An Assignment Theory of Desire (20180715)

Alonzo Fyfe

1. Introduction

> *The right act is the act that a person with good malleable desires and lacking bad desires would have done in those circumstances.*[1]

Ultimately, this is the thesis that I would like to defend.

I am not going to get there in this paper. I will only be making the first step. However, it may help you to understand this step if you have some idea of the destination.

I expect that many potential readers will immediately reject the destination. "That's typical. He has *right act* on the left side of the equation and *good desire* on the right. No doubt I am going to find all sorts of question-begging and viciously circular claims as he inserts into *good desires* exactly the ideas he needs for *right action*."

That is a legitimate concern. It is true, I need an account of *good malleable desire* that does not beg any questions.

Of course, an account *good malleable desire* must begin by answering the question, "What is a desire?" Then we can look about the ways in which desires are malleable and what makes them good (or bad).

So, this is that first step. What is a desire?

I suggest the following hypothesis:

> *A desire is a propositional attitude - a mental state that that can be expressed in the form 'Agent desires that p', where 'p' is a proposition capable of being true or false, which assigns a value V to the proposition p representing the importance to Agent of p being made or kept true.*[2]

What does this mean?

It means that, if you want to go back to bed and get some sleep then you have a mental attitude to the proposition, "I am in bed and sleeping," that assigns to this proposition a certain value V representing how important it is to you that "I am in bed sleeping" be made or kept true. The more important this is

---

[1] This formulation is non-accidentally similar to that which Rosalind Hursthouse presents: "An action is right iff it is what a virtuous agent would characteristically (i.e. acting in character) do in the circumstances." Hursthouse, 1999. This account of "good desire" will differ significantly from Hursthouse's Aristotelian concept of "virtue", but many of Hursthouse's arguments are applicable here, *mutatis mutandis*.

[2] I have toyed with the idea that both beliefs and desires assign a value to a proposition being true. The value that it assigns to a proposition that is the object of a belief - V(B) - represents the credence of that belief - the likelihood given whatever else the agent knows that the proposition is, in fact true. Whereas, as I said, the value it assigns to a proposition that is the object of a desires - V(D) - represents the importance to that agent of that proposition being made or kept true. However, I have not given the issue of belief in this regard a lot of thought. I toss that out here simply as something to think about.

to you, the more power the desire has to override other concerns such as getting the house cleaned up for company this evening or getting that philosophy paper written.

In order to explain this idea a bit more clearly, please allow me an opportunity to present a short glossary of terms and phrases that I will be using describing how I intend to use them. I also want to include some assumptions that I will be making – since a paper such as this cannot be expected to cover everything under the sun.

2. Glossary

Humean or Non-Humean

Many of those who are familiar with the philosophical debate concerning desires may try to understand this in terms of a classic debate between "Humeans" and "Non-Humeans". Labels such as these are often useful. When we identify somebody as Utilitarian or Kantian, Muslim or Atheist, Marxist or Randian, it tells us a lot about a person in a phrase. These terms allow people to quickly identify a person with a set of beliefs in a large set of claims and counter-claims.

Yet, these heuristics sometimes fail, and this is an example.

A Humean theory of desire is generally one that conforms to Hume's claim, "reason is the slave of the passions" and that it is "not contrary to reason to prefer the destruction of the whole world to the scratching of my finger" (Hume, 1896). Basically, beliefs and desires are two different kinds of things. Reason or rationality apply to beliefs. There is no rationality of desires.

In the philosophy of desires, the dominant view among those who identify as "Humeans" propose a "dispositional" theory of desires. Dispositionalism holds that desires are dispositional states – dispositions to act so as to realize certain ends or goals or to engage in certain expressions or to experience certain emotions (Alvarez, 2017).

Standing up against dispositionalism, we see a diverse sent of non-Humean philosophers who hold that there is something to be said for the rationality of desires. These are known as "evaluativists". The paradigm example of an evaluativist theory holds that a desire is a form of belief. Specifically, a "desire that p" is "a belief that p is good". Other versions say that a desire that p is to be understood as perceiving (as opposed to believing) it to be the case that p is good. Other versions suggest that the agent believes or perceives that p "ought to be".

Since beliefs can be held to be reasonable or unreasonable, and perceptions can be accurate or inaccurate, desires can also be reasonable or accurate or fitting (Oddie, 2005, 2017).

I consider myself to be a Humean insofar as I think that "reason is the slave of the passions" and, in a sense, it is "not contrary to reason to prefer the destruction of the whole world to the scratching of my finger."

At the same time, the hypothesis that I gave above is not a dispositional theory of desire. It explains the desire that p in terms of p having a value. Such a theory would appear to be more at home among the evaluativists.

This would imply that I am a Humean evaluativist. This is not going to fit well within the standard categories.

What I will argue reconciles these two positions is that the value assigned to p in virtue of a desire that p is subject to rational requirements. The desire that p is neither accurate (or inaccurate) or fitting. The value that p has in virtue of a desire that p is something that, at the most basic level, cannot be mistaken

because the desire that p gives it its value. It is not contrary to reason to assign a higher negative value to p where p = "there being a scratch on one's finger" than to q where q = "the destruction of the whole world".

Since I will be presenting a Humean/evaluativist theory, somebody assuming attempts to pigeon-hole this account as either a Humean/dispositional or a non-Humean/evaluativist account is likely to be confused.

Perhaps the confusion is mine. I hope to show that this is not the case when I discuss dispositional theories in Section 3 and non-Humean evaluativist theories in Section 4.

Realizing P

Another issue that will be important in what follows concerns the question of what desires aim at. I will be building on the suggestion that a desire that p aims at realizing p. By "realizing" I am saying that the desire assigns a value to it being the case that 'p' is real – meaning that it is true. If p is false and the agent desires that p, then the agent will seek to make it the case that 'p' is true. If 'p' is already true and the agent desires that p, then the agent has a reason to preserve p – to keep it real.

One way to illustrate what this means is by looking at Robert Nozick's (1974) experience machine thought experiment. Nozick asked readers to imagine a machine that will feed the brain with images – nerve stimulations – identical to those that would come from the apparent fulfillment of an agent's desires. So, if one wanted to aid the poor and starving in the most poverty-stricken parts of the world, the experience machine, like some virtual reality device, will give the subject a perfectly believable (and believed) experience of aiding the poor and starving in the most poverty-stricken parts of the world.

