

Manufacturing Reasons (20190503)

Alonzo Fyfe

I. The Humean Theory of Reasons

After some discussion, the Stanford Encyclopedia of Philosophy (SEP) entry, “Reasons for Action: Internal vs. External” settles on the following as the Humean Theory of Reasons (Revised):

HTR (revised): If there is a reason for someone to do something, then she must have some desire that would be served by her doing it, which is the source of her reason (Finlay, Stephen and Schroeder, 2017).

This is false.

It is not false in the sense of “Thus, I refute Hume.”¹ It is false in the sense of, “No Humean should agree to this”. It can easily be proved false on Humean grounds.

Imagine a world in which there are two people: Agent1 and Agent2. Each has only one desire - an aversion to their own pain. With this, a Humean would say that Agent1 has a reason to prevent the realization of any state in which Agent1 is in pain and Agent2 has a reason to prevent the realization of any state in which Agent2 is in pain. Consequently, there are two reasons. Agent1 has a reason, and Agent2 has a reason.

¹ I am also not asserting that this is false in the sense that, “Hume never said such a thing”. He may not have. The Humean Thesis is typically taken to be a thesis inspired by Hume that takes desires to be the source of all reasons. What I write will be consistent with the Humean thesis understood in this way.

From this, we can infer that the antecedent, “There is a reason for Agent1 to refrain from realizing a state in which Agent2 is in pain”, is true. The source of that reason is Agent2’s aversion to pain. At the same time, the consequent, “Agent1 must have some desire that would be served by refraining from realizing a state in which Agent2 is in pain” is false. Agent1 has no such desire.

What a Humean should agree to is this:

(1) If there is a reason for someone to do something, then there must be a desire that would be served by her doing it, which is the source of that reason; and

(2) if an agent has a reason to do something, then that agent must have some desire that would be served by her doing it, which is the source of her reason.

This version pairs up “there is a reason” with “there is a desire” and “has a reason” with “has a desire”. It explicitly denies any necessary direct connection from “there is a reason” to “has a desire.” At the same time, it allows for the construction of artificial or manufactured reasons, which I will discuss in the second half of this paper. I will also argue that moral reasons are manufactured reasons – or, more precisely, reasons that there reasons to manufacture. However, before I get to that, I wish to consider some points that might cause some resistance to this interpretation of the Humean theory.

Some may accuse me of merely playing with words. In Humean terms, Agent2’s aversion to her own pain clearly gives her a reason to prevent P (P = “Agent2 is in pain”). However, it is an abuse of language to say that it counts as a reason for Agent1 to prevent P. I wish to argue that this is not the case.

II. The Misinterpretation Objections

When I say that there is a reason for Agent1 to prevent P, I am literally saying nothing more than that there is a reason that exists to prevent P (in this case, grounded on Agent2’s aversion to her own pain),

and that Agent1 can act to prevent P. If the objection states, “Yes, but this does not imply that Agent1 has a reason to prevent P,” my answer is that this is true. In fact, that is the point. “Agent1 has a reason to prevent P” has been deliberately made false.

Common Practice

The first step that I wish to make in defense of this use of language is to point out that we routinely distinguish between “there is” and “he/she/it has” and “I/you have”. For example, “there is 150 million ounces of gold in Fort Knox” does not imply “I have 150 million ounces of gold in Fort Knox.” The statement “there is something” leaves open the question of who (if anybody) has that something. So, in this case, I am saying that the phrase “there is a reason” leaves open the answer to the question of who has that reason.

A critic will then likely respond that I am ignoring the phrase “for Agent1” in the claim. This reason is a reason “for Agent1” to act in some way. With the introduction of this phrase, the question of whose reason we are talking about is no longer open.

I will admit that there is some ambiguity to the phrase, “There is a reason for Agent1 to do X.” It can be read in one of two ways:

(1) *(There is a reason for Agent1) / (to do X)*

(2) *(There is a reason) / (for Agent1 to do X)*

If a person has interpretation (1) in mind, then the objection is sound. “There is a reason for Agent1” does, indeed, imply “Agent1 has a reason”. Yet, the existence of interpretation (1) does not rule out the concurrent possibility of interpretation (2). The second interpretation is used in claims like, “There is a reason for having the party on the Friday,” or “There is a reason to build a high speed railroad from New

York to Los Angeles". A person acknowledge that these reasons exists without implying that they are reasons that she has.

