CHAPTER 12

Voluntarist reasons and the sources of normativity*

Ruth Chang

In virtue of what does a consideration provide a practical reason? Suppose the fact that an experience is painful provides you with a reason to avoid it. In virtue of what does the fact that it's painful have the normativity of a reason – where, in other words, does its normativity come from? As some philosophers put the question, what is the *source* of a reason's normativity?

This question should be distinguished from two others. One is: Which sorts of consideration ultimately provide practical reasons? That is, are practical reasons given by one's desires, by evaluative facts about what one desires, or by some hybrid of the two?[1] This question concerns which considerations are the ultimate bearers of practical normativity. The question of source, by contrast, concerns *that in virtue of which* the considerations that ultimately bear normativity – whichever they are – do so. Another question is: What is the nature of normativity? That is, is normativity an irreducibly distinct justificatory force, a motivational force, or a volitional force? This question concerns normativity's essential features. The question of source, by contrast, concerns *that in virtue of which* something has normative force, whatever the nature of this force.

The questions of the source, bearers, and nature of normativity are logically distinct but naturally related. If normativity is an irreducibly distinct justificatory force, then it is natural to think, as Plato did, that its ultimate bearers are irreducibly normative facts and that its source is an irreducible normative reality. If normativity is a motivational force, then

[1] Desire-based and value-based views are the main ways in which philosophers have attempted to systematize the ultimate bearers of practical normativity. For a state-of-the-art discussion of these two views, see Derek Parfit (forthcoming). I explore a hybrid view in Chang (2004c).

243

one might conclude, as Hume did, that its ultimate bearers are desires and that normativity derives from a relation involving our desires. And if normativity is some kind of volitional force, perhaps Kant was right: normativity has its source in the will, and the ultimate bearers of normativity are facts about the consistency of willing actions. But none of this clustering of views is forced upon us.

The question of source is sometimes obscured because it gets folded into questions of the nature and ultimate bearers of normativity – hence my concern to distinguish those questions. But it also gets obscured because, for many philosophers, the question of source is illegitimate, or if legitimate, has only a degenerate answer. Such "normative externalists," as I will call them, think there is, strictly speaking, *nothing* in virtue of which a consideration is reason-providing – there are just the irreducibly normative facts that such-and-such considerations provide reasons.[2] Nevertheless, such externalists must allow that there is a sense in which the "source" of normativity is found in irreducibly normative facts. It is in this sense that the fact that an experience is painful provides a reason to avoid it in virtue of the irreducibly normative fact that being painful provides a reason to do so. Normative externalists, then, can be said to locate the source of normativity in a realm of external, irreducibly normative facts.

Normative internalists, by contrast, locate the source of normativity in mental states internal to us, and in particular, in desires and dispositions to which we are for the most part passively related.[3] A consideration has the practical normativity of a reason in virtue of its serving or furthering our procedurally constrained desires or dispositions. So, for example, being painful provides a reason to avoid it in virtue of one's fully informed desire to avoid pain. The desires that provide the source of normativity might be ones we must have in order to be rational or indeed to be agents at all, but they are not themselves states of agential activity.

Finally, normative voluntarists locate the source of normativity in us, but not in our passive states. Rather, normativity has its source in something we *do*, and, in particular, in our active attitudes of willing or

reflective endorsement.[4] By willing something, that is, by actively engaging our volition, we can give a consideration the normativity of a reason. So, for example, if we will a law or principle according to which we avoid pain, the fact that an experience is painful can thereby be action-guiding. According to normative voluntarists, the source of practical guiding is to be found neither in irreducibly normative facts nor in normativity is to be found neither in irreducibly normative facts nor in non-cognitive states towards which we are passive. Normativity is rather borne of activity of the will.

Of these three main approaches to the source of normativity, I think it is fair to say that voluntarism is the runt of the litter. It strikes many philosophers as too bizarre to be taken seriously. How can we magically endow a consideration with the normativity of a reason simply through an act of will? Can we really turn a consideration into a reason – giving it action- or attitude-guiding force – simply by willing something? But some philosophers, most notably contemporary neo-Kantians such as Christine Korsgaard, have offered ingenious arguments in defense of our having such normative powers.[5] And voluntarism has an impressive pedigree reaching back to Hobbes, Locke, Pufendorf, and Kant, among others. Thus, the starkest divide in approaches to normative source is arguably between those who think we can create normativity and those who think that normativity is somehow given to us.

In this essay, I try to make sense of this divide by suggesting a form of voluntarism that straddles it. It seems to me that Hobbes, Locke, Kant, and their progeny offer us an important insight about normativity that we should not ignore. This is the thought that we *can* confer normativity or value on things. At the same time, externalists and internalists are right to be suspicious of the claim that all normativity derives from the will. Surely some normativity is given to us and not created by us.

Voluntarism is most plausible, I suggest, if we understand the normativity we create as hierarchically related to the normativity we do not create. Thus "hierarchical voluntarism" rejects an assumption that normative externalists, internalists, and standard voluntarists all share,

---

[2] Normative externalists include Clarke (1706), Ross (1930), Prichard (1968), Nagel (1970), Moore (1971a), Scanlon (1998), Raz (1999), Dancy (2000), Shafer-Landau (2003), Wallace (2006), Wedgwood (2007), Parfit (forthcoming) among others.

[3] Normative internalists include Hume (1978), Williams (1981), Falk (1986), Railton (1989, 2003), Smith (1994, 1999), Brandt (1996), Broome (1997), Velleman (2000), and, arguably, Rawls (1971).

[4] Since the distinction between "passive" and "active" mental states is notoriously problematic, the line between internalists and voluntarists may not be a sharp one. For paradigmatic examples of voluntarism, see Hobbes (1651), Pufendorf (1672), Kant (1785), and Korsgaard (1996b). Korsgaard has done the most to raise the question of normative source to contemporary philosophical consciousness, though sometimes even her discussion of source is folded in with other issues about normativity, such as its nature.

[5] See Korsgaard (1996b).

namely, that practical normativity has an univocal source. According to this assumption, *all* practical reasons have their normative source in irreducibly normative facts, or in the agent's desires, or in her will. While standard forms of voluntarism maintain that all practical (or moral) reasons have their normative source in the agent's will, hierarchical voluntarism maintains that voluntarist reasons – reasons whose normative source is an act of will – depend on there being *non-voluntarist* reasons – reasons whose normative source is not an act of will. More specifically, it holds that an agent cannot have a voluntarist reason unless her non-voluntarist reasons have "run out." If this is right, then the answer to the question, "In virtue of what is a consideration reason-providing?," is not univocal; sometimes a consideration is normative in virtue of an act of will and sometimes not. In this way, practical reason is marked by a deep duality in its source.[6] This duality suggests a fundamental difference between practical and theoretical reason, but I won't be pursuing that issue here.

My case for hierarchical voluntarism centers on two puzzles about human rationality and agency. The first might be glossed as follows: How can we have *most* reason to do something when our reasons have in some sense "run out"? Sometimes it seems that our reasons fail to determine what we should do, and yet further deliberation determines that we have most reason to choose one alternative over the others. How can this be? The second puzzle begins with the thought that we *make*, through an act of agency, our ideal rational selves – you make yourself into someone who has most reason to spend weekends at wild parties, and I make myself into someone who has most reason to spend weekends quietly reading books at home. But how can we *make* ourselves into agents with these distinctive ideal rational selves if we are rationally required to follow our reasons? I will have more to say about these puzzles in due course, but for now it's worth noting that they are largely assumed to be unrelated. Those interested in the first tend to focus on questions about the "incommensurability" or "incomparability" of reasons or values. Those interested in the second tend to focus on how an agent can be "self-governing" or the "author of her own life." What I want to suggest is that the puzzles are linked by a common, attractive solution found in hierarchical voluntarism.