Nozick then asked us to consider whether we would want to live our lives inside of the machine. A refusal to enter the machine implied a preference for real-world experiences. We would rather risk the failure of perceiving our desires as being fulfilled in the real world, then to live with the illusion of them being fulfilled in the experience machine.

The idea that the objects of our desire being the realization of that which is desired says that the value of the experience machine is found in whether it can make true that which the agent desires. If the agent desires that she have the experience of aiding the poor and starving in the most poverty-stricken parts of the world, then the machine has the capacity to make true the propositions and she will find life in the machine tempting. However, if she desires to actually aid the poor and starving in the most poverty-stricken parts of the world, then the machine has nothing to offer her.

Both agents are seeking to make true the propositions that are the objects of their desire. The difference between them is that they want different things. One of them wants the subjective experience, the other wants to help people.

Desire that p, Desire to φ

There is a convention in this discipline that says that desires can be expressed in the form "agent desires that p" or "agent desires to φ."

The symbol 'p' in the first phrase, as mentioned above, represents a proposition – the meaning of a sentence, which is capable of being true or false. Examples may include, "I desire that I be accepted into the PhD program," or "I desire that my favorite sports team win the championship," or "I desire that my obnoxious neighbor comes down with a dreadful disease so that I can watch him suffer."

The symbol 'φ' in the second phrase represents some action. "I desire to stay home and cuddle with my wife," or "I desire to clean up the house before company arrives this evening," or "I desire to get this philosophy paper written by the end of the month."

That is not to say that these are the only types of desire claims we see in the English Language. We may see expressions such as "I want a beer" or "I would like to have somebody honest run for public office for a change." These do not directly fit into the two models given above. However, they can be rewritten so as to fit the model. "I want a beer" becomes "I desire to drink a beer" or "I desire that I am drinking a beer." Similarly, "I would like to have somebody honest run for public office for a change" becomes "I desire that somebody honest runs for public office."

In fact, the convention states that, in the end, even "desire that φ" statements can be translated into "desire that p" statements. I have already demonstrated this where I equated "I desire to have a beer" with "I desire that I am drinking a beer." It is in virtue of this possibility that I expressed the desire hypothesis at the start of this paper exclusively in terms of a "desire that p". We can rewrite all other types of desire into an expression having this form.

Intrinsic vs. Instrumental Desires

Another convention among those who write in this field is to distinguish what are called "intrinsic desires" from "instrumental desires".

Another conventional way of talking about these two concepts is to talk about ends (or goals, or objectives, or that which is valued for its own sake) and means.

I fear that the term "intrinsic desires" is an invitation to making a serious mistake. The term invites us to relate "intrinsic desires" to "intrinsic values" – a type of "objective, intrinsic prescriptivity" that J.L. Mackie (1977) and Sharon Smith (2006) warned us against and which I hold does not exist.

Many who use this term do not have "intrinsic values" in mind. They are using the term "intrinsic desires" for any type of end (as opposed to means) of human intentional action regardless of what form it might take. However, the author's intentions do not eliminate the ease at which a leader can slide from "intrinsic desires" to "objective, intrinsic values."

For the reason, I prefer the term "final desires" or "end desires" to "intrinsic desires." Here, I am simply specifying that these desires are at the end of the desire chain. The identify the agent's final goal or ultimate purpose. They identify that for the sake of which other things are ultimately done.

Within this same set of terms, I also want to endorse the tradition that an instrumental desire is not some different kind of thing from final desires and means-ends beliefs. Rather, it is a term we use to refer to bundles of final desires and means-ends beliefs. "I want some aspirin" states "I want to get rid of this pain and I believe that if I take some aspirin than the pain will go away." Many people do not actually want to go to work in the morning as a final end. Instead, they have a desire to have a comfortable and secure place to sleep at night and food to eat and a means-ends belief that going to work will result in receiving money that one can use to purchase the opportunity to have a comfortable and secure place to sleep and food to eat.

A few fortunate people actually do work that they value as a final end.

A part of the discussion to come will look at the hypothesis that a desire is a type of belief. One of the key questions we will be addressing here is whether our desires can be mistaken in the same way that our beliefs can be mistaken.

In addressing this question, it will be important to keep in mind that this question is only interesting when we ask it about final desires. Since instrumental desires are partially made up of means-ends beliefs, there is no controversy to the claim that we can be mistaken about instrumental desires. A person who "wants an 8mm wrench to tighten a bolt" might actually want a 7mm wrench. However, this does not imply that we can also be wrong about our ends.

While people generally focus on these issues specifically about talking about intrinsic (final) values and instrumental values, there are more relationships between objects of derivative evaluation and final ends. The peanut butter in a peanut butter cup does not have value in virtue of being a means to the end of having a peanut butter cup. It has value in virtue of being a part of the whole. The same can be said of any other ingredient in the peanut butter cup.

Just as is the case between means and ends, people can make mistakes about the relationship between parts and wholes. A person can be mistaken about whether that which is being evaluated is a part of the whole, or can be mistaken about whether the whole would continue to have its value with or without this part.

Nomy Arpaly and Timothy Schroeder (2014) make a case for realizer values. This is a state of affairs that realizes something that has final value. They use the example of a person who has as a final end his father's health. Such a person may be worried about his father shoveling snow because the effort may be bad for his health. As with instrumental value, the agent's belief about the relationship between the effort of shoveling snow and his father's health may change with a change of beliefs. He may discover that a bit of exercise would actually be good for his father's health.

Desires Narrowly vs. Broadly Defined

Sometimes the word "desire" is used to refer to all pro-attitudes. (Davidson, 1963). At other times the term is used to refer to a subset of pro-attitudes as distinguished from such things as wishes (Doring and Eker, 2017). This potentially leads to confusion as what is true of desires in the narrow sense need not be true of desires in the broad sense – and may well be that which distinguishes desire narrowly construed from desire broadly construed.

In this paper, I will be looking at "desires" in the broad sense. For example, I will not be distinguishing "desiring that p" from "wishing that p." Doring and Eker (2017) distinguished these on the grounds that desiring something means that one can do something to bring it about, whereas wishing that something is the case implies no action and may imply no action. Both of them are examples in which the agent, in some sense, wants it to be that p is the case. That makes them both desires so far as this paper is concerned.

Ultimately, what will distinguish what beliefs from desires is what Elizabeth Anscombe (1963) called "direction of fit". This is captured using the slogan, "If your beliefs disagree with the world, change your beliefs; but if your desires disagree with the world, change the world." I do not want to include this in this section on definitions because I want to devote some substantive discussion to this topic below. However, it is useful in more clearly understanding what I mean to be talking about in using the broader definition of "desire" – any attitude that has the world-to-mind direction of fit.

