

Acting Without Me: Corporate Agency and the First Person Perspective

Herman Cappelen and Josh Dever

The work of Castaneda, Lewis and Perry convinced many philosophers that so-called *de se attitudes* (and the first-person perspective in particular) are the source of deep and wide-reaching philosophical insights. The *de se*, it is claimed, plays an essential role in giving an account of thought and agency. Others claim that the *de se* is essential to understanding morality, perception, or epistemic phenomena such as immunity to error through misidentification.¹

The *de se* deflationists,² on the other hand, have argued that the phenomena Lewis and Perry described tell us nothing of philosophical significance. According to the deflationist, Lewis and Perry used new and confusing terminology to describe old and familiar phenomena.

This paper is a new defence of the deflationary view. Our focus is on the connections between the *de se* and agency. We have two central goals:

1. To provide a new argument that *de se* attitudes play no essential (or explanatorily important) role in understanding actions. Our argument turns on the nature of corporate agency and is, in outline, very simple: corporations, such as Apple and BMW, act, yet we don't and shouldn't attribute to them *de se* states, so *de se* states aren't necessary to explain at least one common and important type of agency.
2. A second central goal is show that corporate and human agency are more similar than is typically assumed: corporate agency can serve as a faithful model of human agency.

We proceed as follows.

¹ The seminal texts are by Perry (1977, 1979), and Lewis (1979), (1986). They both acknowledge the work of Castaneda (e.g. 1966) as an important precursor. More recent fans of the *de se* include Bermúdez (2017), Prosser (2015), Ninan (2010), Torre (2010), while work that puts the *de se* to work in diverse areas of philosophy includes McGlynn (2016), L.A. Paul (2017), and Titelbaum (2017).

² For example Boer and Lycan (1980), Millikan (1990), Spencer (2007), Cappelen and Dever (2013), Magidor (2015). See Ninan (2016) for more references, some quotes, and a useful framing of the debate pro and con.

- In the first section, we give some representative examples of corporations and their actions, with the goals of (a) making plausible that corporate entities can undertake genuine action and (b) showing some ways in which those actions can float free of the intentions and actions of individuals composing and associated with the corporate entities.
- In the second section, we set out a number of considerations suggesting that corporations do not need any sort of proprietary indexical states in order to act. We show that several lines of thought that people have taken to show that action is impossible without *de se* thoughts are not only unconvincing in the corporate case, but turn out to have their grip in the individual case weakened by virtue of seeing why they fail in the corporate case. At the end of the second section, then, the minimal goal is that the reader will be convinced that corporate action does not require *de se* states.
- In the third section, we turn to considering responses that attempt to preserve an essentiality thesis by arguing that corporate action is in some important sense different from “normal human action”, and is best understood as a kind of pseudo-action. We argue that all attempts to differentiate corporate from human action are dialectically ineffective, as they turn out to beg exactly the question at issue in the discussion of essentiality theses.
- In the final part of the paper we address an apparent tension in our view: On the one hand, we need to show that corporations don’t have or require *de se* mental states. In order to do that, we make the case for an important difference between individual and corporate action. On the other hand, the more convincingly we argue in favor of the view that individual and corporate action are the same, the more we inadvertently make the case that corporations *do* have and require *de se* states. But, the objection goes, we can’t have it both ways: either there is a crucial and philosophically interesting distinction here, or there isn’t.

Before we get to all of that, a brief outline of the history of this debate. The tradition we’re responding to can be traced back to influential work by John Perry and David Lewis. Perry’s most famous example—the famous messy shopper from his (1979)—has been interpreted as showing that there’s a deep connection between indexicality and action. In that example, Perry is described as changing his messy behavior only after thinking the thought that would be expressed in words by ‘I’m making a mess’. Recognizing, in a first personal way, that *he* was the shopper making a mess, is, according to Perry, essential for the change in action. Perry says: “When we replace (“I”) with other designations of me, we no longer have an explanation of my behavior and so, it seems, no longer an attribution of the same belief. It seems to be an essential indexical.” (1979:3)

This line of thought has been extremely influential throughout philosophy. Here are two illustrations of how the assumption is now appealed to more or less as common ground in contemporary philosophy:

...practical guidance is, in Perry's phrase, essentially indexical, in the sense that its function depends not only on which of many propositions it expresses but also on how that proposition is determined by the context—specifically, on its being determined in the same way as the reference of indexical expressions such as “I”, “you”, “here”, and “now” (Velleman 2015: 78)

It is widely agreed that agents need information in an egocentric form: they must think of places as “here” and “there”, times as “now” and “then” if they are to be able to act on what they know (Owens 2011: 267)

Our goal in what follows is to sketch a view of human and corporate agency according to which they are fundamentally similar and not tied to the first person or other indexical attitudes.

1 Introducing Corporate Agency

In this first section, we explain what we mean by ‘corporate agency’ and ‘corporate action’ and describe some of its relevant features. By ‘corporate actions’ we don’t mean simply ‘collective’ action. Collective actions paradigmatically involve multiple agents deliberately coordinating on the performance of a single task, in such a way that there is a clear sense in which each individual agent is, via the collective action, involved with and intentionally undertaking that task. Think, for example, of multiple agents cooperating to lift a table. Collective actions, however, are not the case we will focus on. Rather, we will consider cases of what we will broadly call *corporate action*. One particular kind of corporate action is action undertaken by traditional corporations. Apple, for example, chooses to announce a new model of iPhone at a conference. This is an intentional action – it results from deliberation on Apple’s part, as a strategy for satisfying Apple’s desire to win more market share, a strategy selected in light of Apple’s beliefs about what announcement strategy will maximize sales.³ But Apple’s action differs in some important ways from the action of the agents collectively lifting a table:

³ There is of course room for skepticism both about whether Apple genuinely performs actions and about whether, if it does, those actions are underwritten by mental states like belief and desire. For some skepticism about the skepticism, see section 3 below, especially the initial pro tanto case for corporate action and the closing comments on corporations and the mental.