Any defense of voluntarism must confront what is commonly taken to be its fatal flaw. The putatively fatal objection, raised by Clark against Hobbes, and more recently, by Gerald Cohen against Korsgaard, goes something like this: IF the source of the normativity of our reasons is in our willing something, then practical reasons become objectionably arbitrary. The Mafioso, for example, can in principle will the violent deaths of his enemies and thereby create reasons for himself to bring about those deaths. But he has no such reasons. Therefore voluntarism is false. Part of the task here will be to clarify how this objection should be understood. As we will see, while standard forms of voluntarism fall prey to this objection, hierarchical voluntarism does not.

I end the essay with some speculative remarks about how hierarchical voluntarism fares better than its standard relatives in separating views about the source of normativity from views about its nature. As we will see, standard forms of voluntarism are held hostage to two of the three main conceptions of normativity, while hierarchical voluntarism is arguably plausible no matter which conception of normativity is correct. In this way, hierarchical voluntarism can provide the will a secure place in understanding the source of normativity, however normativity itself is to be understood.

I

A voluntarist reason is one whose normativity derives from an agent's act of will. But which act of will? The simplest form of voluntarist reason is one an agent might have by willing – or as I shall say, "taking" – a consideration to be a reason. Suppose you take the dulcimer tones of the harp to be a reason to play the harp. By taking the dulcimer tones to be a reason, you can – under suitable conditions – make this consideration a reason for you to play. You can create a new voluntarist reason to play the harp through an act of will. The bearer of normativity is, we are supposing, the fact that the tones of the harp are pleasant. That fact gives you a reason to play in virtue of your taking it to be a reason – in virtue of your act of will. I am going to take this simple form of voluntarist reason as my working model. The view I want to explore, then, is that under suitable conditions, by taking something to be a reason you can thereby endow that thing with the normativity of a reason. Why should we believe that we have such normative powers?

---

[6] Sidgwick (1907) famously thought that practical reason is marked by a duality, but his was a duality in the *kinds* of reasons or values relevant to action. He thought that reasons of individual prudence were wholly incomparable with reasons of impartial beneficence – that there was no way to put them together by a normative relation such as "stronger than." (I suggest some reasons to think that there isn't this duality in Chang [2004a].) But Sidgwick was a monist about the source of normativity since he located normativity in irreducibly normative facts about value or "points of view." See Sidgwick (1907), book III, chapter 14, book IV, chapter 6.

Sometimes the reasons in a choice situation fail to determine what one should do. It might turn out, for example, that some reasons favor pursuing a career as a lawyer while others favor pursuing a career as a scuba diving instructor, and there is no all-things-considered conclusion about which career one has most reason to pursue. I shall say that in such cases one's reasons have "run out." Still, one must make a choice. In some such cases, when one chooses – say, to become a lawyer – it seems that one has *most reason* to choose it. How can this be?

To understand the puzzle, we need to start with the idea of reasons "running out." Sometimes the reasons for choice *determine* what you should do. In this case your reasons deliver the conclusion that you have *most* reason to choose one alternative over the others. You have most reason to choose x over y if the reasons for x outweigh, trump, silence, exclude, cancel, bracket, or are more stringent than the reasons for y. Your choice of x is rationally determined.

Sometimes, however, the reasons for choice *underdetermine* what you should do. This can happen in one of two ways. Your reasons might deliver the conclusion not that you have most reason to choose one alternative but that you have *sufficient* reason to choose any of two or more alternatives. You might have sufficient reason to choose x or y if the reasons for x are equally as weighty as those for y, or if they are "on a par," or, perhaps, if they are incomparable.[7] In this case, although the choice of either is justified, your reasons don't determine *which* you should choose. You have multiple rationally eligible options but no reason to choose one over the others. Whichever you choose, your choice, though justified, is rationally underdetermined. As some philosophers say, all you can do is "pick" rather than "choose."[8] When you pick, you justifiably choose on the basis of reasons, but your reasons do not determine your choice.

A more radical way your reasons might underdetermine what you should do is by failing to deliver the conclusion that you have most *or* sufficient reason to choose either alternative. When your reasons fail to deliver any justified choice whatsoever, we might say that they *break down*. When your reasons break down, there is no justified choice to be had; whichever alternative you choose, your choice will be beyond the scope of practical reason. Some philosophers think that reasons never

break down. Those who think they do tend to think they do when the alternatives are incomparable. Indeed, some incomparabilists define incomparability as the failure of practical reason to deliver any justified choice among the alternatives. In any case, if reasons break down, all you can do is "plump" for no reason rather than justifiably pick or choose on the basis of reasons.

To summarize. There are two ways your choice can be justified: when you have most reason to choose that alternative or when you have sufficient reason to choose it among others. But there is only one way a choice can be rationally determined: when you have most reason to choose that alternative. When your choice is not rationally determined – that is, when you have sufficient reasons to choose among several alternatives or when your choice is beyond the reach of practical reason – your reasons *run out*. Reasons run out when they fail to deliver a univocal answer to the question, "What should I do?" When reasons run out, it seems that all one can do is pick or plump for no reason.

Some philosophers deny that reasons ever run out; all choice is rationally determined.[9] For the purposes of this essay, I am going to assume that reasons do run out, and that the cases in which they do range from the most mundane – what dessert to eat, what to wear, how to answer social invitations – to the most profound – how and where to live, which career to pursue, whether to have children, with whom to make a life, and so on. If this is right, underdetermination by reasons is a wide and deep phenomenon. If the assumption about the scope and significance of cases in which reasons run out is mistaken, then the scope and significance of voluntarist reasons will have to be adjusted accordingly. Insofar as there are cases in which reasons run out, however, there is room for voluntarist reasons, or so I will now try to show.

2

Suppose that you are faced with a choice between a career as a philosopher and one as a trapeze artist. You have investigated each career from every angle, vividly imagined yourself writing philosophy articles and swinging

---

7 For a description of the distinction between equality, parity, and incomparability, see Chang (1997) and (2002). I believe that one's non-voluntarist reasons must be "on a par" in order for voluntarist reasons to have a role in practical reason. But that is another story.
8 See Ullmann-Margalit and Morgenbesser (1977).

9 Some philosophers who deny that reasons ever run out suggest that they appear to do so only because they are underspecified or not fully articulated. See Taylor (1985); Richardson (1994); Millgram (1997); and Helm (2001: esp. chapter 6). Insofar as some of these authors go on to suggest that one way of further specifying or articulating these reasons is through an act of will, their views are congenial to the spirit of the view presented here.

under the big top, carefully considered and re-considered the reasons for and against each career, thought long and hard about how the reasons for and against each relate, sought advice from people whose judgment you respect, and so on. Suppose that, as a result of careful and thorough investigation, you come to believe that, all things considered, the reasons for and against each career have run out.

What should you do? It seems it would be a mistake for you simply to pick or plump for one career. Even if you have sufficient reason to choose either career and the choice of either would be justified, it seems odd to think that you might simply pick between them as you might between cans of soup. And if the choice between careers is beyond the scope of practical reason, plumping for a career for no reason also seems a mistake. Rather, it seems what you should do – and what many seemingly perfectly rational agents *do* do – is to continue to deliberate in the hopes of coming to a conclusion about which one has most reason to choose. Indeed, it seems perfectly possible that further deliberation might lead to a rationally determined choice.