An alternative way to express (2) is: "There is a reason for Agent1 doing X". This can be true, even though Agent1 has no reason to do X.

Bernard Williams' Distinction

Bernard Williams (1981) saw this distinction when he wrote about the difference between internal and external reasons. However, he did not fully appreciate it. He reported that the phrase "has a reason" tends to be used when an agent "has some motive which will be served or furthered" by performing the action in question. At the same time, he wrote that the phrase "there is a reason" tends to be used when "the reason-sentence will not be falsified by the absence of an appropriate motive." This accurately describes using "has a reason" in the sense of "has a desire" and "there is a reason" in the sense "there is a desire".

Williams then said, "it would be wrong to suggest that either form of words admits of only one of the interpretations". Using the formula I am defending, "There is a reason" is compatible with "has a motivation" when the reason that "there is" is a reason that the agent has. Whether the agent can have a reason without having a motive may be open to some dispute, but this is something I would deny.

Though Williams argued against the existence of external reasons, he must allow for a certain type of external reasons – those being the reasons that motivate other agents. When Agent1 is aware of the fact that Agent2 has an aversion to being in pain, Agent1 knows that Agent2 has a reason to avoid being in pain. This is, from Agent1's perspective, an external reason. However, it is not a reason that implies anything about Agent1's motivation or rationality (which Williams would say makes this type of external reason irrelevant to his argument).

Mark Schroeder and “Fact of the Matter”

Mark Schroeder (2007) offers a different analysis of “there is a reason” that points to the facts of the matter that are relevant to fulfilling the desires of the agent who has a reason to be concerned. He wrote:

For R to be a reason for X to do A is for there to be some p such that X has a desire whose object is p, and the truth of R is part of what explains why X's doing A promotes.

So, “there is a reason” means “there is an R” that is true in some state of affairs S, and the truth of R explains why Agent1, in performing the action, fulfills the desire that Agent1 has. For example, if there is dancing at the party, then there is reason for Ronnie to go to the party, given that Ronnie has a desire to dance. The claim that “there is a reason for Ronnie to go to the party” would be false if not for the fact that Ronnie has a desire to dance.

This may be a perfectly legitimate use of the term “has a reason”. However, this does not exhaust the options. We are free to use the phrase this way while also using it in the way I have described, the correct sense of the term coming from the context in which it is used. Assume that Arnold wants to go to the party and needs a chaperone. Ronnie can do this. Now, there is a reason for Ronnie to go to the party in the sense I have described even though he has no desire that would be served by his going.

The Mackie Maneuver

If none of this works and one is still wanting to insist that “there is a reason for Agent1 to do X” implies “Agent1 has a reason to do X,” there is one more response I can give. Assume that we have locked “there is a reason” so that it is just another way of saying “has a reason”. The set of propositions that I have identified with the alternative “there is a reason” is still real. We have reasons to discuss those

relationships and need a term for them. We already have “has a reason” to use when we want to say that an agent “has a reason”. We do not need two terms that mean the same thing. So, I will appropriate “there is a reason” for my purposes.

J. L. Mackie (1977) used the example of the atom in defending a similar move. He wished to change the meaning of “good” from “having objective, intrinsic prescriptivity” to “is such as to fulfill the desires in question”. To defend this, he pointed out that scientists once defined “atom” to mean “without parts”. At the same time, they had been using this term to refer to the smallest bits of a matter recognized as such. When they came to suspect that these smallest bits of an element had parts (electrons, neutrons, and protons) scientists did not respond by becoming atomic anti-realists. They shifted the meaning of “atom” so that it still referred to the smallest bits of an element recognized as such, but allowed that those bits could have parts (Mackie, 1977, Loc. 1456-1459).