- When Apple announces a new model of iPhone, relatively few *members* of Apple may act. The announcement may be made solely by Tim Cook, while thousands of other employees⁴ go about their usual business.
- Even the deliberative process leading up to the announcement may have been undertaken only by a few members. So there is no helpful sense in which, when Apple announces a new iPhone, all members of Apple are involved with and intentionally undertaking the announcing.
- Indeed, it could be that *no one* does any announcing. One person might write part of an announcement script, someone else another part, a third person might edit it, yet a fourth person might forward the final version to the newspapers, and so on.

Broadly, then, there are many organizations and social structures—entities which presumably in some broad sense are grounded in features of individual people (more or less – see one case to the contrary below)—which can be thought of as independent locuses of deliberation and action.⁵ It is these sorts of broadly corporate agents, and the actions undertaken by them, on which we will focus.

For some of our subsequent purposes, it is worth attending to some features of the expansive boundaries of the category of corporate action. Note the following:

- A corporation can perform action A, as guided by an intention to A, resulting from a desire D and a belief B, even when *every single member* of the corporation is attempting not to do A (or even attempting to block A), as guided by an intention that not A resulting from a desire that not D and a belief that not B. Apple, for example, acts to cut prices on the iPhone, wanting to increase profits and thinking that lower prices will lead to higher sales and more profits. But it turns out that every member of the board actually wants the iPhone price to go up, because they want to sabotage Apple profits (perhaps to trigger a golden parachute clause in their contracts). But each board member thinks (incorrectly) that they are hated by the other board members, so that if they speak out in favor of and vote for the price cut, the other board members will vote in the opposite way. As a result, the vote comes out unanimously in favor of the price cut, and Apple so acts.

⁴ There's a technical sense of 'corporation' where Apple's employees don't count as members of the Apple corporation. That's not how we use the term.

⁵ We start with the paradigm case of traditional corporations, but the category that interests us is not limited to such cases. Traditional corporations have a clear member/nonmember distinction, but we do not insist on this feature in all cases. Traditional corporations come with a hierarchical structure that may include designated decision-makers and dedicated spokespersons authorized to speak on behalf of the corporation, but we do not insist on this feature.

The moral to draw from this is that the determination relations between the mental states and actions of corporate members and the mental states and actions of the corporation can be more or less arbitrarily complex and convoluted.

A corporation can perform an action even at times when the corporation has no members at all:

- Even if every member of Apple were tragically killed in an iPhone presentation disaster, the corporation could continue to act. Perhaps this is because there are authorized agents—not members of Apple—who can trigger further actions. (Lawyers hired by Apple can still bring it about that Apple defends its patent rights.) Or perhaps it is because there are mechanisms already put in place to trigger action in the right circumstances – stock purchasing algorithms, designed and authorized by Apple (although not necessarily by any member of Apple!) bring it about that Apple performs a hostile takeover of Yahoo when Yahoo’s stock drops below a crucial threshold.

The moral to draw from this is that *the determination basis for corporate action needs to be temporally extended*, depending on the actions and intentions of corporate members across a series of times before (and in suitable cases after) the time of the action, and agentially extended, depending on actions and intentions of agents who are not part of the corporation.

What action a corporation performs can be underdetermined even by the totality of facts about all its members’ mental states. Consider the following rather contrived example:

- Apple has two people on the payroll as “possible spokespeople”. Each day, a coin is flipped by a coin-flipping machine in a sealed vault. If the coin lands heads, person P1 is the actual spokesperson for that day, and statements made by that person (and not by P2) count as declarations by Apple. If the coin lands tails, P2 is instead the actual spokesperson for the day, and it is P2’s statements that count as declarations by Apple. On Wednesday, P1 makes a public statement that Apple is raising the price of the iPhone, and P2 makes a public statement that Apple is lowering the price of the iPhone. If the coin (unseen by anyone) has landed heads that morning, then Apple has raised the price of the iPhone, while if it landed tails, Apple has lowered the price.⁶ So Apple’s

⁶ We can separate the action of raising/lowering the price (an action making a normative change in what customers should be charged) from the action of charging customers differently. This second action might also be underdetermined by any fully mental supervenience base, if customer charges are automatically linked to the combination of coin flip outcome and P1 and P2’s statements.

action can be altered simply by changing a non-intentional fact: the orientation of the coin. It is not people's *beliefs* about the coin that determine what Apple has done, but only how the coin actually (again, unbeknownst to anyone) lies.

The brief moral: the determination basis for corporate action can extend beyond the (human) intentional. Of course, in this way corporate actions are actually like human actions, since human actions can be grounded, partially or fully, in non-intentional facts, given that brute physiological features may in some cases help determine what action a person performs.

2 There is no de se requirement on corporate action

Even if, when concentrating on ordinary human action, you find the thought that agency is impossible without mental states of some special de se kind compelling, you shouldn't find it compelling when you pay attention to corporate action. The burden of this section will be to examine a cluster of morals that derives from demonstrating that and considering why there is no plausible case for a de se requirement on corporate action. We show that there is simply nothing that would even count as an interesting de se kind of content for corporate entities in the first place.

2.1. Corporate Agency and The De Se

We are targeting the claim that no action is possible without the possession of some de se mental states, and those de se states are involved in the explanation of the action. We approach this question by asking: When Apple decides to announce the new iPhone, does it need to be thinking about itself in some special way? We answer this question by showing that specific theories of de se content, especially ones which are ambitious in making such content particularly unusual or distinctive, are going to look implausible when we move to the corporate case. Consider the following four theories of de se content:

1. Egocentric Space: Suppose, for example, that one has a roughly Evans-inspired thought that first-personal thought is thought that organizes the world around a spatial representational system that is egocentrically organized, so that to think of the world first-personally is to think

of things in terms of where they are in relation to *you*.⁷ Then it's going to look wildly implausible to think that Apple has such states. Apple doesn't organize its beliefs about the world by locating objects in "egocentric space". It isn't even clear that Apple has a location, and even if it does, there is no reason why Apple should systematically structure its representation of the world in terms of the relation of other objects to Apple. (Apple, for example, doesn't typically gain information about the world by a faculty of perception which delivers "centered information".⁸) So given that Apple does act, Apple better not need *de se* states of this sort as an essential prerequisite for action.

