Do group agents have free will?

Christian List*

This version: April 2023
Published in *Inquiry*, 2023, DOI: 10.1080/0020174X.2023.2218721

Abstract: It is common to ascribe agency to some organized collectives, such as corporations, courts, and states, and to treat them as *loci* of responsibility, over and above their individual members. But since responsibility is often assumed to require free will, should we also think that group agents have free will? Surprisingly, the literature contains very few in-depth discussions of this question. The most extensive defence of corporate free will that I am aware of (Hess 2014) rests on a compatibilist understanding which takes reasons-responsiveness and acting from one's own "actional springs" as key conditions for free will but says little about a condition that free-will libertarians emphasize: the possibility of doing otherwise. In this paper, I will argue that group agents can have free will not only in a (less demanding) compatibilist sense, but also in a recognizably libertarian sense, which includes the possibility of doing otherwise. In developing this account of corporate free will, I will bring together recent work on group agency and recent work on free will.

1. Introduction

It is common to ascribe agency to some organized collectives – from corporations, courts, and non-governmental organizations to states in their entirety – and to treat those collective entities as having goals and beliefs of their own and as taking corporate actions, distinct from the actions of their individual members.[1] When a state declares war against another state, this counts as the action of the state, not just as the action of some individuals. And when a corporation sues another party in the courts, the plaintiff is the corporation, not its managers, employees, or shareholders. Some individuals will have contributed to those actions, but conceptually and legally, it is the state or the corporation as a group agent that is declaring war or entering litigation.

The idea that collectives can be agents over and above their members is useful in several respects. It enables us to describe and explain the behaviour of those collectives more parsimoniously than if we had to unpack all the individual-level actions underpinning them. If we tried to explain the actions of a corporation by enumerating all the contributions of the individual members, or the actions of a state by enumerating all the relevant actions of its citizens, for example, we would miss the forest for the trees. The idea of group agency is also normatively relevant, when we treat group agents as legal persons and when we hold them responsible – morally or legally – for harms they have caused, as in regimes of corporate civil or criminal liability.

---

* This paper brings together some of my work on group agency (especially List and Pettit 2011 and List 2018) with my work on free will (especially List 2014, 2019, 2022). While I will include some references to those works, I will try not to overload the paper with cross-references. Further general literature references can also be found in those prior works. In the concluding chapter of List (2019), I had suggested but not fully developed the application of my account of free will to group agents. For helpful comments, I thank Kendy Hess, Olof Leffler, Lars Moen, the participants of the Research Seminar in Decision and Action Theory at LMU Munich in February 2023, the participants of a Philosophy Colloquium at Yale University in March 2023, and an anonymous reviewer.

[1] For discussions, see, among others, French (1984), Rovane (1997), Pettit (2001, ch. 5, 2003, 2007), Tollefsen (2002, 2015), Erskine (2003), List and Pettit (2006, 2011), Tuomela (2013), Hess (2013, 2014, 2018, 2020), Hindriks (2014), Huebner (2014), and List (2018).

But how far should our commitment to the idea of group agency go? If we think that certain groups are agents of their own and bear responsibility for their actions, should we also think that they have free will? As Kendy Hess has noted, one may object that corporations and other collectives "cannot qualify as moral agents because they lack free will".[2] The worry, she continues, "is that corporations (and other highly organized collectives like colleges, governments, and the military) are effectively puppets, dancing on strings controlled by external forces".[3] Ish Haji, for instance, has argued that group agents face an "uphill battle" if they are to meet the requirement of being "ultimate originators" of their actions.[4] If this is so, the ascription of moral agency to them is at best a metaphor, and there is no basis for treating them as responsible agents on a par with individuals.

