

CHAPTER SIX
MAKING SENSE OF THE RELATIONSHIP BETWEEN SELF-CONCEPTIONS
AND PRACTICAL REASONS

6.1 Introduction

Having achieved the first organizing goal of this project in the preceding chapter (which was to offer a precise and instructive statement of the relationship between self-conceptions and reasons for action), we turn now to the second and third of these goals (to describe the nature of this relationship and to begin¹ to explore some of its more interesting implications). Proceeding on the assumption that self-conceptions and practical reasons are related as the Modified Co-Determination Thesis (MCDT) suggests, my present aim will be to explain *why* a relationship between the two exists in the first place. What accounts, in other words, for our intuitions that our practical reasons and self-conceptions have such an important impact on each other? To jump ahead to the outcome of this inquiry, I will be drawing the following conclusions:

1. The relationship between self-conceptions and practical reasons is not an accident of the way certain questions have been posed; the relationship between the two is rather one of necessity.
2. The necessity in question is conceptual. Our very concept of a reason for action is such that all reasons must be understood as addressed to agents falling under certain descriptions (e.g. we must understand a reason to wait in long lines every week to watch and enjoy soccer matches as addressed to “Soccer Enthusiasts”).

¹ My pursuit of this goal will carry over into the next chapter.

3. This fact about the relationship between self-conceptions and practical reasons has a number of interesting implications. Attending to this relationship can, for example, give us new avenues for critiquing proposed actions (by helping us to see what these actions imply about us) and simplifying the process of deliberation (because many self-conceptions “bundle” and summarize large numbers of practical principles). Thinking about the terms of this relationship can also help moral theorists justify the common but rarely acknowledged practice of grounding moral claims in appeals to self-conceptions.

To establish these claims, I proceed as follows. In Section 6.2 I argue for the view that there is a conceptual connection between practical reasons and self-conceptions. I suggest that this connection is explained by fact that reasons are essentially *functions* relating agents under certain descriptions to the world of possible actions. In Section 6.3 I present a pictorial model of this relationship to help make better sense of the connections between its members. I then conclude in Section 6.4 by cataloging some interesting implications of this relationship.

6.2 Making Sense of the Relationship

As we have seen, there is a definite relationship between normative reasons and licensed self-conceptions. Following Aristotle’s adage that proper understanding requires knowing *why* in addition to knowing *that*, our task will be to explain the basis for this relationship. How can we account for it? What is the precise nature of this relationship?

To make headway on these questions, a good place to start would be to ask whether this relationship is accidental or fundamental. Is this relationship simply an accident of the way in which certain questions have been posed? Or is there a sense in

which it *must* be the case that licensed self-conceptions and normative reasons are closely connected?

Let's first explore the possibility that the relationship is accidental by assuming that this is so. Let's assume, in other words, that the connection between licensed self-conceptions and normative reasons is more indicative of their coupling in our inquiries than of a more fundamental link between them. After all, one can identify an accidental relationship between any two items—including normative reasons and kitchen sinks—simply by asking the right questions about them. For example, in the course of thinking about how normative reasons might be related to kitchen sinks, we might note that every action recommended by a normative reason must be performed in certain proximity to a kitchen sink (e.g. “save the person who is drowning 223.4 feet from that kitchen sink”).

In assuming that the relationship between licensed self-conceptions and normative reasons is accidental, we are assuming that there is no real difference between this “relationship” and the absurd one involving kitchen sinks. We are assuming, in other words, that there is no fundamental, “question-independent” relationship between the two. Our task in what follows will be to determine whether this assumption corresponds with our other views about what normative reasons are like. I will argue that it does not, and hence that the discussion that follows constitutes a *reductio* against this assumption.

The first step is to show that, if there is no interesting relationship between normative reasons and licensed self-conceptions, nor can there be an interesting relationship between normative reasons and *licensed descriptions* of an agent more generally (where a licensed description is a description it would be appropriate to apply to an agent even if she doesn't describe herself that way). Note that this is a perfectly legitimate modification to make; simply altering the perspective of the

person attributing a description to an agent has no bearing on the present question given the nature of our assumption. If there is no link between my licensed self-conception “Jeff’s Friend” and my normative reasons, how could there be a link between the licensed description “Jeff’s Friend”—a description that could be attributed to me rightly by someone else—and my normative reasons?²

To recap, we are now assuming that licensed descriptions of me and my normative reasons are not fundamentally related. On the basis of this assumption, we are now in position to draw the following, rather important inference. If there is no fundamental relationship between licensed descriptions of an agent and her normative reasons, there can be no fundamental relationship between *facts* about this agent and her normative reasons either. This is so because for any given fact about an agent, there will be a corresponding true description of this fact, and vice versa. Facts and true descriptions of these facts, in other words, rise and fall together.

(Note that this category of “facts” is meant to be quite broad; it includes, for example, facts about what one cares about, facts about the projects one has committed oneself to, and facts about one’s attitudes. The only sort of facts that we must exclude are facts about the self-conceptions one holds, simply because including self-conceptions at this stage will lead to recursion problems that will become clear shortly. At the end of this thought experiment I discuss the role that self-conceptions *as such* play in this argument.)

For example, if a description of me as Justin’s brother is unrelated to my normative reasons, then the fact that I am Justin’s brother must also be unrelated to my normative reasons. To see why this is so, simply note that denying the consequent

² Things get more complicated when we are not assuming the absence of link between these terms, since the class of licensed descriptions that are tied to normative reasons might be partially determined by the class of licensed self-conceptions that a person holds (i.e. when this assumption is not in place, the mere fact that a description of an agent is licensed might not entail that it is reason-giving; this will be the case if the agent doesn’t hold a corresponding licensed self-conception and the description in question is “optional”).

requires denying the antecedent. If the fact that I am Justin's brother is indeed tied to a normative reason to be supportive of Justin, then a description of me as Justin's brother will also be tied to this reason. We cannot deny (or affirm) one of these connections without denying (or affirming) the other.