2. Locating the agent's body relative to external objects. Suppose you think that *de se* states play a special role in getting action locally connected to the body (i.e. that *de se* attitudes are involved in locating the objects we act on relative to ourselves). Here is an example from Perry:

Consider a transaction with a fax machine. To press certain buttons on it, I have to move my fingers a certain distance and direction from me. It isn't enough to know where the buttons were relative to one another, or where the fax machine was in the building or room. I had to know where these things were relative to me. (Perry 1998, 87)

If the role of the *de se* states (and the source of their essentiality) is to get the targeted external objects hooked up with the body of the acting agent in the right way, then it's not clear that it even makes sense to talk about corporations having states of this sort, much less that corporations *must* have such states in order to act. Apple doesn't have a body in any normal sense, so it has no need for a special category of mental states that track a body or help coordinate the body with the external world. Apple's lack of a body is in part a further consequence of the fact that Apple lacks a clear physical location. But it is also a consequence of the fact that Apple doesn't have dedicated and stable causal connections to specific parts of the world. Jones's hand is part of Jones's body, one might think, because Jones reliably has causal control over the state of the hand. The body/non-body distinction for Jones is drawn by the distinction between the things Jones can reliably control and the things he cannot. But

⁷ Cappelen and Dever (2013) note that there's a harmless deflationary version of this proposal, in which you think of objects in terms of their relation to you in the same, non-indexical, way that you might think of objects in terms of where they are in relation to the Eiffel Tower. We are targeting the non-deflationary version in the text.

⁸ How then does Apple gain information about the world? It's not our goal here to develop an account of corporate epistemic faculties. Perhaps Apple has faculties that are in some broad sense perceptual; perhaps obtaining contingent information about the world need not route through even broadly perceptual faculties. Whatever the final story is, though, there is no reason to think that it involves distinctively egocentric information.

Apple needn't have any such distinction – the things Apple can control might vary constantly, and there might be a spectrum of degrees of control.⁹

3. *The De Se as Implicit/Unarticulated Representation.* A third view has it that de se states are identified by their special mode of mental representation (see for example Recanati (2013)). Perhaps de se states characteristically don't include "explicit representation" of the self, but do include explicit representation of other objects. When Jones believes de se that he is in Athens, on this view, were we to look inside Jones's "belief box", we would see a mental representation of Athens (a token of a Mentalese term for Athens), but we would not see a mental representation of Jones himself. Again, it is unclear what it would even mean for corporations to have states of this sort. How could we distinguish between what corporations "explicitly represent" and what they merely implicitly represent? It's already unclear how to cash out the "belief box" metaphor with ordinary people, but here it's at least possible that an adequate computational neurobiology might reveal syntactic representation systems that allow us to show that there is no "self-representation" in certain beliefs. But it isn't even clear how to get started in looking for something "belief box"-like in the case of Apple.

4. *Epistemic characterization of the De Se:* Suppose you think that de se states are epistemically characterized. Perhaps de se states are states that originate in internally-directed epistemic capacities, such as proprioception and introspection.¹⁰ (To think of things in this way is to tie de se states to the epistemic phenomenon of immunity to error through misidentification – states are de se because there is no chance of epistemological error via misidentifying who the state is about, because the state comes about via a faculty that can only deliver information about the self.) Now, Apple *might* have such internally-directed epistemic capacities. (Although given the greater flexibility of corporate structure than human biological structure, any such capacities are less likely to be *essentially* internally-directed.) Apple might find out what its profits are in a very different way from the way it finds out what Google's profits are. But it's hard to see why any such difference should play an essential role in Apple's actions. Even if Apple did have to learn its own profits by reading the *Wall Street Journal*, it would still seem quite capable of acting on the resulting information.

⁹ Should Apple employees count as part of "Apple's body"? Perhaps some employees can be controlled by Apple, and others can't. Researchers might have research autonomy included in their contracts, so that their research activity can't be controlled by Apple, while salespeople might be contractually obligated to follow company policy. And employees might frequently and unpredictably move from category to category. Where there is control, it could be idiosyncratic in pattern. Perhaps the spokesperson can be told what policy to announce, but not what specific words to use in announcing it (except perhaps for the adjective "vibrant") or what clothes to wear when announcing. Should Apple's ad agency count as part of "Apple's body"? The ad agency isn't a corporate part of Apple, but it can in certain ways and under certain circumstances be directed by Apple.

¹⁰ "Internally directed" in that they represent things as internal, or that they normally represent the internal. This is not meant to deny that things other than the body can be represented in proprioception.

Obviously we can't exhaustively consider every possible substantive way of characterizing de se mental states, and argue that corporations can act without states of that sort (either because corporations can't even have states of that sort, or because corporate action can proceed without such states, even when corporations can have them). But we think that minimally a serious *challenge* has been offered to the fan of a de se essentiality view: give a substantive account of de se states that is compatible with (a) those states being required for action and (b) corporations being able to act. We think the considerable flexibility in the internal causal structure of corporations make it very hard to believe that that challenge can be met. Let us finish this section with two final points about the challenge.

First, an additional reason for thinking that the challenge is hard to meet is that we can imagine corporations which are set up so as to be simply *incapable* of self-beliefs of any sort. Perhaps Apple by explicit corporate policy never tracks information about Apple. Apple can know what raw materials are needed to produce iPhones, and can want iPhones to sell well. On the basis of these states, Apple can act to purchase materials at low prices. But in doing so, it does not *and cannot* represent itself as acting.

Second, it's sometimes tempting to meet this kind of challenge by deflating the de se: by treating it as evidence that the de se must be something "thinner" than we might pre-theoretically have suspected. We aren't trying to resist deflationary strategies here. One version of a deflationary strategy is a broadly functionalist approach. One cheap way to get an essentiality thesis is to simply *define de se* states to be those states which are efficacious in producing action. If that's all that's meant by de se states, we of course agree that corporations, like people, have de se states, and that such states are essential to action. But it's obvious that there is no philosophical interest in that deflationary conclusion. In endorsing it, we haven't learned anything substantive about the nature of content, the nature of mental representation, or the relation between mental content and action.