Simon Blackburn (1985) argued that Mackie should have introduced a new moral language, using the term “schmoral” in place of “moral” to reflect what would be an error-free morality. Mackie did not need to do this any more than chemists needed to introduce a term “schmatom” as a substitute for “atom” in claiming that what the term “atom” had been referring to had parts.

I need a term to refer to states of affairs like those in which (1) Agent2 has an aversion to P – thus Agent2 has a reason to prevent states of affairs in which P is true, and (2) Agent1 has the ability to act so as to prevent such a state of affair. The language I choose to adopt is, “there is a reason for Agent1 to prevent P” grounded on Agent 2’s desire, though it says nothing about Agent1’s motivation or rationality.

Moral Reasons

At this point one may ask, “Why care about these things? What use do we have for *there is a reason* understood in this way?”

The answer is that moral reasons are reasons of the “there is a reason” variety. There is a moral reason for Agent1 to realize states of affairs in which Agent2 is not in pain. Agent1 may not have a reason to do so, but such a reason exists, grounded in Agent2’s aversion to pain. The proposition, “There is a moral reason for Agent1 to do X” does not imply that Agent1 has a reason to do X, and cannot be falsified by showing that Agent1 lacks a reason to do X. I will argue in the next section that “there is a moral reason for Agent1 to do X” implies that there are many and strong reasons to give everybody – including Agent1 – a reason to do X.

The SEP describes how their interpretation of the Humean Theory of Reasons leads to a problem that, in that entry, is called “the central problem of morality” (Finlay, Stephen and Schroeder, 2017). To generate this problem, we start with the original Humean Theory of Reasons, which says that (in their example) there is a reason for Hitler not to order genocide only if he has some desire that would be served by not ordering genocide. We combine this with a principle that states, “an action (like ordering genocide) is morally wrong for an agent (like Hitler) only if there is a reason for him not to do it.” This appears to be inconsistent with another widely accepted thesis: the wrongness of Hitler ordering genocide does not depend on whether Hitler has a desire that would be served by not ordering genocide.

If we substitute the original interpretation of the Humean Theory of Reasons with the one I propose it remains true that it is wrong for an agent (like Hitler) to order genocide only if there is a reason for him not to do it. There are, in fact, a great many and strong reasons for him not to do it. However, those

reasons are reasons that “there are”, not reasons that “Hitler has”. Thus, the moral wrongness does not depend on the whether or not Hitler has a desire that would be served by not ordering genocide. It depends on whether there are desires that would be served by not ordering genocide – which, indeed, there are.

This, of course, will bring up all sorts of questions such as, “Does the fact that I have a reason to have pizza for supper mean that you are morally obligated to provide me with a pizza for supper?” Or, “Does the fact that Hitler have a reason to order the Jews to the gas chambers imply that the Jews have an obligation to go to the gas chambers?” This suggests some potential absurdities coming from the thesis. To deal with these absurdities, we have to look at what the thesis does imply, which I will do in the next section.

III. Manufacturing and Distributing Reasons

In the original simple model, Agent2 has a reason to prevent P, where P = “Agent2 is in pain”. This is grounded on Agent2’s aversion to pain. This implies that there is a reason for Agent1 to prevent P.

This also implies that Agent2 has a reason to manufacture a reason for Agent1 to prevent P and to give Agent1 that reason. Agent2 has a few options: (1) incentive/deterrence, (2) socializing.

Incentives and Deterrence

Agent2 can give Agent1 a reason to act in ways that prevent P by creating incentives or deterrence.

In our simple model, Agent1 also has an aversion to his own pain. That is to say, Agent1 has a reason to prevent Q where Q = “Agent1 is in pain”. So, Agent2 can make Agent1 an offer. “Agent1, if you prevent P, I will prevent Q.” Agent1 now has a newly manufactured reason to prevent P, one that is grounded on his own desire to prevent Q. Alternatively, Agent2 can say, “Agent1, if you fail to prevent P, I will bring about Q.” Again, Agent1’s reason to prevent Q now becomes a reason to also prevent P.