In using the term "desire" in this sense, I am seeking to approach this subject in the way consistent with those who speak about the "belief-desire" theory of intentional action – what Davidson (1963) and Smith (1993) refer to as the standard view. We could replace this with a "belief-desire-wish" theory, or add more complexities, but that will serve no purpose.

Here, I will bring forth the distinction from the previous section. It is certainly possible that a person can have an instrumental "desire that p" based on p's contribution to the fulfillment of one desire and "desire that not-p" based on the thwarting of a different desire. It is more problematic to imagine a person having a goal or end of p and a goal or end of not-p.

"There Is" vs. "She Has" a Reason

Another area where talk about desires is unnecessarily confusing is exemplified by the Stanford Encyclopedia of Philosophy entry on "Reasons for Action: Internal vs. External."

The entry provides this account of a revised version of the Humean Theory of Reasons:

> *[Humean Theory of Reasons] (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.[3]*

On a plausible interpretation of this, it is easily proved false, even within a Humean framework of desires.

Simply imagine a community made up of two people, Alph and Bet. Each person has an aversion to their own pain. Because of this aversion, each person has a reason to avoid their own pain.

In this community, two reasons exist. There is Alph's reason for avoiding his own pain, and Bet's reason for avoiding her own pain.

In this situation, there are conditions where the antecedent is true, but the consequent is false. Specifically, *there is a reason* for Bet to refrain from putting Alph in a state of pain - a reason that exists. It is Alph's reason to avoid pain, which is grounded on Alph's aversion to his own pain. At the same time, the claim that Bet has a desire that would be served by not putting Alph in a state of pain is false. Bet's only desire is her own aversion to her own pain. A true antecedent (there is a reason; Alph's aversion to his own pain) and false consequent (Bet has a desire that would be served) shows that the conditional is false.

In fact, given that the only desire that Bet has is to avoid her own pain, if faced with a choice between a slight pain on her finger and Alph's excruciating torture, she has a reason to choose to avoid the slight pain on her finger, and no reason to avoid Alph's excruciating torture.

One can suggest avoiding this problem by giving a different interpretation to "there is a reason for Bet to avoid causing Alph pain" to refer only to the reasons that Bet has. That is to say, we should really replace, "There is a reason for Beth to do something" with "Beth has a reason to do something". Using AHTR for Alternative Humean Theory of Reasons, this gives us:

> *[AHTR1] If someone has a reason to do something, then she must have some desire that would be served by her doing it, which is the source of her reason.[4]*

---

[3] Stanford Encyclopedia of Philosophy, "Reasons for Action: Internal vs. External", First published Thu Sep 4, 2008; substantive revision Fri Aug 18, 2017, 7/8/2018, https://plato.stanford.edu/entries/reasons-internal-external/.

[4] Stanford Encyclopedia of Philosophy, "Reasons for Action: Internal vs. External", First published Thu Sep 4, 2008; substantive revision Fri Aug 18, 2017, 7/8/2018, https://plato.stanford.edu/entries/reasons-internal-external/.

We can combine this with the following:

> *[AHTR2] If there is a reason for someone to do something, then there must be some desire that would be served by her doing it, which is the source of that reason*

In fact, this is just the improvement in the language that I want to propose.

This formulation matches "there is a reason" with "there is a desire" and "she has a reason" with "she has a desire," instead of matching "there is a reason" with "she has a desire".

Under this new formulation, it could still be the case that:

> *[AHTR3] 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 that reason.*

However, it is clearly the case that HTR1 and HTR2 can both be true, while AHTR3 is false. In fact, this is what I would argue that the example of Alph having a reason to avoid being in a state of pain (there is a reason) fails to imply that Bet has a desire that would be served by preventing Alph's pain (Bet has a desire that would be served) shows to be false. Only AHTR1 and AHTR2 are true.

There is a further question concerning what Hume would say if we confronted him with these options. While this is an interesting discussion, I fear that we have more and stronger reasons (grounded on more and stronger desires) to avoid this digression than to pursue it.

The part of the discussion where this new terminology will become most relevant is in the case of moral value. Michael Smith (1987) discussed a potential counter-example to his theory of desire that went as follows:

> *Suppose that I am standing on someone's foot so causing him pain, and that I know that this is what I am doing. Surely, we can imagine its being appropriate for an outsider to say that I have a reason to get off his foot even though I lacked the relevant desire, and, indeed, even if I desired to cause him pain.*

I am uncertain about the significance of being able to imagine its being appropriate for an outsider to say that Smith has such a reason even though he lacked a relevant desire. I would hold that it would be false. This outsider can say that there is a reason for him to get off of this person's foot – this being the victim's reason for avoiding pain. But this would not imply that Smith has a desire that would be served – and, *ex hypotheses*, he does not.

Please note that the person on whose foot Smith is standing has a reason to cause Smith to have a reason to get off of his foot. The unfortunate victim has a reason to threaten or bribe Smith, thereby creating a reason for Smith to get off of his foot that corresponds to the reason for avoiding the threatened action or obtaining the offered bribe. Indeed, if the victim had the power, he has a reason to cause Smith to have an aversion to causing him pain. We may, with appropriate caveats and conditions, translate these facts into the claim that Smith should have a reason not to stand on the victim's foot. Yet, none of this implies that Smith does have such a reason, or a desire that would be served.

The moral of this story is that I am wanting to avoid the potential confusion that comes from associating "There is a reason for someone to do something" with "She has a desire that would be served by her doing it." Instead, I am going to relate "there is a reason" to "there is a desire" and "she has a reason" to "she has a desire." I will leave the question of the link between "if there is a reason" and "she has a desire" for a future discussion.

What People Really Want

The last principle that I want to mention is that I will be focusing on what people really want, and not what they may believe will fulfill their desires. To explain this, I want to look at a couple of examples.

The Gin and Tonic Case

Smith presented a second case that the language of desire makes unnecessarily confusing.

In this case, he desires to drink a gin and tonic. Believing that the glass in front of him contains gin, he says that he wants to pour some tonic into it and drink it. However, the glass actually contains petrol. As an objection to Smith's dispositional account of desire (which I will get to later), Smith has a desire to pour tonic into the glass in front of him and drink the contents, he has the requisite beliefs, but he has no reason to pour the tonic into the glass and drink its contents.