3. Corporations Can Act

We have been developing a very simple argument against de se essentiality views:

(Premise 1) Corporations do not have de se mental states.

(Premise 2) Corporations can act.

(Conclusion) De se mental states are not essential for action.

The burden of the previous section has been to defend the first premise (modulo the thought that it is Inflationary, rather than Deflationary, *de se* states that corporations do not have). In this section, we turn to defending the second premise.

A foreseeable reaction to the considerations we set out in the previous section is to say: “Of course corporations can’t have genuine *de se* mental states. That’s unsurprising, because corporations can’t undertake genuine actions.” The critic agrees, of course, that Apple does things. Apple releases the new iPhone, fires employees, and carries out hostile takeovers of competitors. But the critic makes a distinction between *mere doings* and *genuine actions*, and holds that everything Apple does falls into the first category. Our argument, then, depends on rejecting this distinction by the critic.

We will briefly make some comments in favor of the view that corporate doings are genuine actions. But we don’t intend these comments to do more than create a *pro tanto* case in favor of our view, to set the initial burden of proof. The bulk of this section will then be devoted to rejecting arguments against the view that corporate doings are genuine actions.

Why should we think that corporate doings are actions? Two brief considerations. First, our linguistic practice with respect to corporate doings looks indistinguishable from our linguistic practice with respect to ordinary actions by individuals. We describe individuals as asserting, planning, lying, attacking, being insincere, and defending – we equally describe corporations as asserting, planning, lying, attacking, being insincere, and defending. We find headlines such as “Apple acted as a bully in the ebook market” and “Apple acted unfairly in stolen iPhone incident”, which explicitly use the language of *acting* with corporations. Moreover, these intentional descriptions of both individuals and corporations are distinguishable from mere “as if” intentional loose talk (“My new iPhone hates me!”) by being part of a rich practice of using the intentional states in explanation, prediction, and evaluation. We describe corporations just as much as individuals as doing things on the basis of reasons and considerations, and as subject to rational evaluability and praise-and-blame for the things they do. Our ordinary practice, then, suggests a continuity between the individual and the corporate case.¹¹

The second point is that there had better not be a deep distinction between individual “genuine action” and corporate “mere doings”, because for all we could a priori tell, we might have

¹¹ Moreover, these “act” and “action” uses are clearly distinguishable from non-agential uses such as “The refrigerator acted strangely last night”. Actions of both agents and corporations can be modified by various intentional adverbs (“John/Apple acted capriciously/unfairly/compassionately in firing Bob”) while actions of artefacts and other objects cannot (“The refrigerator/cold front acted capriciously/unfairly/compassionately in lowering the temperature”). Actions of both agents and corporations are performed for reasons, while actions of refrigerators and cold fronts are not. (Thanks to [redacted] for discussion on this point.)

turned out to be corporations. Of course we know this to be false, but that's an empirical discovery. We could have discovered that each of us is inhabited by a host of corporate employees, and that we each act by way of an elaborate corporate authority structure. We move our hands not by initiating a sequence of neural firings leading to muscle contractions, but rather by issuing an order to the "hand division" of Me Incorporated,¹² leading to various employees running around inside us pulling various cords and tendons. Now, one *might* take this as reason to think that it's an empirical discovery that we genuinely act, but we think it's greatly preferable to take "genuine action" just to be the kind of thing we do, and thus to think that such action must be compatible with corporate-style realization.¹³

So much for the *pro tanto* positive case. What reason, then, is there to doubt that corporate doings are genuine actions? The general form of the discussion for the remainder of this section will be to consider specific markers of genuine action that are putatively missing in the corporate case. We will consider three types of markers:

- *Epistemological markers*: The thought here is that genuine actions are distinguished by a special epistemological relation to what is done, as in the Anscombian view that in genuine action, the acting agent has non-observational knowledge of what they are doing (and why they are doing it).
- *Structural markers*: Alternatively, some think that genuine actions are structured by the agent's exercise of "direct guidance" or "control" over what is done, and that actions are built up out of a vocabulary of "basic actions".
- *Phenomenological markers*: A third view is that a genuine action is characterized by its seeming a distinctive way to the acting agent.

We will along the way express some skepticism about whether these various markers are reasonable constraints to place on a category of genuine action, but our primary strategy will be to hold that there is no non-question-begging case to be made that corporate doings lack any of these features.

¹² The 'hand division' could consist of little finger gnomes that move the fingers around – as suggested to us by [redacted]. Note the difference between this scenario and Block (1978)'s Homunculi-Headed Robots. We allow the possibility that you, as "CEO" have complex mental states, including intentions to act. It's only the implementation in action of these intentions that are then handed off to the corporate underlings. There's no aspiration here to give a generic functional recreation of a person in terms of homunculi with a "low level of intelligence".

¹³ That is, we take most philosophical opposition to genuine corporate action to be driven by a priori considerations about the nature of action, not by an empirical discovery, through investigation of our own natures, that action must be realized thusly. It can't be simultaneously (a) epistemically possible that we are corporations, (b) epistemically necessary that we act, and (c) epistemically necessary that corporations don't act.

Epistemological Markers: It is a philosophical commonplace that action carries with it certain epistemic privileges. George Wilson and Samuel Shpall's Stanford Encyclopedia of Philosophy article on action observes that "it is frequently noted that the agent has some sort of *immediate awareness* of his physical activity and of the goals that the activity is aimed at realizing." As a more specific spelling-out of the epistemic privilege, Wilson and Shpall note that:

David Velleman [1989] describes knowledge of one's present and incipient actions as 'spontaneous' (knowledge that the agent has achieved without deriving it from evidence adequate to warrant it), and as 'self-fulfilling' (expectations of acting that tend to produce actions of the kind expected). (Wilson and Shpall 2012: §1.1)

Perhaps, then, corporate doings are *mere* doings because they lack this epistemic privilege.

