

Persons, Agents, and Machines

[Extended Dance Mix]

In his recent book, *Legitimacy: The Right to Rule in a Wanton World*, Arthur Applbaum tries to show that a formal institution—in this case, a government—can be what he calls “a free moral agent.” While there are strict qualifications (not every state will function in this way, and most won’t), Applbaum’s argument tries to hold open the possibility that at least some formal institutions could function as moral persons in their own right, over and above the status of their constituent members. The full context for this argument would take me too far afield, but to briefly summarize: Applbaum argues that the only way for a state to legitimately rule free moral agents is for that state to itself be a free group moral agent.¹ Applbaum’s book therefore contains both a summary account of a theory of group agency and the outlines of a theory of free agency, followed by the attempt to show how we might attribute the latter to the former. For those of us who are interested in the issue of the moral agency of formal groups, Applbaum’s theory presents a fresh approach—and, furthermore I think that the refutation of Applbaum’s argument is a helpful new way of clarifying a thesis that I (and others) have defended elsewhere. I want to begin by unpacking a few ambiguities in Applbaum’s account of free agency; then I will go on to lay out the difficulties I have with Applbaum’s attempts to apply this account to group agents. Against Applbaum’s account, I will argue that a formal group agent must be a wanton. I will then conclude this paper by showing how I think we might yet productively recover some of the key insights of Applbaum’s account.

By “free moral agency,” Applbaum says that he means both *internal* and *external* freedom in sufficient amounts. External freedom requires that an agent’s “choice of ends” not be “subject to the control of another person,” and that their “innate powers or properly acquired means” not be “destroyed or unilaterally appropriated by another person’s choices.”² External freedom is in this way to be contrasted with various forms of domination. By “internal freedom,” Applbaum means autonomy, “competent self-governance.”³ This *competency*, for Applbaum consists of three essential capacities “or their functional equivalents.”⁴ These capacities he calls *considering*, *willing*, and *doing*, and they seem to be hierarchized such that having a capacity is also a condition for having the ones that follow it on the list. By “doing,” Applbaum means having one’s actions result from one’s *choice*, and by “willing” he means being able to *choose* in light of “relevant reasons.” In this way, willing in turn rests upon considering, which Applbaum consistently describes as a “capacity to respond to reasons for action, endorsing some and rejecting others.”⁵

This capacity for consideration, which is a necessary condition for the other two capacities on Applbaum's list, is somewhat ambiguous. On the one hand, Applbaum introduces these three capacities, not as the necessary conditions of *moral personhood*, but rather as the necessary conditions for *action*. Now, moral personhood implies a capacity for action—as Applbaum writes, “Anything that can be held to be properly [morally] responsible is . . . capable of action.”⁶ But this does *not* logically imply the inverse, that *anything capable of action can legitimately be held morally responsible*. As Harry Frankfurt warns us, the category of action is not restricted to humans: “We are far from being unique either in the purposiveness of our behavior or in its intentionality.”⁷ Frankfurt gives the example of a spider, for whom there is certainly a difference between *it acting*—say, by moving its leg—and *something happening to it*—say, when a boy ties a string to the spider's legs and moves them. If the capacity for consideration is interpreted in a broader sense—as simply the necessary condition for *action*—then the sense of what it means to “respond to reasons for action” will have to be dramatically different: It seems to me correct to say that a goldfish is *responding to reasons for action* when it sees food pellets hit the surface of the water in the fish tank. Likewise, an animal startled by a loud, sudden noise may freeze and look; when, having either spotted the source of the disturbance or simply failed to detect any further disturbances, it goes back to grazing or sleeping or eating, wouldn't we say that the animal has *rejected a potential reason for action*? By contrast, when the disturbance turns out to be a predator and the animal flees, isn't this the *endorsement* of a reason for action? On the other hand, however, it seems to me that Applbaum is referring to the reflective capacities that distinguish moral persons from lower animals, plants, and so on—and in this sense, he is right to treat his three capacities as both necessary and sufficient for moral personhood. It is for this reason that consideration involves not just reason-responsiveness, but the ability to (reflectively) *endorse* and *reject* reasons, and—as we will see in a moment—it is why Applbaum focuses on *wantonness* as a way of failing to be autonomous.