Where does this leave us? If there is no interesting relationship between licensed self-conceptions and normative reasons as we have been assuming, we can conclude that there is no interesting relationship between facts about an agent and his normative reasons either. We can summarize this chain of inferences as follows:

Licensed Self-Conceptions Are Unrelated to One's Normative Reasons	↔	Licensed Descriptions Are Unrelated to One's Normative Reasons	↔	Facts About One Are Unrelated to One's Normative Reasons
--	---	--	---	--

It is at this point, however, that our assumption begins running into problems. The claim that facts about an agent are unrelated to her normative reasons is simply implausible. How could the fact that I am standing near a drowning child (as opposed to being on a different continent) be irrelevant to whether I have a normative reason to save him? How could the fact that I am not a citizen of Costa Rica be unrelated to facts about whether I have a normative reason to vote in Costa Rican elections? To accept this implication of our initial assumption, we would have to accept that one's normative reasons exist independently of whether one is a Mexican politician or an English bookseller, an aspiring community activist or an aspiring doctor, a human being or an angel, alert or in a coma, etc. And surely this is false. Facts about what one is and how one is situated in the world must be related to one's normative reasons. If we know that certain normative reasons exist, we can infer that certain facts obtain; conversely, if we know that certain facts obtain, we can infer that certain normative reasons (or ranges of possible reasons) exist.³

³ For example, if we know that Sarah has normative reason to swim out to Jamie to save him from drowning, we can infer that Sarah is near a body of water, that Jamie is in this body of water and is at

Based on the chain of biconditionals summarized above, to reject this implication of our initial assumption, we must reject our initial assumption as well. Running backwards through these inferences, we can see that if facts about an agent are fundamentally related to her normative reasons, licensed descriptions of the agent—including her self-conceptions—must also be fundamentally related to her normative reasons.

The upshot of our thought experiment is that there is indeed a necessary relationship here, just as our intuitions predicted. Where there is a reason, there will be a licensed description of the agent who has that reason, and moreover, one that can at least potentially serve as a licensed self-conception. Note that this is not the case with the normative reasons/kitchen sinks “relationship.” Not only is the location of any given kitchen sink incidental to the content of a given normative reason (a different distribution of sinks wouldn’t affect the nature of the recommended actions),⁴ but we could also easily locate recommended actions relative to objects other than sinks. Conversely, a different “distribution” of licensed self-conceptions would radically change the content of particular normative reasons, and nothing could replace the role that licensed descriptions play in our understanding of normative reasons.

Before moving on to the next set of questions, it might be useful to reflect on what precisely we have shown to this point and how far this argument takes us. In particular, one might worry that the relationship between self-conceptions and normative reasons we have established is merely “epiphenomenal.” In other words, based on the preceding arguments, it seems as though the facts about the case are what are directly related to one’s normative reasons, and self-conceptions are just along for

risk of drowning, that Sarah knows how to swim, etc. Conversely, if we begin with this information about the facts of the case (perhaps coupled with practical principles or claims about what is of value), we can infer that Sarah has normative reason to save Jamie.

⁴ Unless of course those reasons involve doing something to or in a kitchen sink.

the ride as possible reports of these facts. Do self-conceptions *as such* figure into this discussion at all? Based on what's been said so far, the answer seems to be "no." I haven't yet shown that self-conceptions as such play a significant role in, for example, determining the content of one's reasons. However, three points are worth noting. The first is simply that, even if the role that self-conceptions play in the story so far is purely epiphenomenal, this doesn't change the fact that the relationship in question is one of necessity. Licensed self-conceptions and normative reasons are indeed related as our intuitions predicted. But this needn't be the end of the story. Now that we have our foot in the door by establishing a necessary (albeit quite general) relationship between the two, we can fill in this account to show more specifically that self-conceptions are related to reasons in a more robust sense as well.

This claim brings us to my second point. While it is true that in some cases the facts of a situation are what determine one's normative reasons—and so the self-conceptions in question simply "report" these underlying facts—this is not always the case. In many situations, the underlying facts will *underdetermine* the possible all-things-considered normative reasons one has. For example, the fact that David is my co-worker and that we often swap amusing stories may leave it open for me to think of myself (a) merely as "David's Colleague," or (b) as "David's Friend" as well. While the (self-conception independent) underlying facts are certainly relevant to the pro tanto reasons I have, if I come to hold the self-conception "David's Friend"—if I take this to be a fitting interpretation of our existing relationship—this self-conception *itself* will play a direct and decisive role in the determination of my all-things-considered reasons. For example, perhaps I only have reason to help David paint his house if I hold this self-conception; mere colleagues, one could argue, don't have compelling

reasons of this sort.⁵ The upshot is that some self-conceptions as such can determine one's normative reasons.

Given the importance of this point, it is worth discussing in a bit more detail. To sum up, the story I've been telling isn't simply that facts give rise to licensed descriptions which (at least potentially) give rise to self-conceptions on a one-to-one basis. The licensed descriptions of me will not map one-to-one onto the self-conceptions I do or could hold because these descriptions (at least in many cases) underdetermine the specific self-conceptions I do or could hold. To fit this claim into the argument for this relationship presented above, the idea is that while the underlying facts of a case are closely linked to associated *pro tanto* reasons—and hence to licensed descriptions of these facts and first-person reports of them—one needn't actually hold all of the self-conceptions one would be licensed in holding. There will be cases in which the underlying facts make it appropriate to hold one of a number of possible self-conceptions—each of which might be linked to different normative reasons—and the agent will have some discretion in choosing which of these potential licensed self-conceptions to hold (and hence some control over his normative reasons).