But is there any good reason to think so? Suppose Apple releases a new iPhone. Does Apple have "immediate awareness" of its having done so? Is that knowledge "spontaneous" in Velleman's sense, and does Apple's expectation of releasing a new iPhone "tend to produce" releases of iPhones by Apple? At least in normal cases, affirmative answers to all of these questions seem reasonable. In normal cases, the mere fact that Apple *has* released a new iPhone is taken as sufficient reason to say that Apple knows that it has released a new iPhone. Much as in the individual case, there is something peculiar about the question of whether Apple *knows* that it has done so (except in certain specifically deviant cases we'll discuss soon.) We normally take that knowledge to be immediate, spontaneous, and non-observational in much the same ways we do in the individual case. Apple might know that Samsung is releasing a new phone by attending Samsung's press conference, or by receiving reports from corporate spies in Samsung factories – but this isn't (typically) how Apple knows that Apple is releasing a new phone. Of course, individual employees at Apple will know that Apple is releasing a new iPhone in a wide variety of ways. But *Apple's* way of knowing need not be the same as the way of knowing used by the Apple employees, even in cases in which Apple's knowing is in some important sense constituted by the knowing of individual Apple employees. And, to stress again points made above, we can easily give cases in which Apple knows that it has released a new iPhone without, or constitutively independently of, anyone at Apple knowing this. In such cases, we may well feel that Apple's knowledge of the release is simply constituted by the release.

The skeptic about corporate action may, of course, feel that these are not genuine states of knowledge on Apple's part. On the skeptic's view, corporations engage in mere mock action, distinguished from real action by the lack of (for example) non-observational knowledge by the acting party of what is done. The skeptic then responds to these considerations by holding that

we have only provided more mockery – we have backed up *mock action* with *mock knowledge*. Well, perhaps. But we're far from convinced. Now we need a reason to think that the corporate epistemic states are mock knowings, rather than real knowings.¹⁴ Since the whole game by the skeptic was to provide a positive reason to distinguish corporations from individuals, we don't think much progress has been made.

Perhaps, though, convincing cases can be described in which Apple releases a new iPhone without thereby suitably knowing that it has done so. The iPhone is ready and in the factories, but no corporate decision to release has been made. But a rogue spokesperson at an industry event announces the new phone. This constitutes a release (because the spokesperson has suitable authority within Apple), but Apple doesn't know that the release has happened, because it didn't know that the spokesperson would make this announcement. It's tricky deciding what to make of such cases, but here are three thoughts:

- (i) Perhaps the spokesperson knowing what he has announced is a way of Apple's knowing what he has announced.
- (ii) Even when construed most favorably to the skeptic, such cases would only show that *some* things we call corporate actions aren't genuine actions. (Still, that would serve to undermine some of our initial *pro tanto* case.)
- (iii) Finally, and most importantly, we think that these cases just cast doubt on the plausibility of epistemic markers for action – because we think that similar cases can be given in the case of individual action. Individuals can (in, for example, a state of distraction) do things (e.g. push a button) that ought to count as real actions, but about which they lack non-observational knowledge, because the relevant executive epistemic systems haven't been monitoring what the individual is doing.

Structural Markers: Consider another characterization of the acting/doing distinction given by Wilson and Shpall:

It is also important to the concept of 'goal directed action' that agents normally implement a kind of *direct* control or guidance over their own behavior. An agent may guide her paralyzed left arm along a certain path by using her active right arm to shove it through the relevant trajectory. The moving of her right arm, activated as it is by the normal exercise of her system of motor control, is a genuine action, but the movement of her left arm is not. That movement is merely the causal upshot of her guiding action,

¹⁴ For an argument that corporations have genuine intentional states, see Tollefsen 2002.

just as the onset of illumination in the light bulb is the mere effect of her action when she turned on the light. (Wilson and Shpall 2012: §1.2)

Perhaps corporate doings don't allow a suitable distinction between what is directly controlled and guided and what is merely a causal upshot of the guidance, and hence doesn't qualify as action.

But why would we think that? On the face of it, a suitable distinction is easily available. In the above example, the difference between the way the right arm moves and the way the left arm moves is spelled out in part in terms of *normal* systems. But corporations, too, have both normal and abnormal ways of doing. Apple has a normal way to move iPhones from China to the United States, in the same way that we have a normal way to raise our left arm. In Apple's case, the normal way is to put the iPhones in a Boeing 747 and fly them across the Pacific. But Apple could move the iPhones in an abnormal way, in the same way that we can raise our left arm in an abnormal way. If the Boeing 747 was unable to fly (was 'paralyzed'), Apple could use other planes to lift that plane and carry it, with its iPhones, across the ocean.

Maybe "normal systems" talk isn't the right way of cashing out the direct control/guidance structural thought. Again, we can't survey every possible way of running the skeptical strategy, so we primarily hope here to raise a challenge to the skeptic – they need to give a convincing way of distinguishing individual genuine action from corporate mere doings. We'll consider two options:

(i) Appeal to Basic Actions construed representationally: Perhaps a better way of capturing "guidance" is via a notion of a *basic action*, where we can initially think of these as *the point at which representation of the means of carrying out the action gives out*. The thought is roughly that there are certain things that we are just *directly* able to do (perhaps typically things like *moving our hands*), and then other more distally-characterized things (like *answering the phone*) count as actions just in case those doings are brought about in the right way by sequences of *basic actions*, in the sense introduced by Danto (1973).

The challenge, then, would be to give a suitable sense of basic actions for corporations. But why would this be difficult? Of course, when Apple does something, like releasing a new iPhone, it is typically *exploiting* underlying action capacities of other individuals (sometimes employees of Apples, sometimes not). The releasing may get done by various people speaking and moving in various ways. But that does not show that Apple's action is not basic *qua* action by Apple, any more than your moving your hand is not basic because you exploit various nerves and muscles doing things in performing your basic action. If basic actions mark the point at which

representation of the means of carrying out the action gives out – if, as Lavin (2013, 4) says, “With basic action, thought about how to realize the end in question gives out; one acts immediately, directly or ‘just like that.’”, then corporations seem rife with basic actions.