Surprisingly, the large scholarly literature on group agency contains very few in-depth discussions of the idea of corporate free will. Hess's paper, in which she also responds to Haji's critique, is a notable exception. She defends the view that corporations and other organized collectives can have free will. It "seems obvious to the point of being a truism", she says, that "[i]t is morally wrong for corporations to lie, cheat, and steal", but for this to be literally true, she notes, those corporations must have moral agency, and this, she suggests, requires free will.[5] Hess argues that if we carefully look at how group agency works and how group agents form their action-guiding intentions, we can see that such agents "act from their own 'actional springs', in Haji's [2006] phrase, and from their own reasons-responsive mechanisms", and she infers from this that "they act freely and are morally responsible for what they do".[6]

Hess offers the most thorough defence of corporate free will that I am aware of, and yet her analysis leaves some questions open. Hess takes reasons-responsiveness and acting from one's "actional springs" as key conditions for free will and says little about another salient condition, namely the possibility of doing otherwise – the idea that an action is not free unless the agent could have acted otherwise. Hess's understanding of free will is in line with the way many *compatibilists* think about free will, for whom the "ownership", "authorship", and/or "intentional endorsement" of one's actions is more central to free will than the possibility of acting otherwise. (This understanding allows those compatibilists to avoid the much-discussed conflict between free will and determinism.) However, Hess's analysis may not speak equally to so-called free-will *libertarians*, for whom the possibility of doing otherwise is essential. Not least because commonsense intuitions often seem to side with libertarianism in presupposing that there cannot be free will without alternative possibilities, it is then natural to ask whether group agents could have free will even in a stronger, libertarian sense.[7]

---

[2] See Hess (2014, p. 241).
[3] Ibid.
[4] See Haji (2006). The quotes are from p. 292.
[5] See again Hess (2014, p. 241).
[6] Ibid.
[7] An empirical study by Sarkissian et al. (2010), for instance, suggests that libertarian intuitions about free will are shared across different cultures. It is worth acknowledging, however, that others have found the evidence about folk intuitions more conflicting and less supportive of libertarianism. See Nahmias et al. (2005).

My aim in this paper is to fill this gap left by Hess's analysis. I will argue that group agents can have free will not only in a (less demanding) compatibilist sense, but also in a recognizably libertarian sense, which explicitly includes the possibility of doing otherwise. So, while group agents do not share all the capacities that individual human agents have – for instance, they arguably lack phenomenal consciousness – free will, even of a certain libertarian kind, is a capacity they may share.[8]

I will begin with a brief introduction to the idea of group agency. I will then explain how I understand free will for the purposes of my analysis. And I will finally address my central question: do group agents have free will?

2. Group agency

Groups agents are collectives that, due to the way they are organized, function as goal-directed agents in their own right, over and above their members.[9] Examples, as already noted, are firms and corporations, multi-member courts, non-governmental organizations, and entire states. When we explain the behaviour of such entities and think about them from a moral or legal perspective, we routinely ascribe beliefs, goals, and intentions to them and view them as capable of acting as if they were persons. The law treats many such entities as legal persons.

When we read, for instance, that "Exxon made 'breathtakingly' accurate climate predictions in 1970s and 80s", "only to then spend decades publicly rubbishing such science in order to protect its core business" (Guardian 2022), we tend to regard this as an act of wrongdoing not only on the part of certain individuals within Exxon (though some individuals may well be blameworthy) but also on the part of the corporation itself. Even if the corporation's membership, which includes its managers and employees, has changed since the time of the wrongdoing, intuitively this does not absolve the corporation of its responsibility.

A skeptic might respond that corporate entities are not literally goal-directed agents like you and me, let alone moral agents who are genuinely capable of bearing responsibility. The ascription of agency and responsibility to a group agent is just a shorthand for a more complicated ascription we could make at the level of the group's individual members. It is a pattern of individual actions that gives rise to certain corporate-level outcomes, and so, the objection goes, when we describe or normatively assess those corporate-level outcomes, what we are really describing or assessing is the pattern of individual actions.