This holds equally true in cases for which the agent can describe underlying facts in different ways; different descriptions can lead to different understandings of underlying facts, and perhaps different normative reasons as well (e.g. the self-conceptions “Invalid” and “Person Who Cannot Use His Legs” are not equivalent—they resonate emotionally in different ways, may have a different valence for different

⁵ There are actually a number of ways of describing what is going on in this case. One might hold that prior to thinking of myself as “David's Friend,” I had a weak *pro tanto* reason to help David paint his house on the basis of other considerations (e.g. this act might increase net pleasure), and that coming to hold this self-conception added enough normative weight to the picture to give me an all-things-considered normative reason to help in this way. Alternatively, one might hold that there was no normative reason whatsoever to help David prior to thinking of myself as “David's Friend.” How we describe this case will be determined at least in part by how detailed we want to make the relevant reason-claims.

agents, etc.). In short, while it is correct to claim that facts (and hence licensed descriptions and licensed self-conceptions) are necessarily related to normative reasons generally, in specific cases some of these potential normative reasons will apply to an agent only because she actually adopts certain licensed self-conceptions.

The third point is that, while self-conceptions as such only determine one's reasons in some cases, it is important to recall that the relationship between self-conceptions and reasons we are examining *isn't exhausted* by a self-conception's capacity to determine one's reasons. The relationship in question is bi-directional. For cases in which a self-conception merely "reports" underlying facts, the relationship is this: the existence of the normative reasons entail the appropriateness of holding certain licensed self-conceptions. Once again, in either case the relationship is between self-conceptions *as such* and normative reasons; the only difference is that in some cases the self-conception as such determines the reasons, and in the other cases the reasons entail the appropriateness of holding certain self-conceptions.

Having established these facts about the necessity of this relationship, we must now ask what sort of necessity we are talking about and what accounts for it (i.e. we have not yet fully addressed Aristotle's adage). One approach to answering these questions involves adopting a specific view about the role that reasons play in our lives. While I can only outline this view here, a sketch of it should demonstrate how it can accommodate one possible answer to these questions.

According to this view, our concept of a normative reason exists because we as deliberative beings unavoidably face (and acknowledge) questions about what to do, and we cannot proceed in our deliberations without believing that some answers to these questions are better than others.⁶ Should I prepare for lecture tomorrow, or

⁶ Korsgaard is a chief proponent of this view of reasons. She claims that we have normative concepts in the first place (to include our concept of a reason) "...because we have to figure out what to believe and what to do. Normative concepts exist because human beings have normative problems. And we have normative problems because we are self-conscious rational animals, capable of reflection about what

should I watch the Red Sox game? Normative reasons—i.e. considerations that actually count in favor of certain actions—are essentially answers to these sorts of questions. As we have just seen, facts about us and our circumstances (and hence the descriptions of us through which we make sense of these facts) must be relevant to these answers. Surely the fact that I teach a certain Philosophy class (and justifiably think of myself as a “Teacher of Philosophy”) must figure into the answer to my question about whether I should prepare for lecture or watch a baseball game.

To put the point another way, if normative reasons offer good answers to practical questions, and if they can do this only if they are properly sensitive to how certain agents are situated, then a consideration will only count as a normative reason if it is sensitive to an agent’s situation. Just as a counselor cannot help me without being properly responsive to the particulars of my case, a reason cannot counsel me in helpful ways unless it too is properly responsive to the particulars of my situation (to include the descriptions through which I make sense of, and give shape to, my situation). This suggests that our very notion of a normative reason is at least partially that of a pro-consideration that directs agents situated in certain ways to perform certain actions. According to this view, we can think of reasons as *functions* relating agents under particular descriptions to the world of possible actions. And it is because reasons are essentially functions connecting descriptions of agents (“Teacher of Philosophy”) to the world of recommended actions (“prepare for lecture”) that the relationship between the two is one of necessity.

we ought to believe and do” (Korsgaard (1): 46). Alan Gibbard discusses a similar view in *Thinking How to Live*. He asks: “Why do we, in a natural world, have non-naturalistic concepts [like that of a normative reason]? Because we think what to do...” (Gibbard: 184). He goes on to remark that “Normative concepts, then, aren’t hocus-pocus, even of a kind we find we just can’t shed. We can’t do without them....[they are vindicated] by assimilating their use to a kind of thinking that needs no vindication: thinking what to do” (Gibbard: 196). It is interesting to note that while this view has been most prominently supported by substantive anti-realists (Korsgaard is a constructivist and Gibbard is at least ostensibly a non-cognitivist), it is a view that substantive moral realists are equally capable of embracing (indeed, one could think of my argument for VBR in Chapter 4 as a discussion of what this general view of reasons implies if one pursues its implications to the end).

Facts about this necessity, operating in concert with this view of the role that reasons play in our lives, suggests the following conclusion. There is a *conceptual connection* between reasons and descriptions of possible respondents to them, such that any full accounting of the former must make reference to the latter. While I have stated this conceptual connection in terms of descriptions of agents more broadly, note that it holds equally between reasons and self-conceptions when agents apply these descriptions to themselves. We can accordingly say that the necessity in question is conceptual; we possess the concept of a normative reason in the first place because we acknowledge our need to figure out what to do, and conclusions about what to do must link us to actions that it makes sense for us—acting under particular descriptions—to perform.⁷

Let us refer to this claim about the nature of normative reasons as the Agent-Addressed Reasons Thesis (AAR):

AAR: Every reason for action must be understood as being addressed to a particular class of agents falling under a particular description.

Implicit in AAR is the notion that it makes no sense to speak of reasons for things to happen *simpliciter*, such that, for example, there is simply a reason for someone's pain to come to an end where this reason isn't addressed to a particular class of possible respondents. AAR holds instead that we should think of reasons as three-place predicates recommending (1) that a certain action be performed, (2) for the sake of a particular end, (3) by a particular sort of respondent.