If there is a worry here, it is perhaps that corporations end up with *too many* basic actions. Basic actions for us, on standard ways of thinking about it, are limited to movements of bits of the body. But basic actions for Apple might include just about anything – Apple might *basically* design a new Macbook, or open a plant in Singapore, or overtake Google in the search engine market. (Of course, Apple could have structured schemes for the doings of any of these, but it also looks possible that at the level of Apple’s acting, that’s all there is to the action, and any further scheming is devolved to (e.g.) Apple employees, in the same way that muscular triggering sequences in hand movements are devolved to the non-intentional anatomy of an acting individual.)

(ii) Appeal to basic actions construed as the domain of autonomous actions: Basic actions are called on to play a dual role. They are *representationally* basic, in that the agent need not further intentionally articulate how to carry things out. But they are also often taken to represent a domain of *autonomy*. When we seek, for example, to answer the phone, we may succeed or fail, depending on how things go in the world around us. But when we seek to perform a basic action, we have a kind of guarantee of success. And perhaps the thought is that that kind of guarantee can’t be spread as widely as the apparent range of corporate basic actions seeks to spread it.

If this is the thought, we think it encapsulates a familiar kind of mistake. To look for a domain of autonomy in this sense is one more instance of trying to find a “cognitive home” in which the possibility of error and failure is eliminated. But there is no cognitive home – failure is always a possibility. Corporations may fail in the exercise of their putative “basic actions”, but that had better not be an argument against their basicness, because we too can fail in the exercise of our basic actions.

Phenomenological Markers: Perhaps actions are distinguished from mere doings by being accompanied by a characteristic phenomenal feel to the acting individual. What might the phenomenal feel be? Elisabeth Pacherie in “The Phenomenology of Action” gives:

A non-exhaustive list of proposed distinctions includes awareness of a goal, awareness of an intention to act, awareness of initiation of action, awareness of movements, sense of activity, sense of mental effort, sense of physical effort, sense of control, experience

of authorship, experience of intentionality, experience of purposiveness, experience of freedom, and experience of mental causation. (2008: 180)

We might helpfully appeal here to phenomena like alien limb syndrome, in which there is a loss of felt ownership of a part of the body, undermining our ability to act with (rather than merely *do* with) that body part.

By now the structure of the dialectic should be clear. The skeptic proposes that corporate doings are *mere* doings because they are not accompanied by the characteristic phenomenology of action. We then ask why we are supposed to think that. When Apple releases a new iPhone, is it not aware of its intention to do so? Does it not have an experience of authority, intentionality, and purposiveness? Does it not, even, have an experience of mental (corporate?) causation?

If the skeptic has any leverage here, it presumably comes from a general skepticism about whether corporations have phenomenal states of any sort. We admit we're not entirely sure ourselves what phenomenal states are and whether we ourselves have them (have them in general, and even more so have them in the specific forms listed above as characteristic of the experience of action). So one response to which we are friendly is a general deflationism about the phenomenal constraints on action.¹⁵

That said, it also worth noting that phenomenal vocabulary is freely used in discussing corporations. We are told that Apple "feels the need to defend Steve Jobs", "feels the pressure with the new iPhone launch", "fears Android", and "hates open-source software". The skeptic, of course, will suggest that this is mere *mock* phenomenology, backing up *mock* action. We then again suggest that the question is being begged. We are told that corporations are importantly different from individuals, but every attempt to cash out that important difference ends up resting on an undefended putative prior difference.¹⁶

¹⁵ This line, in particular as applying to the corporate case, is supported both by philosophical argument and empirical data. Tollefsen (2002) assumes that corporations don't have phenomenal states, as does List (2018). Interestingly, so do the folk: Knobe and Prinz (2008) reports the results of a study showing people are less likely to attribute phenomenal states to corporations than they are to attribute attentional states to them.

¹⁶ Again, this route is also supported by some literature. Schwitzgebel (2015) argues on the basis of cases with some similarities to our corporate examples (the Sirian supersquids and Antarean antheads) that we should take seriously the possibility that full-fledged phenomenal consciousness is much more common than we ordinarily take it to be. (For example, that it might be a feature of nation-states). While we aren't opposed to pushing these considerations that far, for current purposes all we need is that corporations are *phenomenal enough* to meet any phenomenal requirements that there are on action.

More generally: we have been happy to assume throughout that corporations have genuine mental states (beliefs, desires, intentions). However, that assumption is more than is needed for our central argument. Our *de se* deflationism requires only that corporations *act*. If they can act without having mental states, they thereby provide an even more direct counterexample to the claim that action requires *de se* representation. The unconvinced skeptic, of course, may want to link action and mental states, deny that corporations have mental states, thereby deny that corporations act, and then retain an inflationary *de se* requirement on action. But such a skeptic inherits an explanatory burden on their account of the *de se*. Take whatever feature F corporations lack that prevents their states from being mental states. (For example, F might be phenomenal character.) Corporations can, even according to the skeptic, *do things* without feature F, so the presence of feature F is crucial for converting mere doings into actions. Since it's actions that impose a requirement for *de se* representation, feature F had better also explain the appearance on the scene of those *de se* representations. It's thus incumbent on this skeptic to say enough about what *de se* representation is to explain why feature F contributes to giving rise to it. This is a non-trivial requirement – many existing accounts of the *de se*, for example, fail to link *de se* representations distinctively to phenomenal character.

3.1.Objection: Corporations only act because they consist of people with *de se* states

A natural response to the line of thought above is the following:

Suppose that corporations do act without *de se* states. This doesn't show that *de se* states are inessential for action. A corporation is composed of people, and the corporation acts because the people composing the corporation act. But for the familiar Lewis-Perry reasons, the individual people can't act without having *de se* states. Thus *de se* states are essential for the action of the corporation—just not *de se* states *of the corporation*.

Note first that this objection begins by conceding our case against the claim that agency requires *de se* states. The objection grants that corporations act without *de se* mental states, so it isn't true that "agents need information in egocentric form ... to be able to act on what they know."