But now we have a problem. If, as it seems, one's reasons have run out, how can it be appropriate to continue to deliberate about which alternative to choose in the belief that further deliberation might determine which one has most reason to choose? If one has sufficient reason to choose either career, then one does not have most reason to choose either. One should simply pick one, and further deliberation would be irrational. And if the choice between the careers is beyond the scope of practical reason, then again further deliberation would be pointless. One should simply plump for one career for no reason.

As stated, the problem has a ready solution. Although we can grant that reasons sometimes run out, perhaps we can never *know* that they have in any particular case. So it is always appropriate, other things equal, for an agent who believes that her reasons have run out to revisit her reasoning to make sure that they have. Thus when you continue to deliberate about the two careers, you are, as it were, checking your sums. Maybe you failed to give one of the reasons its proper weight or overlooked a detail that might turn out to be important to your decision. In this way, it can be appropriate to deliberate further when you believe that your reasons have run out. Further deliberation, as a corrective measure, can lead to a rationally determined choice.

But we can reformulate the problem in a way that sidesteps this epistemic maneuver. Although we may never be in a position to *know*, in some strong sense of "know," that our reasons have run out in any

particular case, we can, however, be *practically certain* that they have. If you are practically certain that p, it is irrational for you to act on the assumption that not p. It would be irrational, for instance, to revisit your deliberation about whether p. If, for example, you are practically certain that you turned off the lights, it is irrational for you to check to see whether you did. Practical certainty might be understood as knowledge that is relativized to a practical context.[10] What it takes to know that you turned off the lights might be very different from what it takes to know that you turned off the gas oven. Practical deliberation is often like this; we reach a point in deliberation at which it would be practically irrational to second-guess ourselves.

Now suppose that after careful and thorough investigation of your reasons you are practically certain that the reasons for and against the philosophy and circus careers have run out. You conclude with practical certainty that they fail to determine that you have most reason to choose one of the careers. We might say that you know, for practical purposes, that the reasons have run out. Nonetheless, it seems that it could be perfectly appropriate to continue to deliberate about which career you should choose. How can this be? On the one hand, it is practically irrational to revisit the matter. On the other, it can be appropriate to continue to deliberate about which to choose.

Our first puzzle, then, is this. How can it be appropriate to continue to deliberate about which alternative to choose when one is practically certain that one's reasons have run out?

The puzzle can be deepened. Suppose you have a choice between chicken pot pie and roasted squab *au jus* for dinner. After careful and thorough deliberation about your reasons – fill in the details as you like – you are practically certain that your reasons have run out. Unlike in the careers case, however, in this case it seems perfectly appropriate to pick or plump for one of the dinners. Indeed, further deliberation might be irrational. But how can picking or plumping be appropriate in the dinners case but not in the careers case? By hypothesis, the structure of the reasons in each case is the same – you are practically certain that the reasons have run out. Given that reasons run out in a wide variety of cases, how is it that in some cases it can be appropriate to pick or plump while in other cases it is not?

Our puzzle, then, has two aspects. First, how can it be rational to continue to deliberate in the belief that this deliberation may lead to a

---

[10] For a view of knowledge that contextualizes it to the agent's practical interests, see Stanley (2005).

rationally determined choice when one is practically certain that one's reasons have run out? Second, sometimes when one is practically certain that the reasons have run out, it can be appropriate to pick or to plump, but other times it is not. What could explain the difference between such cases given that in both one is practically certain that the reasons have run out?

It might be suggested that when a decision, such as one between careers, is important, it is appropriate to continue to deliberate in the hopes of arriving at a justified choice. Decisions about important matters should not be a matter of picking or plumping. And since a decision between careers is important and one between dinners is not, this explains why it is appropriate to pick or plump in the one case but not in the other.

This suggestion, however, falls short of what is needed. Granting that a decision between careers is typically important, why should this make it appropriate to continue to deliberate if, by hypothesis, one is practically certain that one's reasons have run out? How can the importance of the decision be relevant to the structure of the reasons there are for deciding between two alternatives?[11] And given that the structure of reasons in both the careers and dinners cases is the same, how can the importance or lack of importance of the decision – presumably irrelevant in both cases – explain why picking or plumping can be an appropriate response in the one case but not in the other? That would be like trying to explain why it is appropriate to pick or plump between dinners but not careers by pointing out that one is wearing red shoes.

3

So how is our puzzle to be explained? To keep things simple, let's focus for now on its first aspect. How can it be appropriate to continue to deliberate and to believe that further deliberation might lead to the conclusion that one has most reason to choose one alternative over the other when one is practically certain that one's reasons have run out?[12]

---

[11] Even if the importance of the decision could be made to be relevant to the reasons for and against each alternative, this factor should then be reflected in the structure of the reasons which, by hypothesis, have with practical certainty run out.

[12] It should not be thought that the appropriateness is a matter of instrumental rationality. The appropriateness of continuing to deliberate is not like the rationality of some cases of self-deception or wishful thinking in which being in the state is good in some way. It is, to borrow a distinction made by Derek Parfit, object-given, not state-given. See Parfit (forthcoming).

Two thoughts appear to be in tension. On the one hand, one's reasons have run out, and therefore one doesn't have most reason to choose one alternative over the other. On the other hand, further deliberation might lead to the conclusion that one has most reason to choose one alternative over the other. How could both these thoughts be correct?

The most promising strategy for pursuing a solution to our puzzle, I believe, is to allow that deliberation about what to do can have two distinct stages. At one stage of deliberation, one might be practically certain that one's reasons have run out. Reasons underdetermine which alternative one should choose. But there might be a second stage at which further "deliberation" of a different kind can lead to a rationally determined choice. If this is right, the key to a solution is to determine what these two stages might be.

With this idea in mind, an immediate suggestion might be that the distinction between two stages of deliberation is given by the distinction between comparative and non-comparative justification of choice. Perhaps at one stage of deliberation we determine whether we have most reason to choose an alternative because it is the *best* of all the others, while at a second stage we determine whether we have most reason to choose it because of its intrinsic or deontic – i.e. non-comparative – features. So perhaps in the careers case you are practically certain that neither career is better than the other and thus practically certain that your comparative reasons have run out. You are practically certain that you don't have most reason to choose one career over another because it is better than all the rest. But at the second stage of deliberation you can consider the non-comparative reasons for choosing one of the careers. Perhaps pursuing the philosophy career would display the virtue of nobility and joining the circus would fulfill a promise you made. In this case, further deliberation about one's non-comparative reasons could lead one to a rationally determined choice.

But this misunderstands the puzzle. When you deliberate about the reasons for and against the two careers at the first stage, you do not artificially restrict your deliberation only to that concerning your comparative reasons. If you promised your mother long ago that one day you would run away and join the circus, that is one of the reasons you take into account in the first stage of your deliberation. Your deliberation includes both comparative and non-comparative reasons, and your conclusion is that these reasons have run out. The puzzle is how we can be practically certain that both our comparative and non-comparative reasons have run out and at the same time appropriately continue to

deliberate with the belief that this further deliberation might tell us what we have most reason to choose.

A more promising view is suggested by remarks of Joseph Raz.[13] Raz suggests that when reasons run out, the *will* plays a role in determining what to do. Suppose, to borrow one of Raz's examples, you are trying to decide between a banana and a pear for dessert. The banana has got a lot of potassium, which is good for your heart, and the pear has got a lot of vitamin C, which is good for your immune system, but there is no all-things-considered truth about which you have most reason to choose. Your reasons have run out. (Raz assumes that a rational deliberator can know this.) Raz thinks that you can then *will* to have the banana, which, in the normal course of events, *causes* you to choose the banana. So you choose the banana.