For example, we might claim that a reason to wait in long lines every week to watch and enjoy soccer matches is addressed to soccer enthusiasts, a reason to write a paper on a certain topic is addressed to students in a certain class, and a reason to act

⁷ To put the point in terms of our earlier contrasting cases, the central difference between the reasons/licensed description relationship on the one hand, and the reasons/kitchen sinks "relationship" on the other, is that our concept of a normative reason requires the concept of a licensed description to be meaningful, but does just fine without the concept of a kitchen sink.

so that you use humanity always at the same time as an end and never merely as a means is addressed to all rational agents as such. These examples show how reasons are linked to particular classes of possible respondents; self-conceptions (at least in one application) enter the picture when an agent considers herself under a given self-description as a potential respondent to reason-claims.⁸ The upshot is that if I decide that I have reason to wait in long lines every week to watch and enjoy soccer matches, I will be conceptually committed to a self-conception that renders the existence of this reason intelligible, such as “Soccer Enthusiast,” or something of this sort.⁹

6.3 Modeling AAR

We now have a thesis that explains the necessary relationship between normative reasons and licensed self-conceptions, and an important addition to our understanding of the concept of a normative reason for action. While these products of our discussion are themselves of value, there are also a number of interesting implications that follow from them. Before turning to these implications, however, I would like to introduce a pictorial model of the relationship posited by AAR since doing so will help frame the discussion in the next section.

The model I have in mind is as follows. Picture two adjacent columns. The left hand column contains a list of the self-conceptions a given agent is licensed in

⁸ As discussed above, it is worth recalling that self-conceptions are not limited to merely “reporting” underlying facts. In many cases, underlying facts underdetermine both the self-conceptions available to the agent and the normative reasons associated with them. Accordingly, self-conceptions can play a more “active” role in the relationship in question by influencing the final determination of one’s reasons.

⁹ Note that while most of these examples involve true descriptions of what agents are actually like, AAR indicates that similar conclusions hold in the case of descriptions that agents are trying to make true of themselves. To return to an earlier example, if I decide that it would good to become a “Community Activist” but I am not one yet (and on the assumption that this is a value-commensurate self-conception), by committing myself to this self-conception I may come to have normative reasons to take steps that would make this a true description of me (like signing up for community betterment projects, attending community meetings, etc.). According to AAR, there still must be some description of me that renders these reason-claims intelligible. In this case, that description might be something like “Fledgling Community Activist.”

holding. The right hand column contains a list of general action-types that this agent has normative reason to perform (we can think of each of these action-types as preceded by the statement: “Agent A has a pro tanto normative reason to...”).^{10, 11} Let us refer to these two columns as a given agent’s “conception-action field.”¹²

To illustrate, consider the conception-action field of an agent named Roy. Assume that the left hand column of Roy’s field includes licensed self-conceptions like “John’s Brother,” “Columbia Undergraduate,” and “Pacifist,” and the right hand column includes action-types like “promote John’s well-being,” “enroll in classes that meet Columbia’s distribution requirements,” and “never harm others.” Note that since most self-conceptions are tied to a number of different action-types (promoting a brother’s well-being isn’t the only sort of thing that a brother typically does), there will be many more entries in the right hand column of Roy’s conception-action field than there are in the left.

¹⁰ One might wonder how this talk of “action-types” hooks up with my earlier distinction between two ways that self-conceptions can figure into one’s practical reasoning (i.e. by recommending practical projects or serving as practically relevant considerations). The action-types I have in mind are certainly meant to capture recommended ends. For example, the self-conception “Diabetic” might be associated with the action-type “monitor your blood sugar.” Depending upon how complex we want to make our conception-action fields, we could also add a potentially large number of “conditional” action-types that link self-conceptions to situations in which they function as relevant practical considerations. Sticking with the previous example, if we wanted to include action guidance tied to fairly realistic scenarios, we could add principles like: “avoid situations in which you will be tempted by your favorite dessert.” Perhaps less realistically, if we wanted to capture every possible scenario, we could add thousands of conditional action-types of the following sort: “conceal your health status from neighbors to spare their feelings if they send you what appear to be rather expensive ‘welcome to the neighborhood’ pastries.” In any event, the point is that these actions-types could be understood as capturing both of the ways in which a self-conception can figure into one’s reasons, depending upon how detailed we want to make our conception-action fields.

¹¹ The order of these columns doesn’t really matter; I have assigned sides to simplify the discussion that follows.

¹² While I have described conception-action fields in terms of licensed self-conceptions and normative reasons—such that a given agent’s conception-action field may contain reasons and potential self-conceptions the agent isn’t aware of, we can also speak of “subjective” conception-actions fields that contain only self-conceptions the agent actually holds and reasons the agent thinks he has. The sorts of reflective equilibrium pressures that I discuss below and their “impact” on one’s conception-action field are best understood in reference to a subjective conception-action field.

AAR tells us that each entry in the right hand, action-type column of Roy's conception-action field needs to be linked in appropriate ways to at least one entry in the left hand, self-conceptions column. For example, if Roy has normative reason to enroll in classes that meet Columbia's distribution requirements, there needs to be a licensed self-conception in the left hand, self-conception column that renders this reason-claim intelligible. The self-conception "Pacifist" isn't up to this task. Perhaps the only licensed self-conception that would be is "Columbia Undergraduate." If it isn't appropriate for Roy to hold this self-conception (if, for example, he is an undergraduate at Boston College), we would have to declare this purported reason-claim invalid.¹³

We could represent this demand of AAR pictorially by saying that every right-hand column entry needs to be connected by a line to at least one item in the left-hand column.¹⁴ There can be no general action-types that we have pro tanto normative reason to perform that cannot be appropriately linked to some licensed description of us. It is worth noting, however, that there is an asymmetry here when we turn to the left-hand, self-conceptions side of the field. While most left-hand self-conceptions

¹³ By this point I hope I have made good on my earlier promise to explain why it is that we can pair up self-conceptions with reason claims in all cases, including rather uninteresting ones. While it might seem trivial to claim that a reason to eat chocolate might have to link up with the self-conception "Chocolate Lover," this is what this conceptual connection demands. This actually shouldn't come as much of a surprise; conceptual connections often have uninteresting implications in specific cases. (Perhaps an analogy would help. Granting that there is a conceptual connection between normative reasons and justification, one might get very tired of hearing someone point out "so that act is justified" every time a normative reason claim is made. Similarly, I won't deny that one might become equally tired of hearing someone point out "so there's a self-conception it would be appropriate for you to hold" every time a normative reason claim is made. That having such things pointed out all the time might lead one to think of these claims as trivial doesn't speak against the conceptual claim; it might instead be taken to speak in favor of it). In claiming that it might be useful in some cases to think through the implications of this relationship (as I do below), I certainly don't mean to claim that doing so is useful in all cases. Indeed, there will be many cases in which trying to trace out the implications of a rather banal reason claim (e.g. to apply a Band-Aid, drink some water, scratch an itch) would be a tremendous waste of time.