The point that I want to address here is not he implications this has for Smith's theory, but his concession that it makes sense to say that he desires to pour gin into the glass containing petrol and drink the contents. In contrast, I would say that it would be reasonable to say that Smith thinks he has a such a desire. However, the claim that he has a desire, like the claim that the glass contains gin, is false

Imagine that Smith is with a companion. The companion sees Smith pour tonic into the glass and begin to raise it to his lips to drink the contents. She puts her hand over the top of the glass to stop him and says, "You don't want to do that."

Smith's response is unlikely to be, "Your statement is so clearly false that you obviously have no understanding of the English Language or of what it means to say that I want to drink what is in this glass." On the contrary, it would be perfectly reasonable for Smith to ask, "Why not?" In asking this question, he is admitting to the possibility that "I desire to drink the contents of this glass" is false – and its truth does not depend on his belief that the glass contained gin before he added the tonic. What matters is not his belief about what was in the glass, but what was actually in the glass. Smith is asking for evidence that it is true.

When Smith's companion provides him with convincing evidence, Smith could be expected to say, "You're right. I don't want to do that." This statement, "I don't want to do that," did not become true the instant he acquired the true beliefs about what was in the glass. It was true from the start. Meaning that the assumption that he had the requisite desire then he thought the glass contained gin was false from the start.

The problem with saying that the proposition that Smith wants to drink from the glass is true is that it tends to imply that, if he would drink from the glass, he would acquire something that he wants. However, this is false. We can avoid being seduced into such implications by saying that he does not want to drink from the glass, precisely because drinking from the glass will not yield anything he wants.

The point of this discussion is that "Agent has a desire that p" or "Agent has a desire to φ" in this discussion will be taken to mean that p being the case or φ-ing actually will fulfill the desires in question independent of what the agent believes.

The Picasso Case

Smith's Picasso case is similar to the gin and tonic case in that it concerns the agent's false beliefs. In this case, instead of falsely believing that p will fulfill his desire, Smith falsely believes it will not.

In this case, Smith wishes to purchase a Picasso painting. A stranger offers him an opportunity to purchase a Picasso at a quite reasonable price. However, Smith is suspicious and believes that the painting is not a genuine Picasso painting, so he refuses to buy it.

Does Smith want to purchase this painting?

Smith says, "No." The question of whether Smith wants to purchase the painting is determined by whether Smith believes that purchasing the painting will fulfill his desires.

I say, "Yes". The answer to whether Smith wants to purchase the painting is determined by whether purchasing the painting will fulfill his desires.

We are obviously talking about two different things here. This is not a dispute over which proposition is true and which false. This is a dispute over how to use the word. I wish to use the word "desire" to refer to that which the agent actually wants, not what he thinks he wants. If the agent desires that p or desires to φ then p or φ-ing is such as to fulfill the desires in question.

Smith's way of approaching the problem says:

> [T]he reason that I have to buy the painting in front of me is a normative reason. For it suffices for the truth of the claim that I have such a reason, that there is a requirement-in this case, in the broad sense, a requirement of rationality"-that I buy the painting in front of me.

More generally, Smith seeks to distinguish between motivating reasons (those that cause action) from normative reasons (those that justify action). On Smith's account, it makes no sense to say that he has a motivating reason to purchase the painting, since he wasn't motivated to purchase the painting. However, he does have a normative reason to purchase the painting. This normative reason is grounded on the fact that purchasing the painting would have fulfilled his desire to own a Picasso.

Remember, my objection is not that Smith is wrong, or that he fails to capture normal ways of speaking. My objection is that this way of speaking leads to confusion and error. Smith is saying that there is a rational requirement that calls for buying the painting. However, in refusing to purchase the painting, Smith may be acting in a perfectly rational manner. His desire to purchase a Picasso is not irrational, and his belief that this art seller is offering him a forgery may be perfect rational as well. Indeed, it may be more rational than believing that the painting is an original Picasso.

To explain my "yes" answer, I ask the question of whether Smith has a desire that would be fulfilled if he purchased the painting. He does. Therefore, he wants to purchase the painting.

Smith would say quite emphatically that he does not want to purchase the painting. However, non-coincidentally, he would also say, just as emphatically, that the painting is not a Picasso. Indeed, the strength of his emphasis that he does not want to purchase the painting is proportional to the strength of his conviction that he believes the painting to be a forgery. He is wrong about the painting. I would argue that he is just as wrong about his desire that he does not want to purchase it, and for the same reason.

I am not denying that Smith would say that he does not want to purchase the painting in these circumstances, or that his claim is not a sincere utterance. Instead, I am arguing that, in spite of what he would say, the claim that he does not want to buy the painting is false.

At this point, one can object that I am arguing for a significant change in our use of language. As evidence that this is not such a radical change, let us change the example slightly. The art seller convinces Smith that this is a genuine Picasso, so now Smith believes this. The dealer then invites Smith to play a game. The seller takes a black stone and a white stone, puts his hands behind his back, and puts a stone in each hand. Then, holding his closed hands out, asks Smith to choose one. If Smith

chooses the hand with the white stone, he gets the painting without paying for it. Otherwise, he pays the price of the painting and gets nothing.

Smith clearly wants to choose the hand with the white stone. He does not know which hand it is, but that is beside the point. The clear answer to the question, "Which hand do you want?" is "The one with the white stone." This answer does not look at what Smith believes, it looks at what will fulfill his desire to have a Picasso painting without paying for it. Here, the relationship between desire or want and what actually fulfills that desire becomes a bit more obvious.

Still, if one wants to insist that I am seeking to revise time-honored ways of talking about desire, so be it. In that case, I invoke the Pluto Principle.

The Pluto Proposal

In 2005, when astronomers discovered more objects in the Kuyper Belt about the same size as Pluto, they addressed the question of whether to call Pluto a planet. Specifically, they debated giving a new definition to the term 'planet' in much the same way that I may be interpreted as giving a new definition to the term 'desire'.

In the case of Pluto, astronomers applied the formula that like things should be classified with like. Pluto was more like the other Kuyper Belt Objects being discovered than the eight other planets. To make the language more efficient, they proposed that the definition of 'planet' be changed, and the International Astronomical Union took a vote.