If consideration is a capacity for responding to reasons for action, then it is also important to specify what *sort* of reasons a free moral agent must have the capacity to respond to. There are, broadly speaking, at least two very different ways of interpreting the phrase “respond to reasons,” and they will present us with two very different conceptions of free agency. The reasons in question could be the agent's own; this is the conception of free agency at work in Frankfurt's account of freedom of the will. Here the idea is, in brief, that agents are capable of being moved by various desires—but freedom of the will demands that we be capable of identifying our wills with some desires (endorsing them) and not others (rejecting them), such that we are only moved to act by the desires that we choose to act upon. Inasmuch as this is a problem (or project) for creatures like us, we are *persons*—whereas an agent who

is incapable of either wanting or not wanting to be moved by certain desires is a *wanton*.⁸ Given that Applbaum specifically contrasts free moral agency with wantonism, it seems right to interpret his notion of “consideration” in this way—and I shall primarily do so below. However, the ambiguity here comes from the fact that a set of theories very different from Frankfurt’s about what it means to have “free will” has come to be known as “reasons-responsive theories.”⁹

For a reasons-responsive theory (key references here are Susan Wolf, John Martin Fischer, and Mark Ravizza), free moral agency is a matter of being responsive to what we might call *external* reasons—primarily rational and moral reasons.¹⁰ In stressing the ability of an agent to be *responsive* to reasons like this, a reasons-responsive theorist looks to situations of *brainwashing* and *manipulation* rather than *addiction* as the interesting cases. I think it is significant, for example, that the word “wanton” doesn’t appear anywhere in Fischer and Ravizza’s seminal book *Responsibility and Control*. As Michael McKenna and Chad Van Schoelandt point out, reasons-responsive theories are primarily ways of accounting for free will “in terms of how an agent is externally related to her environment.”¹¹ By contrast, so-called “mesh” theories of free moral agency (like Frankfurt’s) are attempting to account for freedom of the will “in terms of internal states of the agent”—hence an issue like wantonism is a key area of concern.¹²

As I see it, the distinction between these two approaches to free moral agency is important for our purposes because Applbaum seems to equivocate between them. He clearly says, near the beginning of *Legitimacy*, that a free moral group agent must have “a capacity for second-order reflection,” and he early on identifies *wantonism* as a failure of free moral agency. However, in his attempt to show how a group agent might overcome the threat of wantonism, Applbaum will focus most of his efforts on showing that a group agent can be responsive to rational and moral reasons for action (both in taking its actions and in offering reasons to justify those actions). Now, I think that any sufficiently complete reading of Fischer and Ravizza would show that a group agent *cannot* be a free moral agent on their account—that the very features by which they determine that “smart animals, children, and psychopaths” are not moral agents will also rule out formal institutions like governments and corporations.¹³ Nevertheless, that argument will have to wait for another paper. In the limited time that I have here, I will be making a more restricted argument: that a group agent *must be a wanton*. In arguing that Applbaum’s suggestions are unable to overcome wantonism, then, I will therefore simply *set aside* some of his comments on possible reasons-responsive mechanisms in group agents; they will be *irrelevant* for our purposes, whether or not I think that they ultimately hold up on their own merits.

If, as I have just charged, some of the resources Applbaum offers us for establishing free moral group agency simply sidestep the issue of wantonism, then it may be helpful to briefly reconstruct Frankfurt's concept of a wanton, at least as I understand it. Wantons, by Frankfurt's famous description, are "agents who have first-order desires but who are not persons because, whether or not they have desires of the second order, they have no second-order volitions."¹⁴ A first-order desire has as its direct object some object, sensation, activity, or state of affairs—as when I am hungry for Thai food or I want to read a book. A second-order desire has a desire as its direct object—as when I *want to* desire to exercise, or when I *don't want* to desire to smoke. Now notice, importantly, that Frankfurt says that wantons have no second-order volitions, *whether or not they have second-order desires*. And so the capacity for second-order desires will not be sufficient for showing that an agent is capable of free moral agency.¹⁵

What, then, is this *second-order volition* that the wanton does not possess? Frankfurt distinguishes between what *he* calls two different second-order desires:

Someone has a desire of the second order either when he wants simply to have a certain desire or when he wants a certain desire to be his will. In situations of this latter kind, I shall call his second-order desires "second order volitions" or "volitions of the second order."¹⁶