However, explanatory practices in the social sciences suggest not just that this individual-level backstory is often too complicated to tell, but that a satisfactory individual-level explanation may not be available at all. Rather, the ascription of agency to some groups seems at least practically indispensable if we wish to explain certain collective behaviours in the social

[8] Libertarianism about free will has been defended, in different variants, by, e.g., Chisholm (1966), Kane (1999), O'Connor (2000), and Steward (2012). The variant I find most congenial for present purposes is the so-called compatibilist libertarianism defended in List (2019). As explained below, it shares libertarianism's commitment to alternative possibilities as a requirement for free will but offers a distinctly agential interpretation of this notion.
[9] Recall the references in footnote 1.

world.[10] Examples can be found in the theory of the firm in economics, which explains the behaviour of firms and corporations by depicting them as rational agents maximizing expected profits, and in realist theories of international relations, which explain strategic interactions between states through the lens of game theory and rational choice theory. In both areas, there is no neat translation scheme that would allow us to translate talk of corporate-level behaviour into matching sets of statements about purely individual-level behaviour.

Philip Pettit and I have used ideas from judgment-aggregation theory to argue that *even if* the action-guiding attitudes of a group agent are a function of some underlying attitudes of the group's members, this function may be complicated and need not allow a straightforward reduction of the group-level attitudes to corresponding individual-level attitudes.[11] A simple example illustrates this problem.[12]

Suppose the board of a firm is deciding whether to install a new safety device. It is agreed that the safety device should be installed (proposition *S*) if and only if three jointly necessary and sufficient premises are true: the danger to be averted is serious (proposition *P*), the proposed safety device is effective (proposition *Q*), and the costs are bearable for the firm (proposition *R*). Suppose further that the profile of individual judgments among the board members is as shown in Table 1. Here each premise for installing the safety device is supported by a majority of the board members, and yet all board members – each for a different reason – reject the resulting proposal of installing the device. What should the firm do?

Table 1: A profile of individual judgments

	Serious danger? (*P*)	Device effective? (*Q*)	Costs bearable? (*R*)	Install device (*S*)
1/3 of the members	Yes	Yes	No	No
1/3 of the members	Yes	No	Yes	No
1/3 of the members	No	Yes	Yes	No
Majority	Yes	Yes	Yes	No

Suppose, for the sake of our example, the board's decision proceeds as follows. The board takes votes on each of the three premises (*P*, *Q*, and *R*), thereby forming corporate-level judgments that the danger is serious (*P*), the proposed device is effective (*Q*), and the costs are bearable (*R*), and the board then takes itself to be committed to supporting the installation of the new device (*S*), given its prior agreement that *S* should be accepted if and only if *P*, *Q*, and *R* are all accepted. The board thus uses the so-called premise-based procedure for its decision, and this solves the board's decision problem.

However, the group agent has now formed the action-guiding judgment that it should install the new safety device, even though none of the participating members individually share this judgment. Since this is an example of how a group agent might form its action-guiding attitudes

[10] This point has been emphasized by Tollefsen (2002) and List and Pettit (2011), among others. Below I will restate this point in the form of an explicit indispensability argument.
[11] See List and Pettit (2006, 2011).
[12] For the example and a more general analysis, see List and Pettit (2002).

– albeit a simplistic example – it shows that the corporate attitude on a proposition need not be a function of its members' attitudes on the same proposition. Formal results in the theory of judgment aggregation show that this problem is very general and not restricted to specific examples such as the one just given.[13] There may be only a complicated, holistic pattern of dependence between individual and corporate attitudes. We can summarize this observation as follows:

> **The basic non-reducibility thesis:** A group agent's attitude *qua corporate entity* on some issue is not generally reducible to its members' individual attitudes on this issue; it's not generally the majority attitude or some other simple summary of the individual attitudes.[14]

Hess comes to a similar conclusion, stressing the ubiquity of distributed decision-making in organizations:

> "[T]he corporate entity [may] come to believe that x … and to desire that y … in a manner that has no necessary connection to the beliefs or desires of its members regarding x and y. When its members act, they will tend to act in ways that (collectively) express these corporate beliefs and desires rather than their own, possibly contrary opinions. This can be done knowingly and deliberately, as in the case of [a] board vote, or it can be done unknowingly, by masses of people going through the motions of doing their jobs without ever being in a position to see the corporate commitments – the corporate beliefs and desires – that they create [through their distributed decision-making] and then conform to by so doing."[15]

The bottom line is that even if what a group agent does supervenes on what its members do, group-level descriptions of attitudes and actions may not be straightforwardly reducible to corresponding individual-level descriptions. This bolsters the claim that talk about group agency is more than just a shorthand for talk about the participating individuals.