There are two possibilities here.[14] Your willing the banana might make it true that you have most reason to choose the banana or it might not. If it doesn't, Raz's view does not help us to solve our puzzle, for we need to explain how it can be true that you have most reason to choose an alternative when your reasons have run out. Raz's view would simply give a causal explanation of how one comes to choose an alternative when reasons have run out.

Could willing an alternative explain how one could have most reason to choose that alternative? It is hard to see how it could. To think that it could would involve double-counting one's reasons. When you conclude that the reasons for and against the banana and pear run out, you have that the reasons in favor of the banana, namely that the banana is full of potassium and good for your heart. At the first stage of deliberation, then, the fact that the banana is full of potassium and good for your heart does not give you most reason to choose it. At the second stage of deliberation, after you will to have the banana, the very same reason that failed to give you most reason to choose the banana at the first stage gives you most reason to choose it in the second. But this is just to count the reason for choosing the banana twice. Why shouldn't the reason in favor of the pear similarly be counted twice?

[13] See Raz (1997, 1999: chapters 4–5).
[14] I find it difficult to know which alternative best represents Raz's intended view. Raz says that the choice of the banana is "rational," but it was also rational before one willed it because one has sufficient reason to choose the banana or the pear. There may be an ambiguity here in "rational"; a choice can be "rational" in the strong sense of being rationally determined – that is, being what one has most reason to choose – but it can also be "rational" in the weak sense of there being sufficient reason to choose it among other alternatives. Our concern is to explain how a choice can be rational in the strong sense when one is practically certain that one's reasons have run out.

We could try modifying the view to avoid this problem by suggesting that willing an alternative adds normative weight or significance to the reasons that support the alternative one wills. Return to our careers case. When you are practically certain that your reasons have run out, by willing to join the circus, you can give the reasons in favor of the circus career extra normative force. The extra force of these reasons can then give you most reason to choose that career.

The trouble with this suggestion, however, is that it entails that one always has greater reason to choose an alternative after one wills it than before one wills it, and this isn't always the case. Suppose you have a choice between two evils – say, betraying a friend and causing a stranger to suffer physical pain. After careful investigation, you are practically certain that your reasons have run out. Suppose now that you *will* that the that your reasons have run out. Suppose now that you *will* that the stranger suffer. It doesn't follow that you thereby have greater reason to cause his suffering than you had before your act of will. It might be that because you now will his suffering, you have less reason to choose this alternative than you did before you willed it. Willing the suffering of others may make the suffering of others a worse alternative than it was before one willed it. Whether one has more reason to choose an alternative after one wills it is a substantive matter that should not be built into an account of practical deliberation.

Although this last suggestion fails to explain our puzzle, it goes some way towards what I believe is its correct solution. It gives a role to the will in the second stage of deliberation and, more importantly, recognizes that the will can be a source of normativity. Our puzzle, however, requires us to explain how further *deliberation* is appropriate, and it is unclear how willing an alternative can be a form of deliberation. As I now want to suggest, the proper role of the will in the second stage of deliberation is not to will an alternative but to will a reason that supports an alternative. This willing creates normativity by creating new reasons whose normativity derives from the very act of will. And as we will see, creating reasons through an act of will is part of a deliberative process of making oneself into the distinctive rational agent that one is.

4

From our discussion so far, we can extract two general principles for our puzzle's solution. First, in thinking about what to do, there may be different stages of deliberation. If, at one stage, one is practically certain that one's reasons have run out, there may nevertheless be another stage in

which further deliberation yields a rationally determined choice. Second, this further stage of deliberation is one in which the will has some role to play. Our questions then are how to understand these two stages and what the role of the will at the second stage might be. Voluntarist and non-voluntarist reasons, I believe, provide an attractive way of answering these questions.

Our non-voluntarist reasons are the reasons we ordinarily take ourselves to have – reasons whose normativity derives either from normative reality or from our desires, but not from our own act of will. They typically include facts about the alternatives, about ourselves, and about the relation between the alternatives and ourselves. So, for example, your non-voluntarist reasons to choose the philosophy career might include the fact that philosophy is a noble pursuit, that you have a strong desire to better understand a particular philosophical problem, that you are especially suited to abstract thinking, that you would enjoy teaching Plato to undergraduates, and so on. Since these are the reasons given to us and not given by us, we might call them our *given* reasons. Our voluntarist reasons, by contrast, are the reasons we create for ourselves by taking a consideration to be a reason when our given reasons have run out. Thus when your given reasons for choosing one of the careers run out, you can take a consideration in favor of one of the careers as normative for you, thereby creating a new, voluntarist reason to choose that career.

I suggest that given reasons have a role in the first stage of deliberation, and voluntarist reasons have a role in the second. Our given reasons are the only reasons we have at the first stage of deliberation, and when they run out, there can be no rationally determined choice on the basis of those reasons. But it does not follow that a rationally determined choice is precluded. There is a second stage of deliberation in which we can create new voluntarist reasons which, in conjunction with our given reasons, may deliver an alternative we have most reason to choose. In this way, it can be appropriate to continue to deliberate to a rationally determined choice when we are practically certain that our reasons have run out.

To see how the proposal might work, return to the careers case. Suppose, for simplicity, that there are only two given reasons relevant to the choice between them. In favor of the philosophy career is the intellectual satisfaction you would get from a life of contemplating deep philosophical questions. In favor of the career as a trapeze artist is the thrill you would get from nightly daredevil stunts under the big top. Suppose that after careful and thorough investigation of these reasons, you are

practically certain that they fail to deliver an all-things-considered truth about what you have most reason to do. Your given reasons have run out.

At this stage in your deliberations it can be appropriate for you to take a consideration as a reason in favor of one of the careers, thereby giving yourself a new, voluntarist reason to choose that career. You might, for instance, take your cousin's whim that you join the circus, or the secret delight you would get from wearing sequins, or the financial gain that would accrue to the manufacturers of trapeze rope in Korea as a reason to choose the circus career. Although your cousin's whim, your secret delight, and the manufacturers' pecuniary gain are not, by hypothesis, relevant to your choice between the careers, when your given reasons have run out you can, through an act of will, make them reasons that are relevant to your choice.[15] You can take any consideration that counts in favor of an alternative, even if irrelevant to the choice, as a reason for you to choose that alternative. In this way you can create a voluntarist reason you didn't have before that might then give you most reason to choose the circus career.[16]

You might instead take the thrill of performing stunts under the big top as a reason to choose the circus career. In this case the voluntarist reason you create shares its content with a given reason you already have. This might seem strange. How can a single consideration – the thrill of performing under the big top – provide you with both a given and a voluntarist reason? Doesn't this amount to double-counting your reasons? There is double-counting of reasons only if reasons are individuated by their contents alone, but it is unclear why reasons should be individuated in this way. Why shouldn't a single consideration have normativity with two different sources and thus provide two distinct reasons? Consider an

15 Some philosophers think that in any given choice situation, *every* consideration that counts in favor of the alternative is a reason relevant to a choice involving that alternative. So, for example, the fact that ordering steak for dinner in New York will help, in a Rube Goldberg-like way, a child in Argentina with her homework is a reason, albeit a "small reason," to order it. I have offered some considerations against this view in Chang (2004b). But if the view is correct, then the two cases described below ultimately describe a single kind of case. If every consideration that counts in favor of the alternatives is already a given reason at the first stage of deliberation, then the role of the will should be understood instead as creating new, voluntarist reasons that share content with one's given reasons.