A second-order volition is the desire to have a desire be one's will—it is a desire not just to have a certain desire, but for that desire to be the cause of one's actions. Lacking a second-order volition, the wanton "does not care about his will," Frankfurt writes. "His desires move him to do certain things, without its being true of him either that he wants to be moved by those desires or that he prefers to be moved by other desires."¹⁷ This entails either an incapacity or a "mindless indifference to the enterprise of evaluating his own desires and motives."¹⁸ We should therefore distinguish wantonness from akrasia or "weakness of will"; for, as Frankfurt says, "when a *person* acts, the desire by which he is moved is either the will he wants or a will he wants to be without. When a *wanton* acts, it is neither."¹⁹

Personhood—that is, free moral agency—is thus a matter of not simply having an intentional will (the ability to choose in light of relevant reasons), but of also having a reflective and interested relationship with one's own will—such that one can have (or fail to have) *the will that one wants*. Moral personhood is in this sense a matter of having an ability to *identify oneself with* not just some of one's desires and even with one's second-order desires, but with one's "volitional complex"—so much so that J. David Vellman even calls second-order volition a matter of "a third-order desire, for there to be correspondence and causal influence between our second- and first-order desires."²⁰ Akrasia is

therefore a problem for persons to the extent that they are not (perhaps even *cannot be*) wantons. Frankfurt describes this experience:

Since the desire that prevails is one on which he [the akratic person] would prefer not to act, the outcome of the division within him is that he is unable to do what he really wants to do. His will is not under his own control. It is not the will he wants, but one that is imposed on him by a force with which he does not identify and which is in that sense external to him.²¹

One way of putting the argument that I am going to make is like this: It is impossible for a corporate agent (like a government or other formal institution) to *have an experience like* the one Frankfurt describes. It is impossible for a corporate agent to “identify with” its own volitional complex in such a way that it could *desire to have* the will that it does have, or in such a way that it could desire to have a will *other than* the one that it has.

In order to show why a group agent must be a wanton, let me now turn to the question of group agency. Intentional group agency is of course a widely discussed topic within the fields of business ethics and social ontology, and I believe that it is by now well established that some types of groups can have intentional agency over and above the intentions and actions of their constituent members. To be clear, however, this is not asserting much. Daniel Dennett and John Danley both use the example of a chess-playing computer. As Danley says, it is perfectly correct to say that “the computer *intends* to respond P-K4 to my king pawn opening.”²² And furthermore, Dennett writes, “I can calculate—under auspicious circumstances—the computer’s most likely next move, *provided I assume the computer deals rationally with these beliefs and desires,*” that is, its *beliefs* about the rules and location of the pieces and its *desire* to win.²³ Frankfurt, for his part, allows that “a rational creature, who reflects upon the suitability to his desires of one course of action or another, may nonetheless be a wanton.”²⁴ While the rational wanton may make significant use of instrumental reason, “what distinguishes the rational wanton from other rational agents is that he is not concerned with the desirability of his desires themselves.”²⁵ As Danley, Dennett, and Frankfurt all show us, then, rational intentional agency is a necessary condition of moral personhood—but not yet sufficient.

Insufficient as it is for moral personhood, rational intentionality is still a condition that not just any group will fulfill. For a *group* to qualify as a rational agent, it is not enough that it “display[s] a robust pattern of attitudinal and behavioral rationality,” as Philip Pettit puts it.²⁶ Pettit notes that groups constituted through meshed aims and plans will generally exhibit such rationality, thanks to the rationality of their individual members.²⁷ Such a group would behave as an intentional group agent, but

the rationality would not be the rationality *of* the group agent. For a group agent to qualify as a moral agent, its rationality must be autonomous. To put it in Pettit's terms: it needs to be *an autonomous agent*. And a group agent will "fail to be autonomous . . . if the attitudes ascribed to it—if you like, the group mind—[is] just a function of the corresponding attitudes adopted by the members, whether independently or under some scheme of coordination."²⁸ For the group agent to have its "own mind"—that is, its own rationality and intentional agency—it requires a constitution by procedure. And as Applbaum points out, not just any procedure will do:

Indeed, one tempting test of whether a procedure constitutes a shared agent is that the outcomes of the procedure meet some appealing standards of rationality even when the collective choice is at odds with the individual choices appealingly aggregated.²⁹