We do not have an International Philosophical Union we can go to in order to argue for a new definition of desire. Instead, the best we can do is propose the reasons for the change to each other and see how many people take up the suggestion. This section of the paper may be read as making this type of appeal. Our current use of the term 'desire' is confusing. To reduce the confusion, the proposal before the committee is that 'Agent desires that p" or "Agent desires to φ" be understood to mean that p or φ-ing will actually fulfill the desires in question, regardless of what the agent believes.

Furthermore, I propose that we distinguish carefully between the thesis that "If there is a reason for Agent to φ, then Agent has a desire that would be served by her φ-ing" from "If Agent has a reason to φ, then Agent has a desire that would be served by her φ-ing".

This is the language that I would like to use in what follows.

3.   Dispostional Theories of Desire

As mentioned above, dispositional theories of value hold that a desire that p is a disposition to act so as to make it the case that p, and a desire to φ is a disposition to φ.

Michael Smith (1987) provides a classic example of dispositionalism.

> [A]ccording to this conception, we should think of the desire to φ as that state of a subject that grounds all sorts of his dispositions: like the disposition to φ in conditions C, the disposition to φ' in conditions C', and so on.

Sabine Doring and Bahadir Eker (2017) express this position as follows:

> Necessarily, for any agent a, any proposition p, any time t, and any act type φ, if, at t, a desires that p, then a is disposed at t to φ in circumstances where a takes her φ-ing to be conducive to p's being the case.

The Stanford Encyclopedia of Philosophy expresses the position as follows:

> *For an organism to desire p is for the organism to be disposed to take whatever actions it believes are likely to bring about p.*[5]

These accounts, which equate having a desire with having a disposition to act in a way that the agent believes will bring about a particular state of affairs, is one that we can test in two different ways. We can test it by seeing whether we can find any dispositions to act in a way that the agent believes will bring about a state of affairs that is not desired. And we can test it by seeing if there are things that are desired where the agent is not disposed to act in a particular way.

On example of the first type of test is that which Warren Quinn (1993) describes involving a man with a disposition to turn on radios even though turning on radios does not serve as a means to some valued end and is not something the agent values for its own sake. Without it being the case that the end is valued, according to Quinn, it is not desired.

Examples of the latter include cases where the agent desires something he could not be at all disposed to change because it is – and is known to be - beyond his power, such as wishing to have never been born, that pi was an irrational number, or anything that is already the case (since the agent cannot bring about that which has already been brought about) (Schroeder, 2004; Lauria, 2017).

I will discuss these types of cases shortly, but I want to say a quick word about what the dispositional theory gets right.

What Dispositional Theories Get Right

Before offering criticism, I want to focus some attention on what dispositional is gets right. A desire that p does dispose agents to act so as to realize p. This is true. It is also true that the wind is disposed to cause tree branches to sway and create waves on the water.

However, it is one thing to say that desires have these effects and desires are these effects. Once we know what the wind is, we can say something about how it causes branches to sway and waves to form. However, if we leave it as, "wind simply is the disposition for branches to sway and waves to form," our understanding of these phenomena remains limited.

Once we understand what desires are, we will see that they are things that dispose people to act in particular ways. Indeed, they are used to explain and predict intentional actions. They are the causes f the dispositions, not the dispositions themselves.

Radioman

The first case I want to discuss is Quinn's counter-example of Radioman, a person with a disposition to act, but who has no desire. Quinn describes Radioman as somebody who is simply disposed to turn on radios if there is one within reach that he discovers is not on. He has no interest in listening to music or in hearing any type of broadcast. In fact, he would prefer silence. He doesn't value having radios on in any way. It is simply that, as he passes buy such a radio, he will reach out and turn it on. He has, in this case, a disposition to act so as to turn on radios. However, he has no desire to do so. Consequently, having a disposition to act cannot be identified with having a desire – not if there are dispositions to act that are not desires.

Michael Smith responds to this by arguing that the case of Radioman is one that defies explanation, or that Radioman has a desire.

The first option asks us to find a way of explaining Radioman's behavior if it is not a desire. Why is he turning on the radios? We have no answer. It is not a muscular twitch or the effect of a drug or some other external cause. Absent an explanation such as these, we can do nothing but throw up our hands and give up any attempt to answer the question. It simply happens.

On the second option, Smith notes that the behavior is believe-moderated and has a particular state relative to other desires. It is belief-moderated in the sense that the agent must have a belief that there is a radio, that it is off, and in what he must do to turn the radio on. It is subject to the considerations of actual desires in that the agent may be persuaded not to turn the radio on by threatening to shoot him if he did so. This requires weighing the turning on of the radio against other considerations and finding the interest in turning on radios to be insufficiently strong. So, radioman has a desire, though a strange sort of desire that is not related to any view that what he desires to do has any value.

One of the reasons why I dislike thought experiments such as this is because it is easy to manipulate the results and build in the elements that serve the interests of the author. I would like to discuss a real-world case.

For 12 years, when I left work, I would turn right upon exiting the building to go to the bus station to catch the bus home. Then, the city built a new bus station to the left and closed the old bus station. For several days, when I left the building to go home, I turned right. I had a disposition to turn right when I left the building. If a desire is a disposition to act, then it seems to follow that I must have had a desire to turn right. However, this is false. I had no such desire. Thus, a desire cannot be accounted for in virtue of a disposition to act.

One cannot say that my turning right was not an intentional action. It was just as intentional on the first day that the bus station was closed as it was on the last day the bus station was opened. Nor could my behavior be attributed to a false belief since, if anybody would have asked me, I would have told them exactly how to get to the new bus station.

Indeed, there were a couple of times when I kept repeating to myself that the bus station was to the left and, while I was making sure that I was thinking of this fact, I nonetheless discovered after leaving the building that I had turned right. If somebody wants to claim that I did not turn right intentionally, I would again have to ask if I turned right intentionally the day before the bus station closed.

The fact is, I turned right not because I had a desire, but because I had a habit. Habits are dispositions to act that are distinct from desires, and the dispositional theory of desire cannot explain the difference between a desire and a habit. The actual difference is that the desire motivates action that aims to realize something of value, whereas the habit may or may not (merely contingently) is associated with something of value. We distinguish a good habit from a bad habit in virtue of whether it tends to bring about something we value – in whether it operates in the service of a desire.

We can conduct an experiment of this type at home. Reprogram your computer keyboard to switch the 'r' and the 't' keys. After doing so, if you are a reasonably competent typist, you will discover what when you seek to type 'the' you will instead type 'rhe' more often than not. Your act of typing 'rhe' is just as intentional as the act of typing 'the' was before you changed the keyboard. One can scarcely argue that you were thwarted by the false belief that the 'r' and 'h' keys were in their original position because you know that this is not the case. You are disposed to type 'rhe', but you desire to type 'the'. Thus, the desire is not equivalent to the disposition to act.