¹⁴ It is possible that a given action-type could be recommended by multiple self-conceptions, either individually or in concert (i.e. there may be some cases in which an action will be recommended by two or more self-conceptions working together dynamically, even though it isn't recommended by any one of them individually).

will be linked to right-hand action-types, depending upon how detailed we want to make our conception-action fields, not all of them need be. The reason is simply that some self-conceptions don't appear to be associated with general action-types (to include perhaps "Person Who Can Curl Her Tongue"). In the interest of open accounting, it may make sense to include such self-conceptions on our conception-action fields despite the fact that they are rarely of practical relevance.¹⁵

How does this pictorial model link up with our subjective experience of this relationship? Given the nature of the conceptual demand made by AAR, it will manifest itself in our thinking in the form of a rational pressure to accept reason-claims only if they can be linked in appropriate ways to accepted self-conceptions. When in a reflective mood, we tend to be skeptical of reason-claims if we can't see them as appropriately connecting us (under relevant descriptions) to the actions in question.¹⁶ Stating precisely what determines the "appropriateness" of a given connection is not easy, but the need for this qualification is intuitive. Surely we don't want to say that the general action-type of playing a game of checkers can be appropriately linked to the self-conception "Lumberjack." (Bizarre circumstances might provide a link between the two in specific cases, but there doesn't seem to be a general link between the two.)¹⁷

¹⁵ Depending on how detailed we want to make our conception-action fields, we could do away with this asymmetry by including all possible action-types in the right-hand column. We might, for example, want to tether the self-conception "Person Who Can Curl Her Tongue" to general action-types that involve entering tongue-curling contests or showing off this trait to bored children. At bottom, our decisions concerning which self-conceptions we should treat as tied to general action-types will be inevitably ad hoc. At the extreme, all self-conceptions could be tied to action-types. In claiming that self-conceptions like "Person Who Can Curl Her Tongue" might belong on a conception-action field untied to any right-hand entries, I am simply following what I take to be conventional intuitions governing which self-conceptions are "inherently action-guiding."

¹⁶ This is not to say that this pressure comes up often in our lives, but rather that it can come up and sometimes does.

¹⁷ To reiterate an earlier point, the explanation for this lies in the fact that many self-conceptions are closely associated with characteristic activities and concerns (e.g. to understand what an environmental activist is requires an understanding of the sort of practical principles that characterize people falling under this description). On the basis of this observation, we can say (as a first approximation) that the appropriateness of a link between a self-conception and action-type is determined by characteristics

To put this point another way, we could say that AAR can manifest itself in our thinking as a demand that we achieve reflective equilibrium between our self-conceptions and what we take to be our reasons for acting. Holding a self-conception generally disposes one to treat certain considerations as reasons, and to accept that you have reason to do something generally involves seeing yourself as someone for whom it would make sense to perform the action in question. Like good double-entry bookkeeping, we are under rational pressure to ensure that the columns of our conception-action fields are in balance. We need our action-types to be linked to self-conceptions, and with some minor exceptions, we need our self-conceptions to be associated with acting in characteristic ways (because this is often part of what it means to hold a self-conception).

One implication of this reflective equilibrium pressure is that when one's columns don't tally, some change to one's conception-action field must be made. If I accept a reason-claim directing me to perform an action-type that is either "orphaned" (i.e. not appropriately connected to a self-conception that I hold) or fundamentally contrary to a self-conception that I hold, I either need to disavow this reason-claim or change my self-conception profile. For example, if I'm really convinced that I have good reason to join the military to defend my country from terrorists, I either need to stop being and thinking of myself as a "Pacifist," or I need to question the legitimacy of this purported reason (which might involve looking for other ways to help protect my country).

A second, closely related implication is that changes to one column of my conception-action field may necessitate changes to the other. If I come to realize that a certain self-conception does indeed apply to me ("I'm really his father?"), it is likely

internal to the self-conception. If the general action-type figures into our understanding of the self-conception in the right way, the connection between the two in our conception-action field will count as "appropriate."

that some new items will need to appear on the right hand, action-type column to maintain reflective equilibrium. Alternatively, if I come to realize that I have reason to do something that I hadn't recognized before, I may need to add a new self-conception or modify an existing one to reflect this fact. For example, perhaps deciding that I have reason to make a big sacrifice to help Eric out of a jam implies that I should start thinking of myself as "Eric's Friend," and not just as "Eric's Neighbor." The full implication of these reflective equilibrium pressures will become clear in the next section.

Before moving on, however, I would like to discuss one more feature of this pictorial model of AAR. There are generally two sorts of questions we can ask that draw attention to different sides of our conception-action fields. First, we can ask questions about what we ought to do, and second, we can ask questions about how our self-conceptions are related to one another. When we ask the first sort of question about what to do, the natural place to look is at the right hand, action-type side of the field. Focusing on this side of the field might prove useful if we want to identify conflicting practical demands or "multiply recommended" actions (i.e. actions advocated by multiple considerations). Conversely, when we ask questions of the second sort about how our self-conceptions are related to one another, the natural place to look would be the left hand, self-conceptions side of the field. Scrutinizing the items in this column, for example, might help us ensure that our self-conceptions cohere with one another.