16 It might be objected that if one's voluntarist reasons must always be reasons which antecedently to any act of will count in favor of choosing an alternative, the source of their normativity is not truly the will but in whatever it is in virtue of which the reason counts in favor. This is to conflate the content of the reason with the status of being a reason. While a given content cannot count as a reason *as a conceptual matter* unless that content counts in favor of the action, it does not follow that the source of the normativity of that reason cannot be the will.

analogy with physical force. A single object – say, an apple – can exert two distinct physical forces on another object, one gravitational and the other electromagnetic. The forces are distinct because the "sources" of those forces are distinct. Similarly, a single consideration – the thrill you would get – can provide two distinct reasons in virtue of two distinct sources of normativity. It might be a fact of normative reality, for instance, that you have a reason to pursue what thrills you, and it might also be the case that by taking your thrill as a reason, you can endow it with the normativity that derives from your will. Just as an apple can exert two distinct forces on an object, a consideration can provide two distinct reasons for action. The considerations that provide the contents of your given reasons may also provide the contents of your voluntarist reasons.[17] So sometimes what we take to be a reason is a consideration that already provides a relevant given reason and sometimes not.

This view about the possible contents of one's voluntarist reasons fits nicely with the phenomenology of choice in the cases of interest. Sometimes when we are practically certain that our reasons have run out, further deliberation takes the form of *agonizing* over the considerations that provide the contents of our reasons. When you are practically certain that your reasons for the philosophy and circus careers have run out, you may nevertheless find yourself agonizing over the intellectual satisfaction you would get from the philosophy career and the thrill you would get from the circus one. Agonizing over these considerations – otherwise inexplicably irrational – can be understood as deliberating over whether and which such considerations to create as new, voluntarist reasons for yourself. Other times, further deliberation takes a different form – we focus on factors that count in favor of an alternative but aren't by our own lights relevant to the choice. For example, your cousin's whim that you join the circus is, by hypothesis, irrelevant to your choice between the two

---

[17] Although, I believe, reasons should not be individuated by their contents, there is a variant of the view presented here that could avoid relying on this claim. Instead of creating new, distinct voluntarist reasons, perhaps an act of will can transform given reasons into reasons *part* of whose normativity is a creation of the will and *part* of whose normativity is given. It might do this by endowing one's given reasons with special normative weight so that we have the "same" reason but with added normativity that derives from the will.

Note that this variant does not run afoul of the difficulty raised for the modified version of Raz's view. That view had the implication that by willing an alternative, one thereby had more reason to choose it than one had before one willed it. Willing more weight or significance for a particular given reason does not have this implication because it allows that even though one creates greater weight for a particular reason, the reasons overall to choose the alternative after willing may remain unchanged or even be weaker due to organic unities and the like. It correctly leaves the normative relations among reasons a substantive matter.

---

careers. But you might nevertheless take her whim as reason-providing for you and thus make it a voluntarist reason to choose the circus career.

It might be wondered how deciding whether to take a consideration as normative for oneself can be a form of deliberation. Unlike the sheer willing of an alternative, willing a consideration to be a reason is part of the process of making oneself into a distinctive normative agent, that is, creating one's own "rational identity." I will have more to say about this below, but for now we can note that if the creation of voluntarist reasons is a part of rational self-governance, we have a tidy explanation not only of the first but also of the second aspect of our puzzle.

Recall that in some cases when we are practically certain that our reasons have run out, it seems appropriate to pick or plump, while in other cases further deliberation seems to be in order. Typically, it seems appropriate to pick or plump between dinners but not between careers. If creating reasons for oneself is making oneself into a distinctive normative agent, then the difference between the dinners and careers cases can be explained by the role such choices play in one's rational identity. "Who you are," rationally speaking, typically involves which career you pursue but not which dinner you eat. Choices between dinners are not usually occasions on which one "makes" oneself into the distinctive rational agent one is. But they might be. A world-class chef might create reasons for herself to choose one dinner over another when her given reasons have run out. But since most of us have normative identities that are not bound up with what we have for dinner, it seems appropriate to pick or plump between dinners. Careers are a different matter. Your rational identity may very much be a matter of the career you pursue. If it is, then when your given reasons to choose between careers have run out, you have an opportunity to make yourself into the kind of agent for whom there are stronger reasons to fly under the big top than to sit at home contemplating the nature of reasons.

At the heart of our puzzle is a common assumption about practical deliberation. Deliberation is a matter of discovering, recognizing, investigating, appreciating, and engaging with the reasons there are. If one is practically certain that one's reasons have run out, then deliberation seems to have finished its job. How could further deliberation yield the result that one has most reason to choose one of the alternatives? Once we allow that deliberation may also involve the *creation* of reasons – and as we shall see, forming our ideal rational selves – the puzzle disappears. While it can be true at the first stage of deliberation that one's given reasons have run out, one can create new, voluntarist reasons at a second stage that may

then lead to a rationally determined choice. In this way, voluntarist reasons explain how we can have most reason to choose an alternative when our (given) reasons have run out.

5

There is a second puzzle to which voluntarist reasons seem to provide a tidy solution. How can rational agents both be fully responsive to reasons and at the same time exercise control over who they are, rationally speaking? Each of us has the power to make ourselves into a person with a distinctive *rational identity*, as someone who, for example, has stronger reasons to become a scuba diving instructor rather than a lawyer, or to spend her leisure time fly fishing rather than listening to opera. But if being rational is responding appropriately to our reasons, it is not clear how *we* can get enough distance from our reasons to be able to make ourselves into one kind of distinctive rational agent rather than another. The "we" of agency seems to disappear; there is only the rational agent's responsiveness to her reasons but no active *making* of oneself into one kind of rational being rather than another.

To understand this puzzle we need first to understand the notion of a rational identity. Your rational identity is your ideal rational self. It is a function (perhaps a simple conjunction) of the reasons that determine what you have most reason to do in actual and hypothetical choice situations. Consider the whole range of possible choice situations you and I might face. What you have most reason to do in those situations is different from what I have most reason to do. We have distinctive rational identities. This set of rationally determining reasons can be described in a rough-and-ready way by general roles rational agents can play – tinker, tailor, soldier, spy – and general attributes they can possess – extrovert, fun-loving, people-oriented, party-going. Your rational identity, in short, is who you would distinctively be were you perfectly rational. It is the rationally angelic you.

Your rational identity need not be how you conceive of yourself. You might think of yourself as shy and retiring, but if you have most reason to choose bungee jumping and karaoke over book-reading and museum-going, your rational identity will not match your self-conception. In this way your rational identity differs from what Christine Korsgaard calls your "practical identity." For Korsgaard, a practical identity is a "description under which you value yourself, a description under which you find your life to be worth living and your actions to be worth undertaking – it

is how you see yourself."[18] A rational identity, by contrast, is not a self-conception but rather a description of your normatively ideal self – a loosely unified way of understanding the reasons that justify doing what you have most reason to do.