The objector wants to propose a new essentiality thesis to replace the defeated thesis. The exact form of the new essentiality thesis isn't completely clear; perhaps it's something like "agents need something in the full causal explanation of their actions to involve someone's having *de se* mental states". Whether such a thesis, if true, would give the *de se* a philosophically interesting position in our understanding of agency and personhood is far from

clear. But our immediate concern is whether there is any reason to believe the revised essentiality thesis. We outline two reasons for skepticism:

First Reply

If corporations can act without *de se* mental states by having their actions determined by the actions of individual members of the corporations, presumably we too can act without *de se* mental states in the same way. As noted above, we could have been corporation-like: our actions could be the result of the actions of individual agents, call them gnomes, who are members of our corpora. It would work like this:

Humans acting with the assistance of internal gnomes: Suppose A wants button B to be pushed. Inside A there are finger gnomes and arm gnomes. They implement A's desire by moving the arm and finger in an appropriate way so that button B is pushed. So A pushed the button with the help of the arm and finger gnomes (similar to how Apple acts with the help of employees in various departments).

It's hard to resist the idea that this is a genuine action by A: A often whips cream using an electric whipper. The gnomes are just internal versions of the electric whipper. Of course, no humans we know actually have internal gnomes, but it's not unlikely that technological advances in the relatively near future can make available electronic versions of such gnomes.

What we in the first instance take from this example is that if corporations can act without having *de se* states, so can we: we could be organized much as corporations (with the internal gnomes as the analogue of employees). The revised essentiality thesis offers no objection to that conclusion—it merely observes that our ability to act without having *de se* states ultimately rests on the *de se* states of the gnomes.

Now ask: why should we take seriously the thought that the gnomes (that by hypothesis make our actions possible) need some special *de se* mental state in order to act? The original motivation for heading down the *de se* path was specific cases that looked like they called for specific kinds of action explanation. When we are told about the messy shopper's sudden change in behavior, our immediate reaction is to think that the change in behavior calls for (i) a specific kind of belief (ii) in a specific person: Perry. What we're supposed to learn from the example is that the shopper himself comes to believe that he himself is making a mess. That compelling thought then leads to the philosophical puzzle about what the nature of that specific belief could be. But if Perry had ended the messy shopper case by announcing that clearly *someone else* (e.g., one of the gnomes inhabiting him) had a new belief about where he was, or about how many fingers he had, we would have found the announced conclusion

completely baffling. The messy shopper case just doesn't provide any reason to think that *those* sorts of beliefs are required.

So once the fan of the *de se* has conceded that the *de se* states might be in our gnomes rather than in us, and might be about features of the gnomes' environment quite unrelated to the actions we are performing, the original philosophical impetus to propose a new kind of mental state goes away. Absent any positive reason to incorporate *de se* mental states in the gnomes, we'll then propose that the gnomes in fact act without the *de se*. They could, for example, be little computers. Now we have a counterexample to the revised essentiality thesis as well: when we act, there aren't any *de se* mental states *anywhere*. And once we reach that point, we can drop the appeal to the gnomes, and just think that the story about action begins and ends with us, and doesn't involve the *de se* anywhere.

Second Reply

A second point in response to the revised essentiality thesis: we were supposed to think that corporate action involved the *de se somewhere* because the corporate actions depended on the actions of the people making up the corporation. But perhaps we can have corporations not made up of people. We've already noted that corporations can act when they have no members at all, if the corporation is acting on instructions or via algorithms set in place when there were members. And we could imagine a corporation whose members were entirely robots or computer programs. The thought that corporate actions need to be underwritten by the *actions* of corporate members would then be met with the suggestion that corporate actions could be underwritten by the *doings* of the corporate members, whether or not those doings rose to some standard of counting as actions.

Would such non-human corporations threaten our confidence that the corporations were genuinely *acting*? Consider again us and our hypothetical gnomes. If we've already been convinced that we could be acting as we do as a consequence of the actions of gnomes who make us up, would it really be a threat to our status as actors to suggest further that the component gnomes are more like robots than people; that the gnomes do rather than act? That change in the gnomes won't entail any change in our ability to deliberate and subsequently intend on the basis of rationally weighed reasons -- why would it entail a change in our ability to act? We might then suggest that the gnomes have a familiar name: neurons.

4. An Objection

We have now finished the core of our argument: we have argued that corporate action is genuine action, but that in no interesting sense are *de se* mental states required for

corporations to act. Hence, *de se* states cannot in general be a requirement for action, so any strong *de se* essentiality principle must be false. In this final section, we consider an attempt at a transcendental refutation of our argument. The central thought is simple. Our argument, holds the transcendental refuter, requires two contradictory things.

(i) On the one hand, we need corporate action to be real, and not merely mock, action. We thus need there to be no trait that the skeptic can use to argue for a crucial difference between corporate and individual action.

(ii) On the other hand, we need corporate action to be distinguished by the absence of distinctive and interesting *de se* mental states (or by the inessentiality of such states, if they do occur). The goal is then to use that difference to convince the *de se* theorist that *de se* states aren't *essential* to action – they might be a local feature of the way action systems are implemented in beings like us, but consideration of corporations shows that action could be implemented in other non-*de-se*--requiring ways.

But (i) and (ii) pull in different directions:

- On the one hand, to establish (ii) we need to show that corporations don't have or require *de se* mental states. In order to do that, we inadvertently make the case for an important difference between individual and corporate action (and so undermine (i))
- On the other hand, the more convincingly we argue in favor of (i), i.e. that individual and corporate action are the same, the more we inadvertently make the case that corporations *do* have and require *de se* states.

We can't have it both ways – either there is a crucial and philosophically interesting distinction here, or there isn't.

The situation here is one which confronts many broadly deflationary strategies in philosophy. The classic instance is Berkeleyian idealism. Berkeley wants to argue for two things:

- First, the nature of reality differs sharply from what we ordinary believe – the existence of objects is, so to speak, much more metaphysically “lightweight” than we take it to be.
- Second, the “lightweight” nature of reality is deeply indistinguishable from the “heavyweight” reality we took there, in our pre-Berkeleyian naivete, to be.

The problem that Berkeley confronts is then that the success of the second point threatens the plausibility (and even the sensibility) of the first point. If objects being metaphysically

lightweight is a perfect simulation of objects being metaphysically heavyweight, then it's unclear what the *actual difference* is between things being light or heavy. Thus the threat that by "external objects", all we ever meant was "idea in the mind of God", so that we all agreed with Berkeley the whole time, and he made no progress toward convincing us that reality was not heavyweight.

