Suppose they worry about global warming and would feel bad taking a long unnecessary plane trip that pollutes the atmosphere. So instead of not doing it, they take the plane trip, then pay for some environmental organization to clean up an amount of carbon equal to or greater than the amount of carbon they emitted. They're happy because they got their trip, future generations are happy because the atmosphere is cleaner, everyone wins.

We can generalize this to ethics offsets. Suppose you really want to visit an oppressive dictatorial country so you can see the [beautiful tourist sights there](#). But you worry that by going there and spending money, you're propping up the dictatorship. So you take your trip, but you also donate some money to opposition groups and humanitarian groups opposing the dictatorship and helping its victims, at an amount such that you are confident that the oppressed people of the country would prefer you take both actions (visit + donate) than that you take neither action.

I know I didn't come up with this concept, but I'm having trouble finding out who did, so no link for now.

A recent post, [Nobody Is Perfect, Everything Is Commensurable](#), suggests that if you are averse to activism but still feel you have an obligation to improve the world, you can discharge that obligation by giving to charity. This is not quite an ethics offset – it's not exchanging a transgression for a donation so much as saying that a donation is a better way of helping than the thing you were worried about transgressing against anyway – but it's certainly pretty similar.

As far as I can tell, the simplest cases here are 100% legit. I can't imagine anyone saying "You may not take that plane flight you want, even if you donate so much to the environment that in the end it cleans up twice as much carbon dioxide as you produced. You must sit around at home, feeling bored and lonely, and letting the atmosphere be more polluted than if you had made your donation".

But here are two cases I am less certain about.



Suppose you feel some obligation to be a vegetarian – either because you believe animal suffering is bad, or you have enough moral uncertainty around the topic for the ethical calculus to come

out against. Is it acceptable to continue eating animals, but also donate money to animal rights charities?

A simple example: you eat meat, but also donate money to a group lobbying for cage-free eggs. You are confident that if chickens could think and vote, the average chicken would prefer a world in which you did both these things to a world in which you did neither. This seems to me much like the cases above.

A harder example. You eat meat, but also donate money to a group that convinces people to become vegetarian. Jeff Kaufman and Brian Tomasik [suggest](#) that about \$10 to \$50 is enough to make one person become vegetarian for one year by sponsoring what are apparently very convincing advertisements.

Eating meat is definitely worth \$1000 per year for me. So if I donate \$1000 to vegetarian advertising, then eat meat, I'm helping turn between twenty and a hundred people vegetarian for a year, and helping twenty to one hundred times as many animals as I would be by becoming vegetarian myself. Clearly this is an excellent deal for me and an excellent deal for animals.

But I still can't help feeling like there's something really wrong here. It's not just the low price of convincing people – even if I was 100% guaranteed that the calculations were right, I'd still feel just as weird. Part of it is a sense of duping others – would they be as eager to become vegetarian if they knew the ads that convinced them were sponsored by meat-eaters?

Maybe! Suppose we go to all of the people convinced by the ads, tell them “I paid for that ad that convinced you, and I still eat meat. Now what?” They answer “Well, I double-checked the facts in the ad and they’re all true. That you eat meat doesn’t make anything in the advertisement one bit less convincing. So I’m going to stay vegetarian.” Now what? Am I off the hook?

A second objection: universalizability. If *everyone* decides to solve animal suffering by throwing money at advertisers, there is no one left to advertise to and nothing gets solved. You just end up with a world where 100% of ads on TVs, in newspapers, and online are about becoming vegetarian, and everyone watches them and says “Well, I’m doing my part! I’m paying for these ads!”

Counter-objection: At that point, no one will be able to say with a straight face that every \$50 spent on ads converts one person to vegetarianism. If I follow the maxim “Either be vegetarian, or donate enough money to be 90% sure I am converting at least two other people to vegetarianism”, this maxim *does* universalize, since after animal suffering ads have saturated a certain percent of the population, no one can be 90% sure of convincing anyone else.

As far as I can tell, this is weird but ethical.



The second troublesome case is a little more gruesome.

Current estimates suggest that \$3340 worth of donations to global health causes saves, on average, one life.

Let us be excruciatingly cautious and include a two-order-of-magnitude margin of error. At \$334,000, we are *super duper sure* we are saving at least one life.

So. Say I'm a millionaire with a spare \$334,000, and there's a guy I *really* don't like...

Okay, fine. Get the irrelevant objections out of the way first and establish the least convenient possible world. I'm a criminal mastermind, it'll be the perfect crime, and there's zero chance I'll go to jail. I can make it look completely natural, like a heart attack or something, so I'm not going to terrorize the city or waste police time and resources. The guy's not supporting a family and doesn't have any friends who will be heartbroken at his death. There's no political aspect to my grudge, so this isn't going to silence the enemies of the rich or anything like that. I myself have a terminal disease, and so the damage that I inflict upon my own soul with the act – or however it is Leah always phrases it – will perish with me immediately afterwards. There is no God, or if there is one He respects ethics offsets when you get to the Pearly Gates.

Or you know what? Don't get the irrelevant objections out of the way. We can offset those too. The police will waste a lot of time investigating the murder? Maybe I'm very rich and I can make a big anonymous donation to the local police force that will more than compensate them for their trouble and allow them to hire extra offi-

cers to take up the slack. The local citizens will be scared there's a killer on the loose? They'll forget all about it once they learn taxes have been cut to zero percent thanks to an anonymous donation to the city government from a local tycoon.

Even what seems to me the most desperate and problematic objection – that maybe the malarial Africans saved by global health charities have lives that are in some qualitative way just not as valuable as those of happy First World citizens contributing to the global economy – can be fixed. If I've got enough money, a few hundred thousand to a million ought to be able to save the life of a local person in no way distinguishable from my victim. Heck, since this is a hypothetical problem and I have infinite money, why not save *ten* local people?

The best I can do here is to say that I am crossing a [Schelling fence](#) which might also be crossed by people who will be less scrupulous in making sure their offsets are in order. But perhaps I could offset that too. Also, we could assume I will never tell anybody. Also, anyone can just go murder someone right now without offsetting, so we're not exactly talking about a big temptation for the unscrupulous.