This issue—which has been discussed extensively by Philip Pettit and Christian List (both separately and in their joint work)—is the condition of "robust group rationality."³⁰ It requires that procedures for determining group attitudes and intentions be structured such that "by the given organizational design, consistent and complete group attitudes are guaranteed."³¹ Rather than merely tying the group agent's "will" to the vector sum of constituent member inclinations—as through a direct majoritarian rule on each issue as it arises—the organizing structure contains hard-coded logic constraints. When "the propositions on the agenda have non-trivial logical connections," these constraints ensure that the group will is logically self-consistent.³² By way of such organization, the group agent itself can be regarded as rational, rather than simply relying parasitically on the rationality of its constituent members. In List and Pettit's terminology, the mind of the group agent "supervenes" in this way on the beliefs, attitudes, and actions of its members—the actions and attitudes of the members are still the raw materials of the group agent's actions and attitudes, but are no longer causally determinate of them.

List and Pettit have—mostly for the simplicity of exegesis—focused on binary logical constraints (ensuring that, for example, if the group determines both P and $P \supset Q$, then the group agent will also determine Q , even if a majority or even all constituent members of the group believe $\sim Q$). Applbaum quite rightly expands the scope of the logical constraints, I think, arguing that an autonomously rational group agent would also have to entail a grammar condition (the text of any decisions or legislation issued by the group agent "must meet the current standards of grammar") and an arithmetic condition (decisions and legislation issued by the group agent "must comply with the axioms of arithmetic").³³ It is

certainly possible that we could come up with several more constraints of a similar nature—as Applbaum says, “Other basic requirements of rationality similarly follow.”³⁴

When rational constraints are hard-coded into the structure of the formal institution, then, the rationality of the group agent successfully supervenes on the rationality of its constituent members. The group agent can now be said to be a rational, intentional agent—“distinct from the agents who are its members,” which is (as Pettit has put it) the “salient, if not the only” way in which a group agent can be an autonomous agent.³⁵ But—as impressive as all of this organization is—the group agent has in this way merely been established as an “Intentional System,” taking its place among chess-playing computers, thermostats, and perhaps even apple trees.

Having shown that being an Intentional System is not sufficient for moral agency, Dennett introduces the idea of a “second-order intentional system” to begin to capture the missing conditions: “Let us define a *second-order intentional system* as one to which we ascribe not only simple beliefs, desires and other intentions, but beliefs, desires, and other intentions *about* beliefs, desires, and other intentions.”³⁶ Dennett plausibly suggests that one really important aspect of such second-order intentional systems is their capacity to *treat others* as intentional systems—the ability to *reciprocate* the ascription of intentionality. But it should be clear that self-reflection also relies on a second-order intentional capacity, the ability to have intentions, beliefs, and desires about *our own* intentions, beliefs, and desires. Frankfurt, meanwhile, suggests that reflexivity seems “much more fundamental and indispensable . . . than that of hierarchy” for explaining the difference between a wanton and a moral agent.³⁷ In other words, a mere hierarchy of desires is not the same as having second- or third-order desires. “The mere fact that one desire occupies a higher level than another in the hierarchy seems plainly insufficient to endow it with greater authority or with any constitutive legitimacy.”³⁸ The difference between a reflective (moral) agent and a wanton is that a reflective agent, by reflectively *identifying herself with* her “volitional complex,” becomes, in Frankfurt’s words, “a genuine participant” in the internal conflict between her various desires. She does not derive her agency from the outcome of this conflict; she rather applies her agency to it. This requires more than the algorithmic sorting of desires within a rational, intentional system; it requires a *sense of one’s own agency*.

Reflectiveness, the ability to distance ourselves from (and take stances towards) our own desires, beliefs, and intentions, is not merely necessary for moral agency—it is the condition that makes morality both a possibility and a necessity in the first place. “A lower animal’s attention is fixed on the world,” writes Christine Korsgaard:

Its perceptions are its beliefs and its desires are its will. . . . But we human animals turn our attention on to our perceptions and desires themselves, on to our own mental activities, and we are conscious *of* them. That is why we can think *about* them.