Again, the difference is that typing 'the' serves a desire, and typing 'rhe' – the act you will be disposed to perform until you lost the old habit and acquired a new one – does not. This 'act in service of an end that has value to the agent' distinguishes a desire from a disposition to act. However, when we assign a value to an end, we leave dispositional theories and enter the realm of evaluativist theories of desire.

Here, then, is a real world case of a disposition to act that is not a desire. Desires produce a disposition to act. The same is true of habits. The difference is that desires dispose people to act in ways that serve an end that has value, whereas the relationship between habits and ends that have value is merely contingent.

The Death of Desire Principle

Dispositional theories of desire have problem when there is nothing the agent can do to bring about that which, according to the dispositional theory, they are disposed to bring about. It seems to imply that it is impossible to desire something that it is impossible to bring about – or, at least, the agent thinks that it is impossible to bring about. Chief among these is the notion that a person cannot be disposed to bring about that which the agent believes is already true. Consequently, when somebody obtains that which they desire, when there is nothing they can do to bring about such a state, the dispositional theory says that the desire disappears.

Lauria (2017) calls this the "death of desire principle". He describes it as follows:

> *Imagine that Sam desires to see Niagara Falls. Mary offers to take him there. There they are, enjoying the breathtaking panorama. At some point, Sam says, "I want to see Niagara Falls." "Sam, you are seeing Niagara Falls," replies a quite surprised Mary. We understand Mary's astonishment. It is strange to express a desire to see something while in the midst of seeing it. Sam might express a desire to continue seeing the false, but this is a different desire than a desire simply to see the Falls. How could he desire simply to see the Falls while he is seeing them and is aware of his doing so? It appears that desire is incompatible with the representation that its content obtains.*

Dispositionalists seek to account for this "death of desire" principle by the fact that, if p is true, the agent cannot be disposed to act so as to realize p. While Sam is looking at the waterfall, Sam cannot be disposed to act to make it the case that he is looking at the waterfall. Therefore, he cannot desire to look at the waterfall. He can be disposed to act so as to continue to look at the water. Therefore, he can desire to continue looking at the waterfall. But, as Lauria states, this is a different desire.

Is it, really?

It seems quite odd to think that Sam had one desire – a desire to look at the waterfall – that motivated him to get into a car, drive to Niagara Falls, walk to the lookout where he could see the Falls, and – the very instant he sees it – his desire to see the Falls disappears. And then, at that very same instant, a different desire emerges that is a desire to continue to look at the falls.

What if that new desire did not emerge? Imagine Sam, walking up to the lookout, getting his first glimpse of the Falls, then turning around and leaving again. His desire to see the Falls is dead, and no desire to continue to see the Falls takes its place.

Or imagine the husband, wanting to cuddle with his wife, goes home with a plan to sit on the couch, invite her to sit beside him for a snuggle, and the instant the cuddling begins, he gets up and leaves so

that he can play a computer game. His desire to cuddle with his wife having died the instant the cuddling began, no different desire to continue cuddling with his wife arose to take its place.

I consider the person with a desire to eat. He gets groceries, fixes himself a fine meal, sits down, picks up a fork, and then gets up to watch television. His desire to eat, like Sam's desire to see the Falls, fails to survive the realization that he is eating, and no new desire to continue eating rising up to take its place.

The case of the desire to eat is particularly relevant. We know something about the biology of hunger. One of the factors relevant to the desire to eat is the level of ghrelin in the blood. High levels act on the brain to promote hunger. As one eats, the stomach expands, and ghrelin levels drop. The hunger that motivates the person to fix a meal to eat, motivates the agent to continue eating until full.

This, then, provides an alternative account of what happens in these scenarios. There is a single desire. There is a desire "that I see the Falls" or "that I cuddle with my wife" or "that I eat." This desire motivates the agent to act so as to realize the desired state, and then continues to motivate the agent to preserve that state until the desire fades as a result of natural causes. Evolution would favor our having desires that faded over time. Evolutionary fitness is not well served by having us permanently in a single maximally desired state, whether it is eating, having sex, relaxing in the shade, or cuddling with a spouse. We value it for a while, then it is time to move on to something else.

We can account for this oddness that Lauria mentions by noting that it would be just as odd for Sam to say, "I want to continue seeing the falls" before he even arrived. There is only one desire, but "I want to see the Falls" is the phrase we use for the context where one is not seeing the falls, and "I want to continue seeing the falls" is the phrase we use in the context where one is seeing the Falls.

It is like talking about "the car I will have" up until one buys a car and "the car I have" after buying it. It may be odd to talk about "the car I will have" when one owns the car, and "the car I have" before purchasing it. However, that does not imply that these are two different cars.

On this alternative account, the state "I am seeing the Falls" has value for the agent. It is important. This importance motivates the agent to make or keep the important proposition true. This gives the agent a disposition to act, both to realize the state of affairs, then to preserve it, because it has importance.

The dispositional theorist has a response to this. She can say, "I was wrong. There is a "desire that p" can survive the realization of p because there are still dispositions to act. Before p is realized, the agent is disposed to make p true. After p is realized, the agent is disposed to preserve p. The desire that p is still a disposition to act.

Unfortunately, this does not work in a case where that which is done cannot be undone. We can use Graham Oddie's (2017) example of Clinton desiring to be first female President of the United States. If she had won the election, and made it through the inauguration, she would have been the first female President. The dispositional theorist would have to say that her desire that she be the first female President dies at that instant. She could not desire to continue to be the first female President because there was nothing she could do to prevent it from being true.

However, there appears to be no reason to treat her desire "that she be the first female US President" any different from desire "that he see the falls" or my desire "that I cuddle with my wife". The fact that the former cannot be undone does not change the fact that she continues to value being the first female US president in the same way that Sam values continuing to look at the falls and I value continuing to cuddle with my wife.

Now, the dispositionist can respond that he is going to distinguish between desiring and valuing. Clinton can continue to value being the first female President but she can no longer desire to be the first female President, because this makes no sense. A desire requires at the least the possibility of action.