AAR tells us, however, that we needn't answer these questions by looking exclusively at the columns of a conception-action field that naturally correspond to these questions. We can, in other words, "translate" discussions involving one side of the field into discussions involving the other. When trying to decide whether two self-conceptions cohere in the right way, I can focus my attention on the practical

implications of these self-conceptions by looking at their associated action-types. Can I be both an “American Patriot” and a “Pacifist” at the same time (at least on a certain understanding of what it means to be an “American Patriot”)? One way of answering this question would be to look at the action-types associated with these two self-conceptions. If it turns out that they are in fundamental conflict, I might decide that I must either abandon or significantly rethink one of these self-conceptions.¹⁸

Conversely, when trying to decide how I ought to act in a given case, I can think of my situation in terms of what certain self-conceptions require of me, rather than in terms of the sometimes dizzying number of action-types I appear to have reason to perform. Instead of trying to weigh a large number of reasons for performing diverse action-types, I might be able to more easily determine what to do by simply focusing on related self-conceptions (for example, by thinking more abstractly about what a “Close Friend” should do in this situation).

6.4 Implications of AAR

In addition to being a thesis of interest in its own right, AAR has a number of interesting implications, four of which I’ll discuss in the remaining pages.

(1) AAR Gives Us A New Avenue For Critiquing Proposed Actions. The first noteworthy implication of AAR is that this thesis gives us a new avenue for assessing and critiquing candidate reason-claims. When we deliberate about what to do, we

¹⁸ For example, if I come to think that an “American Patriot” must defend his country in times of crisis by joining the military, I might decide that I cannot be both an “American Patriot” and a “Pacifist.” This restriction to cases of fundamental conflict is important; given local circumstances, it is always possible that any two self-conceptions could recommend conflicting actions. The self-conception “Justin’s Brother” might recommend spending Sunday afternoon helping Justin study for a big test, while the self-conception “Good Samaritan” might recommend helping a severely injured motorist that I come across on my way home to help Justin study. Such accidental conflicts, however, wouldn’t necessitate abandoning one of the two self-conceptions. A case of fundamental conflict exists when the higher-level, general practical principles associated with two self-conceptions are in direct tension, absent any consideration of specific circumstances (e.g. even in the abstract, I can see that the practical principles “use force to defend your country” and “never perform an act of violence” are in fundamental conflict).

typically focus on the practical demands that we take ourselves to be under. We ask ourselves questions like: how can I best keep my promise to Susan without letting another submission deadline slip? AAR, however, suggests that there is more than one general approach to evaluating reason-claims. Instead of focusing simply on the ramifications of particular practical demands, AAR tells us that we can evaluate reason-claims by looking at what they imply about the agents who accept them. More specifically, AAR tells us that any time an agent claims to have a reason to act, we can demand that this agent produce a self-conception that renders this reason-claim intelligible (and moreover, a self-conception that the agent is willing to accept). If the agent cannot produce one—if he cannot cite a plausible description of himself that makes it clear how this reason-claim applies to him—according to AAR, he would be rationally compelled to withdraw his reason-claim.¹⁹

For example, suppose that I am thinking about breaking an important promise to my friend Jeff. Perhaps I am afraid that fulfilling this promise would cause me to suffer something mildly unpleasant. Furthermore, suppose that I am beginning to see a promise-breaking pattern in my behavior; over the past few months, I have reneged on promises to avoid mildly unpleasant outcomes on a number of occasions, many of which involved promises to Jeff. Regardless of how I actually deliberate in this case, AAR suggests that a decision to break my promise will have two notable implications. First, insofar as frequently breaking promises to someone is contrary to the norms of friendship, this pattern of behavior would undermine my warrant for thinking of myself as “Jeff’s Friend.” Second, perhaps the only self-conception that would appropriately give purchase to a reason to renege on promises of this sort is that of

¹⁹ This of course does not mean that an agent’s inability to think of or articulate a corresponding self-conception implies that the reason-claim in question is invalid. The idea is rather that, from the first-person perspective, it is problematic to believe both (a) that I have reason to do something, and (b) that there are no appropriate descriptions of me that render this reason-claim intelligible.

“Untrustworthy Egoist,” or something of this sort.²⁰ To put the point in terms of conception-action fields, it may be the case that no other (left hand) descriptions of me could appropriately link up with a (right hand) reason to break promises whenever I find it mildly convenient to do so.²¹

With these two implications of promise breaking this promise in mind, we can now look at how my deliberations might go depending upon whether or not I attend to my self-conceptions in the course of making a decision. First, if I focus solely on action-types (which is what we as deliberating agents typically do in such cases), I might decide that the best thing to do is to break this promise to Jeff. Avoiding an unpleasant occurrence is, after all, a good thing. I might also think that given our long history, Jeff will probably forgive me for not keeping my word, even if I have to do a lot of apologizing to make this happen.

Alternatively, if I go with an AAR-inspired approach to this problem and consider how this proposed action (perhaps in concert with my past behavior) bears on the appropriateness of holding certain self-conceptions, I might come to a different conclusion. The dual realization that breaking this promise would make it inappropriate to hold the self-conception “Jeff’s Friend” and rationally compel me to start thinking of myself as an “Untrustworthy Egoist” (since this would be an increasingly apt description of me) might be enough to convince me that I ought to keep this promise. Unless I am willing to part with the self-conception “Jeff’s Friend” and start thinking of myself as an “Untrustworthy Egoist,” this purported reason to

²⁰ I am, of course, not speaking of normative reasons in this particular application. I am speaking more generally about the sorts of descriptions that naturally hook up with certain characteristic forms of behavior, regardless of whether they are justified.

²¹ This claim is admittedly implausible, but the oversimplification is necessary. Constructing an example that clearly demonstrated that all possible justifying self-conceptions were so unpalatable that the agent would feel compelled to withdraw his reason-claim would require adding an enormous amount of detail in terms of the agent’s situation, motivating interests, etc.

break my promise can have no purchase on me. The reflective equilibrium pressures implicit in AAR thus compel me to abandon this reason-claim.