Nor is your rational identity your personality or character. You might be a self-aggrandizing, malevolent egomaniac and thus lie, cheat, and steal for your own advantage. But since, presumably, you don't have most reason to lie, cheat, and steal, the reasons you have to perform these actions are no part of your rational identity. Moreover, the reasons you might mistakenly think you have – even to do what you in fact have most reason to do – are no part of your rational identity. Your rational identity is constituted by the reasons you objectively have that rationally determine what you should do. In this way, your rational identity is not given by what Michael Bratman calls your "self-governing policies." A self-governing policy is "a policy of treating a desire as providing a justifying end in motivationally effective practical reasoning" – in short, a policy of treating certain considerations as reasons.[19] You might have a self-governing policy of treating your desire to take advantage of others as a justifying reason to lie, cheat, and steal. But if you don't objectively have such a reason, this desire – and the self-governing policy that underwrites it – are no part of your rational identity. Your rational identity is the ideally rationally agent within, not the rational agent you subjectively take yourself to be.

Most philosophers concerned with explaining "self-governance" and being "the author of one's own life" focus on how we make ourselves into creatures with the foibles and peculiar irrationalities or personality that we each actually have.[20] But the sort of self-governance of interest here is not a matter of how we make ourselves into the actual, rationally flawed creatures that we are, but of how we make ourselves into creatures – however flawed – with certain *ideal rational identities*. Through self-governance, I make it true that I have most reason to spend my Saturdays collecting stamps and you make it true that you have most reason to spend your

---

[18] Korsgaard (1996b): 101. Korsgaard argues that our "local and contingent practical identities" derive their normativity from our "human identity" – the practical identity we cannot but help have as self-conscious, reflective agents. Our practical identities in turn determine our reasons. According to Korsgaard, we have the reasons we do because of our practical identities. The view of voluntarist reasons developed here is close to the reverse. Our rational identity is something we create through taking considerations to be reasons. We have the rational identity we do because of our voluntarist reasons. Our reasons determine our rational identity, not the other way around.

[19] Bratman (2007): 40 and part I generally. See also Bratman (1999).

[20] See, e.g., Bratman (2007).

Saturdays playing the harp. We are authors not only of our actual lives but also of our ideal rational lives – of the best we can be, rationally speaking. The governing of our ideal rational selves is arguably the central – and most exalted – exercise of rational agency.

But, granting that each of us has a rational identity, why should we think that we "make" those identities? Perhaps our rational identities just happen to us – we come born with them, or they are shaped by our environment. On such a view, the reasons that rationally determine what we should do always come to us unbidden; perhaps they are a causal function of the desires and dispositions with which we are born and the environment in which we act. There is no genuine agency involved in crafting one's rational identity; instead we sit back and relax – our ideal rational selves, like our toenails, grow without any agency on our part. Indeed, perhaps agency itself is nothing more than the "proper [causally specified] functioning of thoughts, beliefs, emotions, and desires" – there is no "you" or "me" apart from the complex causal–functional operation and structure of our various mental states.[21]

This line of thought may, in the end, be correct as part of a reductive account of agency, just as sub-atomic physics may, in the end, provide the ultimate explanation of everything. What is true at this deep level of theorizing is something about which we should probably keep an open mind.[22] But at the level of "shallow" philosophical theorizing about rational identity, such views ride roughshod over the pre-theoretical conviction that there is such a thing as agency, and that agency is somehow involved in our best rational selves being what they are. At this level of theorizing, it is hard to believe that everything I have most reason to do – pursue a life of quiet contemplation on philosophical issues – and everything you have most reason to do – pursue a frenetic life of family, social relationships, and political activism – is just passively given to us,

---

[21] Railton (2003): 200. See also Velleman (2000).

[22] The challenge to reductionists is to provide an account that gives us something recognizably like activity of the will. It will not be sufficient to provide a naturalistic explanation that simply corresponds to or subvenes genuine activity of the will. Nor will it be sufficient to succeed in reducing something that is not recognizably genuine agential activity. My own view is that these two constraints don't leave enough room for the requisite reduction to succeed and that "proper functioning" will be an irreducibly normative matter. It is perhaps worth noting that most naturalists about agency or the will find themselves having to appeal to something that seems itself or through the role it plays suspiciously non-natural, such as "identification," "wholeheartedness," "satisfaction," "endorsement," "adopting a self-governing policy," "treating as a reason," and so on, for the will to do the work it does. In any case, what I say here does not turn on whether agency can be naturalistically reduced.

---

with no agency on our part. Of course the desires we are born with and the environment in which we live play large roles in determining the reasons we have and thus what we have most reason to do. But there is nevertheless some sense in which *we*, apart from the bubbling cauldron of our properly functioning beliefs, desires, and the like, can exercise genuine agency in "making" ourselves into agents with one rational identity rather than another. We need to explain this sense.

6

The problem arises in trying to find a sense in which we can "make" our own rational identities that is compatible with our being fully responsive to our reasons. To see the difficulty, consider three ways you can be fully responsive to your reasons.

Your reasons might give you most reason to choose one alternative; they might give you sufficient reason to choose either of several alternatives; or they might break down, putting your choice beyond the scope of practical reason.

If your reasons break down, then by hypothesis you can only plump for no reason. Plumping for no reason does not involve the sort of agency that goes into making one's own rational identity. This is because making yourself into an agent who has most reason to do one thing rather than another is a practically rational activity and thus within the scope of practical reason.[23]

If your reasons determine that you have *most* reason to choose a particular alternative, then being fully responsive to your reasons requires that you to choose that alternative. How can you then make your own rational identity if, to be rational, you must follow your reasons and do what you have most reason to do? It seems that any "making" of your rational identity in this case would simply be a matter of doing what is rationally required of you. It seems that you don't have enough distance from your reasons in such cases to make yourself into an agent with one set of rationally determining reasons rather than another.

Suppose instead that your reasons give you *sufficient* reason to choose either of two alternatives. Are you now in a position to make your own rational identity? While you are not rationally required to choose one

---

[23] Existentialists, and non-existentialists who make room for "existential commitment," think that when reasons break down, one can existentially plump for an alternative for no reason. But they do not think that by doing so, one is creating one's *rational* identity. See Sartre (1994); cf. Gibbard (1990): 166–168.

alternative over the others, it is not clear how, by picking, you can make yourself into an agent who has most reason to do one thing rather than another. Merely picking what you have sufficient reason to choose does not give you most reason to choose it. By picking, you don't make it true that you have most reason to choose one thing rather than another and thus don't "make" your rational identity.

Of course by picking what you have sufficient reason to choose, you can change what you *subsequently* have most reason to choose. When you have sufficient reason to choose either x or y, you can justifiably pick either. Suppose you pick y. Your choice of y is justified even though you don't have most reason to choose it. Now suppose you are faced with a choice between y and x-minus, only slightly less choiceworthy than x. Again you have sufficient reason to choose y or x-minus. But given that you have previously chosen y, you now, arguably, have most reason to choose y over x-minus – if you had sufficient reason to choose x-minus given your previous choice of y, you could be money-pumped.[24] By picking, then, it seems you can change what you subsequently have most reason to do. And by changing what you subsequently have most reason to do, you change the reasons that comprise your rational identity. Perhaps by picking y you have "made" yourself into the sort of agent who has most reason to y instead of x-minus.

The fact that by picking an alternative you change what you subsequently have most reason to do does not, however, by itself show that you have exercised the sort of agency involved in making yourself an agent with a distinctive rational identity. You can do many things that change what you subsequently have most reason to do – you can have a child, tell a lie, or run round the neighborhood in your underwear. But if you have most reason to do these things, or if doing these things is beyond the reach of practical reason, then even though doing these things changes what you subsequently have most reason to do, you haven't exercised the relevant sort of agency. What you subsequently have most reason to do is grounded not in your will but in something beyond your will. Similarly, if you pick one thing over another, that may change your subsequent reasons. But you don't exercise the relevant sort of agency – you don't make it the case that you have most reason to do one thing rather than another when what you have most subsequent reason to do is grounded in the rational equivalent of a coin flip.