However, establishing that desire requires the possibility of action is not the same as establishing that desire is a disposition to act. At this point, we have "valuing" doing all the work. It provides the motivation where action is possible and the other symptoms of valuing where action is not possible. Everything that is important in a theory of desire now applies to a theory of valuing, with desiring being "valuing when the possibility of action exists". The dispositionalist can then say that he is going to distinguish between valuing and desiring. Clinton, perhaps, would have continued to value being the first female President, but could no longer desire it.

However, this would only be a semantic difference. It would have been the valuing that motivated action when it seemed that action could have an effect, and that would have persisted past the inauguration. "Desiring" is just a word used for valuing under conditions where relevant action is possible. At this point, we would need to ask whether it makes sense to have a dispositional theory of valuing.

Maria Alvarez (2017) may think that this is possible. She provides a broad definition of dispositions that goes beyond dispositions to act. It includes dispositions to express and to feel. These may be called "the other symptoms of valuing" that I spoke about above. She wants to argue for a dispositional theory of desire in this broad sense, suggesting that desires, in fact, are "manifestation-dependent". They must show themselves in some way, or they do not exist.

Broad Dispositional Theory

When Alvarez spoke of the dispositions of desire, she did not limit herself to actions she did not limit herself to dispositions to act in ways so as to realize what is desired. Behavior was just one of the dispositions of desire. She included changes in psychological states. A desire can manifest itself as frustration when attempts to realize the desire fail or one discovers that it is now beyond one's ability to obtain. Swearing or pounding one's fist on the table are actions that may not help to realize any particular end, but they are expressive of a desire nonetheless. Desires manipulate our thoughts and focus our attention. People on a diet can generally tell you that the desire to eat, even when it does not manifest itself as eating (particularly when it does not manifest itself as eating) has a number of other manifestations.

Alvarez could try to use these facts to salvage a dispositional state from the objections I raised to the death of desire principle above. The death of desire does require more than it being the case that there is nothing one can do to bring about that which is desired. Instead, an agent has a desire so long as it manifest itself in some way – not just in action that aims at bringing it about, but expression or effects on the agent's psychological states.

Above, I argued against the idea of distinguishing desiring from valuing, using the term "desiring" only where action is possible, and reserving the term "valuing" for other cases where the possibility of action has been removed but the other symptoms of valuing remain. Using this broad concept of disposition, Alvarez can argue that, so long as there are symptoms, there is desiring, where desiring can be understood as the disposition to create these symptoms. This handles counter-examples where dispositions to act are not possible, but the other manifestations of desiring remain.

Alvarez goes a bit beyond this and argues that desires are manifestation—dependent entities. There are no desires without at least one manifestation.

> *[T]here seem to be dispositions that are what might be called "manifestation-dependent": the absence of the manifestation over the lifetime of the object implies the absence of a disposition. . . My claim is that it is part of the concept of desire that someone has a desire at time t, only if the desire has been manifested in any of the various ways I described above at some point up to and including time t.*

She asserts that attributing to somebody the disposition of being a smoker if they never smokes, or of being generous when they have never given anything of value for the sake of others is not a matter of the person having the desire but it failing to manifest. It is a case of the agent not having the desire.

I believe that we have two reasons to reject this thesis.

First, it seems quite possible for a person to have a desire or aversion that never manifests itself. Imagine, if you will, a species of intelligent life that has a deep fear of spiders. Spiders on their home world are extremely deadly, and this strong aversion to spiders was such an aid to survival that their reaction to spiders was genetic. They acquired an advanced level of technology, where they killed all of the spiders. Then their civilization collapsed – due to war or perhaps some environmental catastrophe that resulted from killing all of the spiders. When their civilization emerged several millennia later, they had lost all knowledge of spiders.

We can plausibly suggest that these people, born into a spiderless world, still have a fear of spiders. We can test our hypothesis by bringing a spider to this world and dropping it into the middle of a birthday party and seeing how the natives react. If they are repulsed by the sight of the spider, we can conclude that they still have this aversion to spiders – and that they have this aversion even if we do not drop a spider into the middle of this birthday party. Indeed, this innate aversion to spiders is what explains their reaction when we drop the spider into the middle of the birthday party.

In response to this type of counter-example, Alvarez writes:

> *First, if the object of desire is something the agent was not at all familiar with, then it is wholly implausible to suggest that the agent's reaction is a manifestation of a desire that was there all along.*

It is precisely this statement that the case of the spiderless world refutes. Here is a case in which the agents are entirely unfamiliar with spiders. However, it is written into their genetic code that the instant they were to perceive a spider, they would perceive it as something to get away from. The aversion would then manifest itself in the form of fleeing behavior.

We can say the same thing of a child who, by good fortune, has never experienced pain – not even once. Alvarez would have us believe that this child has no aversion to pain. Yet, we know that the instant the child is cut or burned or slapped the first time, that the child will have an adverse reaction to the sensation that results. This first experience of pain did not create the aversion.

Second, how would a desire manifest itself the first time if there is no desire before the first manifestation? What caused the first manifestation? If the child mentioned above did not have an aversion to pain until its first manifestation, then what was it that caused this first manifestation? The standard explanation is that humans are wired for an aversion to pain well before its first manifestation, and that it is the presence of this wiring that explains and predicts our first reaction to pain. It has to exist before the first manifestation, or there will not be a first manifestation.

We can ask other questions about this account. For example, is a single manifestation enough? Do we say that an agent has a desire that p if, for one fleeting moment, he had a manifestation and then it

ceased? If a desire can go extinct, when does it go extinct? Does it go extinct with its last manifestation, or when it can on longer be triggered? Indeed, does a desire exist between manifestations? Or does it go extinct and then re-emerge with each new manifestation?

For my purposes here, I do not need to address these questions. It is sufficient that a desire is a disposition for certain types of behavior, expression, or psychological state to result from certain circumstances. The individuals of the alien species described above have an aversion to spiders just in case they are disposed to have an aversive reaction in the presence of spiders, even though they never see a spider. A child has an aversion to pain just in case if something caused the child to experience pain the child will not like it.

This will still yield a response to the death of desire principle. A desire does not need to involve a disposition to act so as to realize that which is desired. It can also involve a disposition to engage in expressive actions or to feel or think certain things. In this way, it is possible for a dispositional theory of desire to explain how a desire can survive when there is no more opportunity to act.

The problem with this move is that it is so broad it is applicable to virtually anything we can know about. From black holes to God to muons to free will, we can say of anything that it does not exist – or at least we have no reason to believe in its existence – unless it manifests itself somehow. The evaluativist theory of desire also says that desires must manifest themselves somehow, at least counterfactually, or we have no reason to say that they exist. The way to distinguish dispositional theories from evaluativist theories is not in virtue of the fact that one has manifestations and the other does not, but that certain manifestations are indicative of one theory and different manifestations indicative of the other.