Group agency must not be confused with three related, but distinct phenomena. One is *collective action*, the cooperation of several individuals in the pursuit of some aggregate effect, perhaps in the presence of incentives for individuals to free-ride. Examples are industrial strikes or the cooperation of several farmers or fishermen in solving a tragedy of the commons.[16] A second distinct phenomenon is *joint action*, the coordinated action among several individuals supported by a structure of shared, common, or collectively oriented intentions.[17] Examples are carrying a sofa together or going for a walk together, where the participants "jointly intend" to undertake this activity. A third distinct phenomenon is *team reasoning*, a mode of collectively directed reasoning in interactions with others, in which participants ask themselves

---

[13] See List and Pettit (2002, 2011). For a survey, see List and Puppe (2009).
[14] A version of this thesis is central to the arguments in List and Pettit (2006, 2011, chs. 2 and 3).
[15] See Hess (2014, p. 248).
[16] See, e.g., Olson (1965) and Ostrom (1990).
[17] See, e.g., Bratman (1999, 2014), Gilbert (1989), and Tuomela (2007).

"what is the best course of action for us collectively, not merely for myself individually".[18] Team reasoning can help groups who engage in it to achieve successful collective action and to solve coordination problems.

Collective action, joint action, and team reasoning each involve groups of people acting together in certain ways, but they need not involve a group *agent* in the sense discussed in this paper. The primary *locus* of agency in cases of collective action, joint action, or team reasoning arguably remains the individual.[19] To be sure, the three phenomena may be ingredients in the *implementation* of group agency. For a corporation or some other organization to get its act together, its members may sometimes need to engage in collective action, joint action, or team reasoning. Nonetheless, group agency, which corporations, states, and other organized collectives may sometimes achieve, goes beyond these other phenomena.[20]

3. Free will

Having sketched what I mean by "group agency", I will now turn to "free will".[21] Free will, informally described, is an agent's capacity to choose and control his or her own actions. For instance, when I choose coffee over tea for breakfast, I tend to think that this is a free choice. I choose coffee if this is what I want to drink, but I could also choose tea instead. The choice is up to me. Similarly, when we are faced with more significant choices, many of us think that we sometimes exercise our own free will. Examples are choices between different career paths or whether to get married.

Different scholars differ in how demanding they take free will to be and what they regard as the precise conditions for it. I find it useful to identify three requirements for free will:[22]

(1) intentional agency,

(2) alternative possibilities, and

(3) causal control over one's actions.

More precisely, I will assume the following. First, only systems or entities that are intentional agents could have free will. So, a mere physical automaton, which lacks intentions and does not act in a goal-directed manner, would not qualify. Secondly, any bearer of free will must sometimes have alternative possibilities to choose from. This is the famous "could have done otherwise" requirement. According to it, someone has free will only if, at least in the case of

[18] See Bacharach (2006) and Gold and Sugden (2007).
[19] Note, however, that Gilbert (1989) invokes the notion of a "plural subject" in her analysis of joint action, and theorists of team reasoning (such as Bacharach 2006) sometimes speak of an "agency transformation". Still, I think that we need not posit a group agent in each instance of a joint action, and I also think that what is being transformed in cases of team reasoning is the agency of each participating individual. Arguably, there is no need to postulate an additional corporate-level agent over and above the participating team-reasoners.
[20] Hess (2020) makes a similar point, recognizing that phenomena such as shared intentions, joint commitments, and collective ends may be ingredients in the formation of a corporate "rational point of view (RPV)" (using a term from Rovane 1997) but the corporate RPV goes beyond those other collective phenomena.
[21] For a general overview of the philosophical debate on free will, see Kane (2002).
[22] The present analysis of free will draws on List (2014, 2019, 2022), where further references can be found.