We are, of course, constantly involved in manufacturing and distributing these types of reasons. When we buy and sell things we manufacture the first type of reason: “If you give me a bottle of aspirin for my headache I will give you some money you can use to fulfill some desire of yours.” Criminal law is substantially built on the second type of offer. “If you take somebody else’s property without their consent, then we will realize a state of affairs in which you are spending 3 to 5 years in prison”.

Manufacturing reasons in the form of incentives and deterrence comes with some serious limitations. If Agent1 can avoid the penalties or obtain the rewards without performing the action, then the promised rewards and punishments fail to give Agent1 any reason to act. This suggests that there would be an advantage to having an invisible friend who (1) knows everything so that nobody could hide a wrongful act from it, (2) has the power to give everybody a reason to do what there are reasons for that person to do, and (3) was disposed by its character to give people sufficiently strong reasons to do that which there are reasons for them to do. Unfortunately, the absence of such a being means that it is up to us to create these reasons.

Fortunately, there is another way to manufacture reasons.

Modifying Desires

Let us assume that, by praising those who avoid causing pain to others and by condemning those who cause pain to others, Agent2 can create in others an aversion to causing pain. Others, in this case, do not refrain from causing pain as a means to acquiring praise and avoiding condemnation. Rather, they form a some “desire that I not cause pain”, which then becomes a reason not to cause pain. Within our simple model, Agent2 has a reason to praise those who refrain from causing pain and condemn those who do not refrain from causing pain, and Agent1 acquires a reason to refrain from causing Agent2 pain as a result. Since Agent1 also has an aversion to pain, Agent1 also has a reason to praise those who

refrain from causing pain and to condemn those who do not. If we had a community of agents like Agent1 and Agent2, everybody would have a reason to praise those who refrain from causing pain and condemning those who do not so as to promote, universally, an aversion to causing pain to others.

This method works among human beings. Our brains contain a system that processes rewards and punishments – including praise and condemnation – in such a way so as to generate new and to modify existing desires and aversions. What is originally valued for its ability to generate a reward (e.g., praise) or avoid punishment (e.g., condemnation) comes to be valued for its own sake. These desires and aversions become reasons that those who acquire the desires and aversions then have. People generally have both the motive, and the means, to manufacture and distribute these reasons.

Fortunately, this method of generating reasons works on people other than those who are rewarded or punished (praised or condemned). It appears to have an effect on others who witness the reward or punishment, and even on those who merely hear about it. In fact, the rewards and punishments one hears about need not have really occurred. That praise or condemnation might be delivered to a character in a story or parable or the condemnation of a hypothetical person who would perform a hypothetical act.

This method of manufacturing and distributing reasons also has its limitations. First, it is a slow process – one cannot manufacture and distribute these types of reasons on the fly, as one can with incentives and deterrence. Second, it is imprecise. We cannot, for example, use praise and condemnation to create an aversion to causing pain except on the third Thursday following a full moon in autumn. Third, it requires a great deal of cooperation within a community, since conflicting messages are not going to have much of an effect.

Among the reasons that there are reasons to manufacture and distribute there are some that people generally have many and strong reasons to promote universally – among the population at large. It seems reasonable to expect that this would include such things as a universal desire to help those in need, a desire to repay a debt, an aversion to breaking promises, an aversion to taking the property of others without their consent, aversions to violently assaulting or killing the innocent, and the like.

We can contrast these reasons with other reasons that people generally have little or no reason to promote universally – reasons governing such things as what to eat, where to live, what profession to go into, which books to read, what to wear, and what to do for entertainment. On these matters, people have little to no reason to promote universal interests. In some cases, people generally have reasons to promote a variety of interests, such as in the selection of a profession, where we benefit from people having a variety of reasons to take on a variety of tasks.

Just as with incentives and deterrence, manufacturing and distributing reasons in the ways described here is another common practice.

The Central Problem Reconsidered

With the possibility of manufacturing reasons in mind, I would like to revisit “the central problem of morality” that I originally discussed at the end of Section II. That problem involved squaring the Humean Theory of Reasons with the principle that for it to be wrong for Agent1 to do X, there must be a reason for Agent1 to do X, and that reason for Agent1 to do X must be independent of Agent1’s desires. When combined with the original Humean Theory of Reasons, this generated a conflict since, on that version of the theory, for there to be a reason for Agent1 to do X, Agent1 must have a desire that is served by doing X.