Similarly in many other areas. The ethical noncognitivist might try to leave everything in our moral talk and practice unchanged, showing that he can make noncognitivist sense of ascriptions of truth, belief, and facts to moral matters. But if his success is complete, we are at risk of losing track of the difference between cognitivism and noncognitivism, and the deflation that the noncognitivist wants becomes indistinguishable from inflation. And if the success is incomplete, then the cognitivist is left with a mismatch between noncognitivist theory and ordinary practice that can be leveraged into an argument for cognitivism. The truth relativist might try to leave everything in our ordinary talk unchanged, by showing how a monadic truth predicate can be constructed out of relativistically-acceptable resources. But if his success is complete, we are at risk of being able to say things like "It is true simpliciter, and not merely true relative to my standards of taste, that the joke is funny", and thus of losing track of the difference between the relativist and the absolutist. Deflationary perfect simulation of inflationary phenomena looks an awful lot like genuine inflation.

Thus the challenge for us is this: If corporate action is a perfect simulation of genuine action, then in what sense have we *deflated* the conditions for real action, rather than *inflated* the understanding of corporate action?

We are not convinced that we are genuinely subject to a perfect simulation challenge. We deny one prong of that challenge, because we deny that the more convincingly we argue that individual and corporate action are the same, the more we inadvertently make the case that corporations *do* have and require de se states. We think individual actions don't require de se states, so we don't think bringing out similarities between corporate and individual actions will create pressure to treat corporate actions as de se involving. Furthermore, we don't think that response is question-begging in context. It's not that we need already to assume de se deflationism in order to avoid the perfect simulation argument. Rather, we think that there's just a real similarity between corporate and individual actions which can be spotted by people independent of their commitments on de se matters, and then reflection on that perceived similarity can be leveraged into evidence against the essentiality of the de se.

Perhaps we are wrong about that. If so, we do indeed face a perfect simulation challenge to our ability to draw a meaningful distinction between the deflationary and the inflationary account

of action. We can then at least take comfort in knowing that our company is good. As we've noted, a broad range of deflationary projects face challenges of this sort. We think there are promising routes for responding to perfect simulation arguments where they do have bite, but this is not the place to pursuing a project of that size. For now we rest content with the possibility that the de se deflationist does not outright win the day, but merely makes the case that there is no substantive difference between the presence and absence of de se mental states.

Bibliography

Bermúdez, José Luis (2017). Yes, essential indexicals really are essential. *Analysis* 77 (4):690-694.

Block, Ned (1978). Troubles With Functionalism. *Midwest Studies in the Philosophy of Science* 9:261-325.

Boer, Steven E. & Lycan, William G. (1980). Who, Me? *Philosophical Review* 89 (3):427 - 466.

Cappelen, Herman & Dever, Josh (2013). *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press.

Castañeda, Héctor-Neri (1966). 'He': A Study in the Logic of Self-Consciousness. *Ratio* 8 : 130–157

Danto, Arthur C. (1973). *Analytical Philosophy of Action*. Cambridge University Press.

Knobe, Joshua & Prinz, Jesse J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and the Cognitive Sciences* 7 (1):67-83.

Lavin, Douglas (2013). Must There Be Basic Action? *Noûs* 47 (2):273-301.

Lewis, David K. (1979), "Attitudes De Dicto and De Se," *Philosophical Review* 88 : 513–543 .

— (1986), *On the Plurality of Worlds* . Oxford : Blackwell .

List, Christian (2016). What is it Like to be a Group Agent? *Noûs*:295-319.

- Magidor, Ofra (2015). The Myth of the De Se. *Philosophical Perspectives* 29 (1):249-283.
- McGlynn, Aidan (2016). Immunity to Error Through Misidentification and the Epistemology of De Se Thought. In Manuel Garcia-Carpintero & Stephan Torre (eds.), *About Oneself: De Se Thought and Communication*. Oxford University Press. pp. 25-55.
- Millikan, Ruth Garrett (1990). The Myth of the Essential Indexical. *Noûs* 24 : 723–734 .
- Ninan, Dilip (2010). De Se Attitudes: Ascription and Communication. *Philosophy Compass* 5 : 551–567
- Ninan, Dilip (2016). What is the Problem of De Se Attitudes? In Manuel Garcia-Carpintero & Stephan Torre (eds.), *About Oneself: De Se Thought and Communication*. Oxford University Press.
- Owens, David (2011). Deliberation and the First Person. In Anthony Hatzimoysis (ed.), *Self Knowledge*. Oxford : Oxford University Press, 261–277
- Pacherie, Elisabeth (2008). The Phenomenology of Action: A Conceptual Framework. *Cognition* 107 (1):179 - 217.
- Paul, L. A. (2017). First personal modes of presentation and the structure of empathy. *Inquiry: An Interdisciplinary Journal of Philosophy* 60 (3):189-207.
- Perry, John (1977). Frege on Demonstratives. *Philosophical Review* 86 : 474–497 .
- (1979). The Problem of the Essential Indexical. *Noûs* 13 : 3–21 .
- (1998). Broadening the Mind. *Philosophical and Phenomenological Research* 58 (1):223-231.
- Recanati, François (2013). *Mental Files*. Oxford: Oxford University Press .
- Schwitzgebel, Eric (2015). If materialism is true, the United States is probably conscious. *Philosophical Studies* 172 (7):1697-1721.

Spencer, C. (2007). Is there a problem of the essential indexical? In M. O'Rourke and C. Washington, editors, *Situating Semantics: Essays on the Philosophy of John Perry*. MIT Press, Cambridge.

Titelbaum, Michael G. (2017). One's own reasoning. *Inquiry: An Interdisciplinary Journal of Philosophy* 60 (3):208-232.

Tollefsen, Deborah (2002). Organizations as true believers. *Journal of Social Philosophy* 33 (3):395–410.

Velleman, J. David (2005). *Self to Self: Selected Essays*. Cambridge University Press.

Wilson, George and Shpall, Samuel, "Action", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/action/>.