And this sets us a problem that no other animal has. It is the problem of the normative. For our capacity to turn our attention on to our own mental activities is also a capacity to distance ourselves from them, and to call them into question. . . . Now the impulse doesn't dominate me and now I have a problem. Shall I act? Is this desire really a *reason* to act?³⁹

The argument, in short, is that without reflective self-consciousness, not only is moral agency impossible—morality as such is unnecessary. And so moral *agency* requires reflexivity because *morality* requires reflexivity. (This, I take it, is the fundamental and important difference between wantonness and *akrasia*; weakness of will is only a problem and a possibility for a moral person, while wantonness is the incapacity for moral agency.)

This is why Applbaum is correct to focus on wantonness as the exemplary failure to attain moral agency; for it is at the heart of the issue. But what would it require for a group agent to overcome wantonness? What would it take for a group agent to be a *free moral agent*? A group agent would have to be shown to be capable of not merely hierarchizing its own desires and impulses, but of identifying itself with its own volitional complex. It would require the group agent to have a sense of its own agency, such that it could conceive of itself as a participant in the struggle between various competing impulses. And that would require the group agent to have an internal, reflective distance from its own impulses, beliefs, and desires, such that it would be possible for that group agent to have second-order desires, *wanting* to be motivated by some desires and not by others. Is there a procedure, structure, or device by which a group agent can accomplish any of this? I believe that there are good reasons to say that there is not.

Instead of reflectiveness, as I have already indicated, Applbaum focuses on reasons-responsiveness; but we can examine these mechanisms for signs of a group agency that is capable of second-order volition. And it is true that the kind of responsiveness that Applbaum demands does itself require an ability for reflection. The governance that Applbaum calls for requires the direct legislative participants “to be able to articulate the reasons they are being responsive to, both to each other and to their co-citizens In short, participants have a duty to deliberate in good faith, answer good faith objections, and explain to each other why they find arguments for one proposal stronger than other

arguments for another proposal.”⁴⁰ This certainly requires reflection and the ability to form second-order beliefs and desires—but it is a task of the constituent members of the group. To use the issue of group rationality as our guide, we might ask: Does Applbaum’s proposed scenario turn the group into an *autonomously reflective agent*? Does reflectiveness *supervene*?

Recall that, for a group agent to be autonomously *rational*, the rationality must be structurally wired into the procedures that constitute the group will. If the rationality of the group agent is merely parasitic on the rationality of the individual constituent members, then the group agent itself is not a distinct, rational entity. Likewise, then: In order for a group agent to be *reflective*, the reflective distance must not be merely parasitic on the reflective agency of the constituent members; “internal reflection” and a sense of its own agency must somehow happen at the procedural level of the group as an institution. While rationality can be hard-wired into group constitution in a way that makes it *supervene*, I believe that reflection cannot.

I have fully accepted the claim that a group agent (at least properly so-constituted) can be a rational, intentional agent, capable of holding beliefs, having desires, and acting on inclinations. But what, specifically, do we mean when we say this? Some comments by Philip Pettit are incredibly instructive:

If we are to recognize the integrated collectivity as an intentional subject, then we must admit of course that it is a subject of an unusual kind. It does not have its own faculties of perception and memory, for example, though it may be able to register and endorse facts perceived or remembered by others: in particular, by its own members. Under our characterization it is incapable of forming degrees of belief and desire in the ordinary fashion of animal subjects; its beliefs are recorded as on-off judgments, its desires as on-off intentions.⁴¹

What Pettit observes here is that, in order to know what the desires of a group agent are, we have to look to its active intentions. Other desires may be proposed—at the level of discussion between constituent group members—but these desires do not *become* the desires of the group agent itself *until they are ratified as intentions*. In the absence of such intentions, it is incorrect to say that the group agent itself actually has the desire; but this means, of course, that it is impossible to ever say that the group wants to have a different intention than it has, or that the group does not want to desire a thing that it desires. Because, as Pettit points out, these intentions are binary—on or off—then the best we can do is to see the group change its mind: yesterday it wanted X, but today it no longer does. Because the group’s beliefs only exist as active judgments, and its desires only exist at the level of active

intentions, it is incapable of reflectively separating itself from its beliefs and desires. It therefore cannot conceive of itself as a participant in the struggle between internal desires, for in a sense we can say that the group agent *is* its beliefs and desires—its existence as an agent is *determined by* them.