I made “there is a moral reason for Agent1 to do X” independent of the desires that Agent1 has by making it dependent on the desires that there are – the desires that other agents have. Here, I argued that the reasons “there are” are reasons to give to Agent1 a reason to do X. In a sense, the moral reasons for doing X are not grounded on the desires Agent1 has, but on the desires that Agent1 “should have” (in the sense that they are the desires that others have reason to give him).

So, does every desire that exist create an obligation on the part of others?

Answer: No. Moral reasons are limited to the subset of reasons that people generally have many and strong reasons to promote universally. A person may want pizza for supper, and that person has a reason to give others a reason to deliver a pizza to his door (e.g., by agreeing to pay for the delivery of such a pizza). However, it is not the case that people generally have many and strong reasons to promote, universally, a desire to deliver pizzas.

There are many reasons for people to enter into certain occupations, and even reasons to encourage more people to enter into certain professions where society would benefit from more practitioners. However, people generally have no reason to promote, universally, an interest in any particular profession. Consequently, what profession to enter into remains largely in the category of moral permissibility. Hit man, art thief, and grifter tend not be on the “permissible” list, but those exceptions are consistent with people generally having many and strong reasons to promote, universally, an aversion to key features of those occupations.

Was Hitler wrong to order the Holocaust?

There certainly are many and strong reasons to give people, universally, an aversion to genocide. This means that there are many and strong reasons to condemn – and to condemn in particularly harsh terms – people who order genocide, even as we talk about them as historical figures. Those reasons are

also reasons to condemn hypothetical agents of genocide and fictional agents – characters made the villains of stories.

Are there other cases where this account would produce counter-intuitive results?

It is a mistake to underestimate the power of human imagination. However, investigating those options is outside of the scope of this paper. We would have to look at cases such as that where Agent1 desires that P and Agent2 desires that not-P, where a substantial number of people desire to torture a single individual, and whether or not it makes sense to cause people to desire slavery. These are interesting cases worthy of further investigation, but that will have to wait for a different opportunity.

IV. Conclusion

Even within a Humean theory of reasons – where all of the reasons that exist are grounded on desires that exist – it is not the case that, when there is a reason for somebody to do something, that the person has a reason to do that thing. Sometimes, the reason that exists belong to somebody else. When the reasons that exist are not those of the agent, they do not provide the agent with motivation, nor do they figure into the agent's rationality. However, they are reasons that others have to give the agent a reason to do that thing. They can give the agent a reason by offering the agent an incentive or threatening deterrence. They can use praise and condemnation to manufacture within the agent (and others) a desire or aversion that would then be the agent's reason to do that thing. There are some reasons that people generally have a great many and strong reasons to manufacture and distribute universally – reasons to keep promises, to repay debts, to help those in need, and to refrain from assaulting and murdering the innocent. In these cases, to say that there are reasons for the agent to perform or forebear from performing some act is to say that there are reasons to praise or condemn those who would perform such an act, thus creating the reasons that there are reasons to create. We can have all

of this, all without setting a foot anywhere near even the edge of what is traditionally known as Humean theories of reason and action.

References

Blackburn, Simon (1985). "Errors and the Phenomology of Value". In Ted Honderich (ed.), *Morality and Objectivity: A Tribute to J.L. Mackie*. Oxford University Press. pp. 1-22.

Finlay, Stephen and Schroeder, Mark, (2017). "Reasons for Action: Internal vs. External", *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2017/entries/reasons-internal-external/>.

Hume, David, (2010a) *A Treatise of Human Nature*, Project Gutenberg, <http://www.gutenberg.org/files/4705/4705-h/4705-h.htm>.

-----, (2010b). *An Inquiry Concerning the Principles of Morals*, Project Gutenberg, <http://www.gutenberg.org/files/4320/4320-h/4320-h.htm>.

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Penguin, UK.

Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press, UK.

Williams, Bernard (1981). "Internal and External Reasons", in *Moral Luck: Philosophical Papers 1973–1980*. Cambridge University Press.