As this example illustrates, the point here isn't simply that AAR gives us new avenues for critiquing proposed actions. AAR also opens up the prospect of deliberating in a way that has notable heuristic benefits. The realization that a certain description of me is true can often have a more profound motivational effect on me than the realization that I have been thinking about doing something untoward. In this particular case, the charge that the stigmatizing label "Untrustworthy Egoist" rightly applies to me hits much closer to home (somewhat literally) than the charge that it is wrong to break promises. We may accordingly think of AAR as a potentially useful tool for combating akrasia—or weakness of will—insofar as it can help make the troubling implications of some proposed actions more apparent. If Plato and Aristotle were correct in claiming that stock cases of akrasia involve epistemic deficiencies, deploying AAR as an "epistemic highlighter" of sorts could indeed prove beneficial.²²

(2) AAR Can Help Simplify The Task of Deliberation. This next implication follows from three claims discussed earlier. The first is that we can "translate" questions about what to do into questions about what we are like. The second is that any given conception-action field will typically have many more right hand action-types than associated left hand self-conceptions. The third—which is represented by the modeling device of drawing lines connecting action-types to self-conceptions on

²² It is worth noting that this heuristic benefit applies primarily in cases in which an agent has already failed to acknowledge considerations that he should have acknowledged. Accordingly, when this heuristic benefit does prove effective, it will often be unfortunate that the agent had to rely upon it to see something that he should have seen already. In the example I discuss, it is unfortunate that I am more motivated by considerations about how I think of myself than by any consideration of what I am doing to my friend. (In this particular case, it is ironic that what helps me see that I should think of myself as an "Untrustworthy Egoist" is an appeal to a potentially egoistic interest I may have with my self-image.) With this said, applications of this heuristic device needn't always be this "unfortunate." The simple fact is that we sometimes have a hard time properly weighing the value of our actions, and in many cases thinking about what our proposed actions say about us can help us better appreciate the propriety (or impropriety) of these actions.

our conception-action fields—is that many self-conceptions can be thought of as “bundles” of practical demands that have been summarized under a single description. (For example, the self-conception “Assistant Manager” can be thought of as “bundling” action-types like “process time cards every Friday” and “boost staff morale.”) The resulting implication is that some self-conceptions can make the task of deliberation easier by minimizing the number of items we need to consider when thinking about what to do.

How might this work? Since some self-conceptions summarize large numbers of associated action-types, they can help agents figure out which action a large number of practical demands recommend without forcing them to consider all of these demands individually. An agent who has become familiar with the practical demands implicit in a self-conception can, at least in some circumstances, simply reflect on what this self-conception requires to quickly conclude that a given action would be appropriate.

Take the self-conception “Brother.” Unlike the self-conception “Environmental Activist,” the list of action-types associated with this self-conception may represent a very large number of practical demands that interact in complex ways. For example, one might think that being a brother warrants both (1) looking out for your brother’s well-being regardless of his personal preferences (thus licensing semi-disrespectful actions, such as trying to subvert imprudent life choices), and (2) supporting your brother’s personal goals. Figuring out precisely which action these practical demands recommend can be quite difficult even when we focus only on these two; things get only worse when we consider that the self-conception “Brother” is surely linked to many other practical demands as well.

However, there may be situations in which simply reflecting on the fact that I am Justin’s brother makes clear what I am supposed to do. Suppose I am debating

about what to do in a difficult situation, but fail to reach a settled conclusion because the numerous considerations I am pondering prove too hard to balance. Suppose my thinking in a particular case runs as follows: “On the one hand, what sort of example would I be setting for my brother if I bent the rules for him? On the other, how can I just let him suffer this disappointment given that this isn’t his fault,” etc. After struggling for some time, perhaps a friend’s declaration that “You are his brother, for heaven’s sake!” makes the solution immediately clear. Perhaps after my attention is focused on the demands of this self-conception, it becomes “simply evident” (to me at least) that a brother should bend the rules in this case to save Justin embarrassment.²³ The basic idea here is that balancing a large number of competing action-types can be difficult work. Successful deliberation for us—especially given our epistemic limits—is at least partly a numbers game; AAR tells us that focusing on the self-conceptions side of the divide can help us minimize the number of deliberative items we need to consider.²⁴

(3) AAR Justifies The Practice of Grounding Moral Claims in Appeals To Self-Conceptions. The third implication concerns an often unappreciated fact about

²³ Explaining precisely how these sorts of situations work is something that I cannot hope to do here. To provide a quick analogy, consider the case of a seasoned football coach who can determine what his team should do in any particular situation without having to consult a long list of practical rules for making such choices (e.g. “when it’s fourth and long, punt the ball,” “when it’s late in the fourth quarter and the other team is winning by more than a field goal, go for a first down,” etc.). We might say that this coach has this ability because he has internalized these rules in a certain sense. Similarly, we might say that the person with the self-conception “Brother” might be able to determine what a brother should do by consulting this self-conception instead of the long list of action-types associated with it, perhaps because the relevant practical demands have been internalized “under the guise” of this self-conception. Once these practical demands are properly internalized in this sense, the mere reflection that I am Justin’s brother might be enough to indicate the right course of action.

²⁴ A related point about how personal descriptions can help us deliberate is implied in Elizabeth Anderson’s discussion of what she refers to as “ideals.” “[T]he grounds of a person’s reflectively held values...lie in her conceptions of what kind of person she ought to be, what kinds of character, attitudes, concerns, and commitments she should have. I call such self-conceptions *ideals*....Ideals set the standards of conduct and emotion people expect themselves to satisfy with regard to other people, relationships, and things....Ideals give us perspectives from which to articulate and scrutinize the ways we value things” (Anderson: 6). Anderson’s point mirrors my own that self-conceptions can help us frame choice situations, perhaps in part by giving us useful “perspectives” from which we can better see how certain values should be pursued.

the way many moral theorists have sought to ground their moral projects. Reflecting on the moral theories that have been offered in the past, it appears that many of them—either implicitly or explicitly—rely on appeals to particular conceptions of what moral agents are like. To help justify their moral theories, many philosophers have sought to convince the reader that their moral theories have purchase because they “speak properly” to a particular, true conception of what people are like.