[24] See also Chang (2005).

So the challenge is to explain how doing what one has sufficient reason to do can involve the sort of agency exercised in making one's own rational identity.

Voluntarist reasons provide an attractive answer to this challenge. When you have sufficient reason to do either of two things, you have sufficient *given* reasons, and instead of picking between those alternatives, you can create for yourself a voluntarist reason that may then give you all things considered most reason to choose one alternative over the other.[25] If, for example, your given reasons for the philosophy and circus careers run out, you might take the challenge of understanding a difficult philosophical problem as normative for yourself, thereby giving yourself a voluntarist reason to choose the philosophical career. This voluntarist reason, in conjunction with your given reasons, may now give you most reason to choose that career.[26] In this way, you can make yourself into an agent who has most reason to choose the philosophy career by creating for yourself reasons that rationally determine that choice. In short, you make your own rational identity by creating for yourself some of the reasons which determine it.

Our given reasons leave us what we might call *a space of rational freedom*. This is the freedom to choose one alternative over the others on the basis of reasons, without acting contrary to our all-things-considered given reasons. We might think of our given reasons as drawing a line in the sand around our agency. We can do or feel whatever we like, but we can't go past *here*. Beyond this boundary, we are free to create voluntarist reasons that may give us most reason to choose one alternative over the others.[27]

The sense in which we "make" our rational identities, then, is the sense in which, when our given reasons have run out, we create for ourselves reasons that make it true that we have most reason to choose one alternative over another. We can be fully responsive to our reasons – always doing what we have all-things-considered most or sufficient reason

[25] This is not to say that whenever you have sufficient reason to do either of two things, you can always either pick an alternative or create a voluntarist reason in favor of one. As I've already suggested, I believe there is a specific way in which one has sufficient given reasons that gives rise to the possibility of voluntarist reasons. This is when alternatives are "on a par."
[26] Of course, your voluntarist reasons don't guarantee delivery of a rationally determined choice. If both your given and voluntarist reasons run out, then you can only pick or plump for no reason.
[27] A full story of voluntarist reasons will include an account of further constraints on the creation of such reasons. Sometimes the creation of such reasons is constrained by considerations deriving from the coherence and unity of agency. Rational requirements might be another constraint on the creation of voluntarist reasons. I cannot explore these issues here.

to do and plumping when our reasons break down – and nevertheless make ourselves into agents with distinctive rational identities. We are slaves to some reasons and masters of others.

Imagine two worlds, one with me and the other with my Doppelgänger. Suppose we start out having the same given reasons. And yet I have most reason to spend my free time train-spotting while my Doppelgänger has most reason to spend her free time collecting stamps. It is plausible that our given reasons don't determine how we should spend our leisure time – our hobbies and personal projects are within the space of rational freedom in which we can make ourselves into an agent with one kind of rational identity over another. While I take the sighting of a rare steam train to be normative for me, my Doppelgänger might take the sighting of a rare stamp as normative for her. The same goes for weightier choices. My financially bereft Doppelgänger and I might each find that we have sufficient reason either to keep our only child or to give him up to a loving, wealthy family who will give him a wonderful life neither of us can provide. I might take the importance I attach to my child as a voluntarist reason to keep him.[28] My Doppelgänger may instead take the wonderful life in store for her child as a voluntarist reason to give him away.

When given reasons run out, it is through *taking* considerations to be reasons that we make our rational identities. If, when your reasons run out, you tend to take only considerations that make essential reference to you as normative for you, your ideal rational self might be self-involved. If you take only considerations that make no such reference, perhaps your rational identity is self-effacing. If you tend to take the features of chocolate as reason-providing, your ideal rational self might be a chocoholic. And if, as many philosophers have argued, we often have sufficient given reasons to choose either a partial act or an impartial one, we craft our rational identities as "partialist" or "impartialist" by the reasons we create for ourselves in such cases. If this is right, some of the most intractable ethical problems may have no universally rationally determinate solution; rather they belong in the space of rational freedom enjoyed by each rational agent.[29]

---

[28] Indeed, I might find myself unable to do anything other than take our relationship as a reason not to give him up. In this way, hierarchical voluntarist reasons may provide a general framework for understanding Frankfurt's "volitional necessities." Cf. Frankfurt (1988); and, relatedly, Williams (1995).

[29] This view of voluntarist reasons and rational identity, I believe, also has applications to groups and, in particular, to institutions such as the judiciary and other government-like organizations which can create voluntarist reasons when their given reasons run out. See Chang (2009).

In short, we make our rational identities by giving ourselves voluntarist reasons. This crafting of our distinctive rational identities is, in a way, what life is all about. And voluntarist reasons are, I believe, at the heart of our doing so.[30]

7

Creating for ourselves voluntarist reasons is a fundamental exercise of our normative power. Unlike existential plumping, it is a rational power – one creates a reason that may give one most reason to choose a particular alternative, thereby making one's own distinctive rational identity. And, unlike picking, it is a power to change what we have most reason to do in the very choice situation and beyond. It is by the creation of voluntarist reasons that we have a solution to our two puzzles: we can sometimes have most reason to do something when our reasons have run out, and we can be fully responsive to our reasons and yet nevertheless make ourselves into distinctive ideal rational selves.

But can agents really *create* normativity? At issue here is not the ersatz "normative power" involved in, say, making a promise to meet someone for lunch. There is quite plausibly a normative principle according to which, very roughly, if you successfully communicate to me your willingness to meet me for lunch, you now have a reason to meet me for lunch. By willing to meet me for lunch, it seems that you have given yourself a reason to do so. But this is not the sort of willing involved in a genuine

---

[30] Other philosophers have held views which I take – perhaps in some cases with a dose of wishful thinking – to be broadly sympathetic to aspects of the view presented here. Some related ideas and discussions include Korsgaard's notion of a practical identity from which one's reasons are derived (Korsgaard 1996b: 100–128); Margaret Urban Walker's thought that "narrative understandings of the moral construction of lives" help to explain the responsibilities we have to those with whom we stand in some relation (Urban Walker 1998: 107–129, 109); Jennifer Whiting's claim that we make our metaphysical identities through what we care about (Whiting 1986); Peter Railton's account of agential activity needed to stop a regress in reasons (Railton 2003); Michael Bratman's appeal to self-governing policies in explaining self-governance (Bratman 2007); Raz's view of the role of the will among "incommensurables" (Raz 1997); Thomas Nagel's discussion of "reasons of autonomy" (Nagel 1986: chapter 9); David Velleman's exploration of the "freedom" we have to do something by intending to do it such that the intention causes its own fulfillment (Velleman 2000: 200–220, 213–215); Robert Nozick's claim that "the (precise) weights to be assigned to reasons is 'up to us'" (Nozick 1981: 294–316); and Samuel Scheffler's claim that some values involve our "seeing" things as valuable in order for them to have that value (Scheffler 2001: chapters 6–7). For many of these philosophers, appeal to voluntarist reasons will be uncongenial, but only because they believe they need not make such an appeal to make the claims they want to make. I am unsure whether this is so, and am even less sure if we are to take seriously the two puzzles discussed in this essay. In any case, some philosophers offer what I take to be clearly competing views about the domain I claim is covered by voluntarist reasons. The most interesting, to my mind, is Sartre (1946).