If a group agent isn't a wanton, then it's not an autonomously intentional agent at all. Either the intentional agency of the group is merely parasitic on the rational agency of its constituent members, or the group agent is a rational wanton. Whether a state, a corporation, or any other institutionalized group, our model for independent group agency should not be that of the moral person, but rather the chess-playing computer. We program institutions to behave in certain ways, and this includes hard-coding certain values into them—most clearly in the cases discussed by Applbaum, List, and Pettit, where we place certain rational, mathematical, and grammatical constraints on the decision-making procedures of the group. However, it should also be clear that formal organizations embody—I might even say *institutionalize*—other values as well, such that the very rational operation of such organizations acts to uphold and further these values. Organizational rationality is not just a matter of group metaphysics and consistency with reality, after all—it is at the heart of organizational efficiency. The rational autonomy of a formal organization does not endow it with moral autonomy, but instead helps to institutionalize certain pre-given values and make the organization more efficacious in pursuing them. In this sense, the “rationality” and “reasons-responsiveness” of an organization is not a complete goal, as John Ladd has pointed out:

In the case of organizations, “a decision is ‘organizationally’ rational if it is oriented to the organization’s goals.” Rationality is consequently neutral as to “what goals are to be attained.” Or to be more accurate, “rationality” is an incomplete term that requires reference to a goal before it is completely intelligible.⁴²

Consequently, I think, we might productively re-read Applbaum’s work on *legitimate* institutions as part of an effort to provide the missing normative context for rational organizational agency, with his “two-pronged test of legitimate government” as an account of some of the purposes that any legitimate institution must serve. Furthermore, if we reject Applbaum’s assertion that such formal organizations might themselves have (albeit even tentative) moral rights, then we can also incorporate two more suggestions from Ladd: First, dropping the idea that an organization might be a moral agent (and hence beholden to moral obligations), it follows that “the only way to make the rights and interests of individuals or of the people logically relevant to organizational decision-making is to convert them into” external constraints, which “would then be introduced into the rational decision-making as limiting

operating conditions.” Applbaum suggests that “a government is legitimate only when it can and does act to secure and protect a minimally adequate list of rights and freedoms on behalf of its free (enough) constituent individuals,” which is to say that “a political authority legitimately governs subjects only if the authority’s governance realizes and protects the freedom of those subjects over time.”⁴³ Rather than worrying about ensuring that the government is itself a “free moral agent,” we are better off ensuring that the governmental institutions we have are machines built to pursue and protect the right kinds of rights, in the right way.

Furthermore, Ladd writes, since a formal organization (like a government) cannot be a moral person, it has “no *moral* right to freedom or autonomy. . . . Hence, the other side of the coin is that it would be irrational for us, as moral persons, to feel any moral scruples about what we do to organizations.”⁴⁴ The question of whether or not to interfere with (or alter) the operation of a governmental formal organization is therefore not a moral question of loyalty or duty to a fellow moral agent; it is a practical question about the improvement of an organizational machine. To this extent, I think that we should tweak Applbaum’s account of legitimacy. An institution’s legitimacy is not to be thought of as the moral right to rule, but rather as the serving of the right kinds of purposes. I think that this reading of Applbaum can productively put to work much of the rest of his account, even if he himself would disagree with it. Most of *Legitimacy*, after all, is a *critical* account of the various ways in which a government can *fail* to be legitimate. And Applbaum suggests that “we should take questions of legitimacy to be primarily practical questions about what to do,” asked *in media res* by people who have no real option to live outside of political institutions. The question, for us, is not how to build a morally free group agent, but rather how to change, preserve, or defend specific, artificial arrangements in the form of formal organizations.

¹ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 77.

² Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 78.

³ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 78.

⁴ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 78.

⁵ *Ibid.*

⁶ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 77.

⁷ Harry Frankfurt, “The Problem of Action,” in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 78.

⁸ Harry Frankfurt, “Freedom of the Will and the Concept of a Person,” in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 19.

⁹ See, e.g., Michael McKenna, “Contemporary Compatibilism: Mesh Theories and Reasons-Responsive Theories,” in Robert Kant, ed., *The Oxford Handbook of Free Will*, 2nd. Ed. (Oxford University Press, 2012), pp. 175-198.

¹⁰ I recognize that the language of “internal/external” is imperfect, but I’m bending the stick a little bit here in order to emphasize the difference between a reasons-responsive account of free moral agency and a “mesh” account like Frankfurt’s.