While the adoption of this approach is more pronounced in ancient ethical theories, it is certainly present in many modern theories as well.²⁵ We tend to see, for example, claims of the following sort. Given that you are a rational animal of a certain sort, you ought to pursue activities that lead to the development of moral and intellectual virtues. Given that you are a certain sort of rational agent, you ought to abide by the Categorical Imperative. Given that you are a temporally extended agent that can legitimately occupy both subjective and objective points of view, you should act only on the basis of “objective” reasons (since only objective reasons have purchase for both points of view).²⁶

While some of the moral views I have in mind don’t rely on overt claims of this sort, they all seem to be at least tacitly committed to such claims. Furthermore, many of these theorists seem to assume that once the reader accepts the particular conception on offer as a true description of what she is, there should no longer be any question of being motivated to act in accordance with the particular moral theory. We might sum up the central moves underlying many moral theories as follows: you (the reader) have compelling reason to abide by proposed moral theory M because you are a certain sort of thing (and once you see that you are such a thing—and hence come to

²⁵ Notable exceptions include moral theories that derive from claims about universal human desires (of the sort offered by Hobbes and Hume), and some forms of utilitarianism. While these moral theories have built into them some conception of what moral agents are like, these conceptions don’t play the particular sort of justificatory and motivating role that I have in mind.

²⁶ I offer these three claims as very rough glosses on the moral theories of Aristotle, Kant, and Thomas Nagel, respectively.

hold the relevant self-conception—it will become clear why you have reason to abide by M).

Given the prevalence of this pattern of argument, it would be a good idea to ask whether it is sound. Is this a good way to build a moral theory? Once again, AAR provides useful guidance. Given the conceptual connection between reasons and self-conceptions posited by AAR, it would seem that AAR does indeed license this sort of approach; if you can show that a given description ought to be accepted by an agent (i.e. that it belongs on the left hand side of her conception-action field), you can also show that she has normative reasons to act in ways associated with this self-conception.²⁷ One might even go so far as to say that AAR *recommends* developing moral theories in this way. Indeed, if the sorts of heuristic benefits discussed above are on target, it might be a very good idea to build moral theories upon arguments directing agents to think of themselves under certain descriptions. Furthermore, it might be a good idea both to do this explicitly (rather than leaving the reader to infer this pattern of argument), and to spend quite a bit of time arguing for the relevant description of what agents are like (since this is a rather important step of the process, despite the fact that it rarely gets the amount of discussion it requires).

(4) AAR Give Us Another Way of Understanding the “Schizophrenia” of Modern Moral Theories. In an influential paper, Michael Stocker has argued that modern moral theories fail to properly integrate one’s motives and one’s reasons, leading to an affliction that Stocker refers to as “moral schizophrenia” (Stocker: 66). Stocker’s basic worry is that modern moral theories tell us that we have reason to act

²⁷ Matters become a bit complicated at this point. In most cases we determine whether a given self-conception legitimately recommends normative projects by looking at whether the self-conception is value-commensurate. However, in some applications of the argument I am considering here, the idea is to show that the truth of a given description *makes it the case* that certain things are valuable (i.e. certain things like food, a good education, and the preservation of my rational faculties are valuable because a certain description is true of me). Instead of calling these self-conceptions “value-commensurate,” it might be more accurate to call them “value-determining.”

in ways that run contrary to our “natural understanding” of what we have reason to do. For example, while people typically think that they have reason to love their family and friends for their own sakes, some utilitarian theories tell us that we should care for family and friends as a means of maximizing pleasure instead. Stocker claims that “Not to be moved by what one values...bespeaks a malady of the spirit. Not to value what moves one also bespeaks a malady of the spirit” (Stocker: 66). The problem is that this is precisely what modern moral theories ask us to do: namely, be moved by something other than what we typically value.

AAR allows us to recast this worry as a more literal form of schizophrenia. Applying AAR to Stocker’s central objection, we can rephrase the issue by saying that modern moral theories tell me I need to act out the life of someone who I am not.²⁸ The resulting disharmony becomes one between two conflicting descriptions of what I am; I might hold a certain self-conception about what I am essentially, but modern moral theories tell me that another description of me is true as well, despite the fact that it—and its associated reason-claims—strike me as alien.

To bring out a further implication of this point, allow me to first suggest that perhaps not all moral theories are as open to Stocker’s critique as he thinks. Stocker’s (AAR-translated) claim is that modern moral theories posit descriptions of agents that are incompatible with “normal” human self-views, and hence must be rejected. But certain moderate versions of Kantian moral theory—to give one example—arguably rest on a view of what agents are like that is quite compatible with our everyday understanding of ourselves.²⁹ Perhaps the real problem is thus that modern moral

²⁸ To be clear, the self-conceptions at issue in these cases involve essential descriptions of what moral agents are like (such as “Rational, Political Animal”), not idealizations or descriptions of what perfect moral exemplars are like (such as “Moral Saint” or “Aristotelian Phronomos”). This is important to note because if one has idealizations in mind, one could interpret every moral theory as recommending a certain sort of schizophrenia (since they all recommend becoming a different person).

²⁹ This is no doubt a contentious claim, but Kantian moralists have done a reasonably good job of defending it recently. Barbara Herman in particular has offered some rather convincing arguments to this effect in *The Practice of Moral Judgement* (see especially chapters 2 and 9).

theorists haven't done enough to convince their audience that their theories rest on immanently acceptable descriptions of what moral agents are like. Until they have done so, AAR predicts that their associated reason-claims will fall on deaf ears.

This claim ties nicely into my earlier point about how one ought to construct a convincing moral theory. The lesson to be learned here may *not* be (a) that modern moral theories recommend adopting alien views of what our reasons are, and hence must be rejected (which is what Stocker argues), but rather (b) that modern moral theories haven't done enough to convince readers of the plausibility of the basic, moral agent defining self-conceptions they recommend their readers adopt. What Stocker may have really identified, in other words, is an explanatory shortcoming rather than an objectionable attempt to browbeat the reader into accepting a mysterious set of reason-claims. To reiterate a point I made earlier, the issue is really that moral theorists ought to spend more time explaining why readers ought to accept the self-conceptions that underwrite their moral theories.