normative power. In willing to meet me for lunch, you simply satisfy the antecedent of a principle that holds when such-and-such conditions are satisfied, one has a reason to do so-and-so. The source of the normativity of this principle may be an irreducible normative reality. If so, your willing to meet me for lunch does not create normativity; rather, the normativity is already given by the source of the normativity of the principle whose antecedent condition you have satisfied through an act of will.[31]

Many philosophers reject the idea that we can create normativity because they think that if the normativity of our reasons derives from our wills, then reasons may be objectionably arbitrary. This was Samuel Clarke's attack against Hobbes' voluntarism and, more recently, Gerald Cohen's objection to Korsgaard's.[32] As Cohen points out, if reasons are voluntarist, then a Mafioso can will himself a reason to murder his rival. But it seems clear that the Mafioso has no such reason.[33]

This objection, usually taken to be fatal, is directed against voluntarist reasons *per se*. But it is unclear how the objection finds its intended mark. Consider the following toy scenario. Suppose the universe is fundamentally chaotic and that the *only* practical reasons there are concern how to tie one's shoelaces – whether to use the single-loop or the double-loop method. All other actions and attitudes not concerned with the tying of shoelaces are a matter of plumping and beyond the reach of reason. Suppose that your given reasons for tying your shoelaces in either of these methods have run out – you either have sufficient reasons to tie in either method or how you should tie is, like everything else, beyond the reach of practical reason. By taking the efficiency of the single-loop method as a reason, you can thereby create a voluntarist reason for yourself to tie in the single-loop method. You now have a voluntarist reason to tie in the single-loop method which, in conjunction with whatever given reasons you have for tying in that method, might give you most reason to tie in that

method. You might instead take the beautiful symmetry of the double-loop method as a reason to tie in the double-loop method, perhaps thereby giving yourself most reason to tie in that method. Or you might take and then "untake" these considerations as reasons if you change your mind about them.

The voluntarist reasons you create are arbitrary in the sense that there are no reasons for you to have created those rather than others. But the fact that you could have, so far as your reasons go, created a reason to tie in the double-loop method rather than in the single-loop method does not, it seems to me, make the reasons you create to tie in the single-loop method *objectionably* arbitrary. To be *objectionably* arbitrary the reasons you create should be ones that lead to substantively objectionable conclusions about what you have all-things-considered most or sufficient reasons to do. We intuitively think that you have most or at least sufficient reasons not to harm people for pecuniary gain. If you could, like the Mafioso, just will yourself a reason to shoot the kneecaps off your rivals that would justify your doing so, the reason you create would be objectionably arbitrary. But the reasons you create to tie your shoelaces one way rather than another don't seem like this. If voluntarist reasons were objectionably arbitrary *per se*, they should be objectionably arbitrary even when what is at stake is only how your shoe laces get tied.

I suggest that those who think that voluntarist reasons are objectionably arbitrary do not object to voluntarist reasons *per se* but rather to the objectionable conclusions about what we are all-things-considered justified in doing that voluntarist reasons might entail. This is indeed a problem for the standard forms of voluntarism that derive from Kant and Hobbes. If voluntarist reasons are the only sorts of practical reasons there are, then we can in principle create reasons that justify our doing things that we aren't justified in doing.

If voluntarism is understood in the hierarchical form suggested here, however, the voluntarist reasons we give ourselves will not be objectionably arbitrary. This is because the hierarchical relation between one's given and voluntarist reasons guarantees that we can never create a voluntarist reason that goes against our all-things-considered given reasons. You can't create a voluntarist reason unless your non-voluntarist reasons have run out. So, for example, the Mafioso can't create a reason to harm his rival since his given reasons require him not to do so. Thus he can't give himself a reason that justifies his doing so. We can create voluntarist reasons only in the space of rational freedom afforded by our given reasons. While our hierarchical voluntarist reasons are arbitrary in

[31] Cf. Scanlon (1998: chapter 7).

[32] It is also Jane Eyre's objection to Mr. Rochester's declaration that he has passed a law unto himself to improve his character. Jane says, "The human and fallible should not arrogate a power with which the divine and perfect alone can be safely intrusted." Rochester: "What power?" Jane: "That of saying of any strange, unsanctioned line of action, – 'Let it be right.'" Charlotte Bronte (1847: chapter 14).
   See Clarke (1706) and Cohen (1996): 167–188. The objection, to put it another way, is that standard voluntarist views like Korsgaard's really amount to what Korsgaard calls "heroic existentialism." See Korsgaard's epilogue (1997: 215–254).

[33] See Cohen (1996): 183–4. Kantians attempt to block this objection by putting constraints on rational willing so that all rational willing turns out to be moral willing, but such constraints are notoriously problematic.

the sense that we have no reasons to have one set rather than another, this arbitrariness does not lead to any substantive objection concerning what we are justified in doing. In this way, hierarchical voluntarism avoids what is widely considered the most fundamental problem with voluntarism.

8

There is another way in which hierarchical voluntarism steals a march on voluntarism in its standard forms. Here is a sketch of what I have in mind.

Three fundamentally different conceptions of the nature of normativity loosely underwrite the arguments among externalists, internalists, and voluntarists about the source of normativity. If normativity is an irreducibly distinct justificatory force, then its source is very plausibly in irreducibly normative facts. If instead it is a kind of motivating force, then where else to find its foundations but in desires and other motivational states of the agent such as the will? And if it is volitional force, then normativity plausibly has its source in the will or perhaps in other motivational states.

Standard forms of voluntarism depend for their plausibility on understanding normativity as something other than an irreducibly distinct justificatory force. For how can we, by a mere act of will, create such force? As some voluntarists have argued, normativity must be a volitional or motivational force since it is hard to see how an irreducibly justificatory force could "get a grip" on agents.[34] If normativity turns out to be an irreducible justificatory force, standard forms of voluntarism are in trouble.

Hierarchical voluntarism, by contrast, is not held hostage to the debate about normativity's nature. Suppose practical normativity is an irreducible justificatory force. While it is hard to believe that mere willing can create this force, it is not so hard – at any rate, less hard – to believe that willing *under the condition that one's all-things-considered given reasons have run out* can create this force. A rough analogy might help. Suppose you are given a blank piece of paper and some watercolors, and are told to create a beautiful forest scene. You will likely fail (if your artistic skills are like mine), painting lopsided trees with badly-proportioned features. If instead you are given the outline of a forest scene drawn by a master and are asked to color in the lines, you may well produce a beautiful forest scene. In this way, a constrained normative power may be more powerful than an unconstrained one. Coupled with the thought that one's willing

<hr>

[34] See, e.g., Korsgaard (1997). For a good reply to this objection, see Parfit (forthcoming).

under the requisite conditions constitutes one's rational identity, it becomes all the more plausible that the reasons that are part of one's rational identity could in principle have an irreducible justificatoriness about them.

If, on the other hand, normativity is a kind of motivational or volitional force, then the voluntarist reasons one creates will naturally bear motivational or volitional force. If the reasons that constitute your best rational self are made by you, they naturally bear motivating force. And since they are the product of your will, they naturally bear volitional force. In this way, hierarchical voluntarism can in principle straddle all three conceptions of the nature of practical normativity. If hierarchical voluntarism is right, then whatever the nature of normativity turns out to be, there is a place for the will in understanding its source.

My aim here has been to sketch the beginnings of what seems to me the most plausible view of voluntarist reasons by showing how such reasons can solve two puzzles about rational action and agency. If the arguments here are right, such voluntarist reasons – hierarchically related to non-voluntarist reasons – are able to capture the main insight of voluntarism – that we can confer normativity or value on things – without falling prey to its main difficulty. They put the will in its proper place.