The difference between the two major theories is not found in the fact that one involves dispositions and the other does not. It is found in the fact that so-called dispositional theories are understood in terms of dispositions to act, whereas evaluativist theories are understood in terms of dispositions to display the manifestations of value. Attempts to defend a dispositional theory by defining dispositions so broadly that they include dispositions to value will be self-defeating.

Multiple Desires

The fourth problem with dispositional theories has to do with how it handles multiple desires.

Dispositional theories tell us what it is to have a single desire. It is a disposition to realize a given this state. Theorists recognize that these dispositions depend on other states. The desire to see the movie disposes one to go to the theater depending on what else one might want to do with the money. But when it comes to actually using the system to explain and predict actual behavior of agents with multiple desires, it becomes complicated very quickly.

Imagine that you have a scale with 100 weights of different and sometimes variable values. Try to develop a counterfactual account that will allow you to predict which way the scales will tilt as different weights are placed on the scale. "Weight B2 will tilt the scales in its direction is combined with weight H9 and I6 against C3 and D3, but not against C3 and D4, unless combined with T2 or F14." One can see how unwieldy this would be.

We can see the practicality of assigning a value to the weight and simply saying that the scale will tilt in the direction of the total weight. We would not need to consult an impossibly large array of counterfactual conditions to predict the effects of the weights on the scale, we could compute the results.

This is comparable to taking seriously the force metaphor of desire. A force is expressed as a vector. A vector has two components; direction and magnitude. When speaking about a force, one speaks of it as pushing in a particular direction and having a particular strength. It would s typically represented as an arrow that points in the direction of the force, with its length determining its magnitude. The heuristic for computing the effect of multiple forces is to line them up, nose to tail, then draw a force vector from the tail of the first force to the head of the last. This gives the direction and magnitude of the total force – the behavior of an object having these forces acting on it.

Dispositionism can handle the direction component of motivational force easily enough. A desire that p pushes the agent to act so as to make it the case that p becomes or remains true. However, it has no easy way to accommodate magnitude. Magnitude is expressed as a list of counterfactual conditions specifying how the object will behave in the context of other possible desires. Though this is theoretically possible, it is entirely impractical.

Assigning a value to a desire is an evaluativist trick. It tells us how good the object of the desire is, or how important, or the strength of its ought-to-be-ness, depending on the version of evaluativist theory one adopts. One can then compute the effects of multiple desires on an agent by combining the strengths and directions of these individual desires. A person with a strong desire that p and weak desire that q will prefer p&q if available, p over q if she must choose, and q (as compared to not-q) if p is impossible.

Indeed, this absence of an assignment of a value to the magnitude of the desire is exactly what is missing from Quinn's Radioman. Furthermore, the importance of p is what explains an agent's motivation to realize p where p is false, and to preserve p where p is true. It is because "that p" matters to the agent with a desire that p.

Summary
Dispositional theories fail because they cannot account for the strength of a desire. One may be tempted out of a love is symmetry to say that evaluativist theories cannot handle the direction of desire. That would be poetic in a way, but this is not the problem with evaluativist theories. They can handle both the direction and the magnitude of desire vectors without any difficulty. Their problem is in handling what it is the vector stands for. I will examine this issue in the next section.

4. Evaluativist Theories of Desire

What Evaluativist Theories Get Right
Theories of goodness and duty.
Direction of fit: importance.

5. Assignment Theory of Desire

6. References

Alvarez, Maria (2017), "Desires, Dispositions, and the Explanation of Action". In Deonna J. & Lauria F. (eds). *The Nature of Desire*. Oxford: Oxford University Press.

Anscombe, G.E.M. (1963), *Intention (Second Edition)*, Basil Blackwell, Oxford: Oxford University Press.

Arpaly, N. and Schroeder, T. (2014). *In Praise of Desire*, New York: Oxford University Press.

Davidson, Donald, 1963, "Actions, Reasons, and Causes", *The Journal of Philosophy*, Vol. 60, No. 23, American Philosophical Association, Eastern Division, Sixtieth Annual Meeting (Nov. 7, 1963), pp. 685-700

Doring, Sabine A. and Eker, Bahadir (2017), "Desires without Guises". In Deonna J. & Lauria F. (eds). *The Nature of Desire*. Oxford: Oxford University Press.

David Hume, 1896, *A Treatise of Human Nature* by David Hume, reprinted from the Original Edition in three volumes and edited, with an analytical index, by L.A. Selby-Bigge, M.A. (Oxford: Clarendon Press, 1896). 7/8/2018. http://oll.libertyfund.org/titles/342.

Hursthouse, R. (1999). *On virtue ethics*. Oxford: Oxford University Press.

Lauria, Fredrico (2017), "The 'Guise of the Ought-to-Be': A Deontic View of the Intentionality of Desire". In Deonna J. & Lauria F. (eds). *The Nature of Desire*. Oxford: Oxford University Press.

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

Nozick, Robert (1974). *Anarchy, State, and Utopia*. New York: Basic Books

Oddie, Graham (2017), "Desire and the Good: In Search of the Right Fit". In Deonna J. & Lauria F. (eds). *The Nature of Desire*. Oxford: Oxford University Press.

Quinn, Warren (1993), "Putting Rationality in its Place", in *Morality and Action.* Cambridge, UK: Cambridge University Press.

Schroeder, Timothy (2004). *Three Faces of Desire,* New York: Oxford University Press.

Smith, Michael (1987), "The Humean Theory of Motivation", *Mind*, Vol. 96, No. 381, pp. 36-61.

Smith, Michael (2012). "Four Objections to the Standard Story of Action (and Four Replies)", *Philosophical Issues, Vol. 22, Action Theory*, pp. 387-401.

Stanford Encyclopedia of Philosophy, "Desire", First published Wed Dec 9, 2009; substantive revision Thu Apr 9, 2015, https://plato.stanford.edu/entries/desire/.

Stanford Encyclopedia of Philosophy, "Reasons for Action: Internal vs. External", First published Thu Sep 4, 2008; substantive revision Fri Aug 18, 2017, 7/8/2018, https://plato.stanford.edu/entries/reasons-internal-external/.

Street, Sharon, "A Darwinian Dilemma for Realist Theories of Value," Philosophical Studies 127(1):109-166