-
- ¹¹ Michael McKenna and Chad Van Schoelandt, "Crossing a Mesh Theory with a Reasons-Responsive Theory: Unholy Spawn of an Impending Apocalypse or Love Child of a New Dawn?" in Andrei Buckareff, Carlos Moya, and Sergi Rosell, eds., *Agency, Freedom, and Moral Responsibility* (Palgrave Macmillan, 2015), p. 46.
- ¹² Michael McKenna and Chad Van Schoelandt, "Crossing a Mesh Theory with a Reasons-Responsive Theory: Unholy Spawn of an Impending Apocalypse or Love Child of a New Dawn?" in Andrei Buckareff, Carlos Moya, and Sergi Rosell, eds., *Agency, Freedom, and Moral Responsibility* (Palgrave Macmillan, 2015), p. 45.
- ¹³ John M. Fischer and M. Ravizza, S.J., *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge University Press, 1998), p. 76.
- ¹⁴ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 16.
- ¹⁵ In a footnote, Frankfurt writes, "Creatures with second-order desires but no second-order volitions differ significantly from brute animals, and, for some purposes, it would be desirable to regard them as persons." (Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 16n. 5; emphasis added.) The desirability of treating some kinds of non-persons as persons in some ways should be familiar to anyone who has occasionally reasoned with children. It is also perhaps a way to think about a theory running through authors like David Copp and Patricia Werhane, who suggest that we might usefully treat some kinds of corporate or group entities as "secondary moral persons." However, I leave that idea aside for another time.
- ¹⁶ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 16.
- ¹⁷ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 16.
- ¹⁸ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 19.
- ¹⁹ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 19.
- ²⁰ J. David Vellman, "The Way of the Wanton," in Catriona Mackenzie and Kim Atkins, eds., *Practical Identity and Narrative Agency* (Routledge, 2008), p. 171.
- ²¹ Harry Frankfurt, "Identification and Wholeheartedness," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 165.
- ²² John R. Danley, "Corporate Moral Agency: The Case for Anthropological Bigotry," in M. Bradie and M. Brand, eds., *Action and Responsibility* (Bowling Green State University, 1980), p. 144; emphasis added.
- ²³ Daniel Dennett, "Conditions of Personhood," in Michael F. Goodman, ed., *What Is a Person?* (Humana Press, 1988), p. 149.
- ²⁴ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 17.
- ²⁵ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 17.
- ²⁶ Philip Pettit, "Responsibility Incorporated," *Ethics* 117 (2007), p. 178.
- ²⁷ Philip Pettit, "Responsibility Incorporated," *Ethics* 117 (2007), p. 179.
- ²⁸ Philip Pettit, "Responsibility Incorporated," *Ethics* 117 (2007), p. 180.
- ²⁹ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 84.
- ³⁰ Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford University Press, 2011), p. 67.
- ³¹ Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford University Press, 2011), p. 69.
- ³² Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford University Press, 2011), p. 68.
- ³³ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 243.
- ³⁴ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 244.
- ³⁵ Philip Pettit, "Responsibility Incorporated," *Ethics* 117 (2007), p. 180.
- ³⁶ Daniel Dennett, "Conditions of Personhood," in Michael F. Goodman, ed., *What Is a Person?* (Humana Press, 1988), p. 151.

³⁷ Harry Frankfurt, "Identification and Wholeheartedness," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 165n. 7.

³⁸ Harry Frankfurt, "Identification and Wholeheartedness," in *The Importance of What We Care About* (Cambridge University Press, 1998), p. 166.

³⁹ Christine Korsgaard, *The Sources of Normativity*, ed. Onora O'Neill (Cambridge University Press, 1996), p. 93.

⁴⁰ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), p. 244.

⁴¹ Philip Pettit, "Groups with Minds of Their Own," in Frederick F. Schmitt, ed., *Socializing Metaphysics: The Nature of Social Reality* (Rowman & Littlefield, 2003), p. 182.

⁴² John Ladd, "Morality and the Ideal of Rationality in Formal Organizations," *The Monist* 54:4 (1970), p. 497; Ladd quotes here from Herbert Simon, *Administrative Behavior* (Free Press, 1965), pp. 77 and 14.

⁴³ Arthur Isak Applbaum, *Legitimacy: The Right to Rule in a Wanton World* (Harvard University Press, 2019), pp. 137, 141.

⁴⁴ John Ladd, "Morality and the Ideal of Rationality in Formal Organizations," *The Monist* 54:4 (1970), p. 508.