some of their actions, it would have been possible for them to do something else. Thirdly, any bearer of free will must have causal control over their actions. That is, their actions must be caused in the right way by their intentions, not by any sub-intentional physical processes. If my bodily state alone, not my intentional mental state, causes me to do something, it cannot be ascribed to my free will. Think of a reflex, for example. Thus, this third requirement for free will is mental causation.

Some philosophers do not insist on all three requirements. In particular, some free-will compatibilists, as already noted, think of a free action as one performed "authentically", with one's own authorship and intentional endorsement, and they do not require alternative possibilities. They say, for instance, that someone whose actions are sufficiently well supported by his or her character is acting freely, even if it wasn't possible for him or her to act otherwise. A much-cited example is that of Martin Luther, the 16th century church reformer who, when asked to renounce his criticism of the Roman Catholic Church, reaffirmed his criticism, allegedly saying "here I stand; I can do no other".[23] Since many of us would view Luther as autonomous, some commentators interpret the example as evidence that someone can act freely even without alternative possibilities. My preferred interpretation, however, is that Luther didn't literally lack alternative possibilities but that what he said was simply meant to convey the strength of his commitment to his chosen action. Acting otherwise wasn't strictly impossible for him, but it would have meant betraying everything he stood for. Perhaps he felt under some kind of "normative necessity" to do what he did, though it would still have been physically possible for him to do something else.[24]

Another famous challenge to the alternative-possibilities requirement for free will comes from Harry Frankfurt, who asked us to imagine someone – call him Jones – with a brain implant making it the case that *if* Jones doesn't voluntarily choose to do a particular action, say action X, his brain makes him do it anyway.[25] Now suppose that Jones voluntarily chooses to do X. Frankfurt suggests that Jones's action can be attributed to his own free will, and that he can be held responsible for it, even though not doing X was not an option for him. The implant would have made him do it anyway. I think, however, that this example doesn't refute the claim that free will requires alternative possibilities; at most, it prompts us to tweak that requirement. Even Frankfurt's contrived case presupposes that there are distinct possibilities: the one in which Jones voluntarily chooses to do X and the one in which the implant makes him do it. Jones is considered responsible for doing X only if he makes the choice voluntarily, not otherwise. These are alternative possibilities at the level of Jones's mental acts.

Independently of our views about these examples, I will here understand free will in a libertarian sense, according to which the possibility of doing otherwise is central, and so I will assume that free will does require all three properties: intentional agency, alternative possibilities, and causal control. I want to show that group agents can possess free will even in

[23] For discussion, see Dennett (1984) and Kane (2002, introduction).
[24] I thank the editors for suggesting the "normative necessity" interpretation.
[25] See Frankfurt (1969).

this relatively demanding sense. *A fortiori*, they will then also possess free will in the weaker sense compatibilists have in mind.

Of course, my characterization of free will in terms of the three properties still leaves some room for fine-tuning the interpretation of each property. Since the variant of libertarianism that I find most compelling differs somewhat from the most well-known variants of libertarianism, I must comment briefly on how I interpret the three properties.[26]

3.1. Intentional agency

An *intentional agent*, as I understand the notion, is an entity that acts in a goal-directed manner, on the basis of certain intentional mental states, such as beliefs and desires. The standard belief-desire model of agency illustrates how this might work. According to it, an agent is an entity that has beliefs, defined as representations of its environment, and goals/desires, defined as target specifications depicting what the entity "would like to achieve" in the environment, and then intervenes in the environment in pursuit of its motivating specifications in line with its representations.

This analysis is compatible with different ways of spelling out the details, but generally, I assume that a good *test* for intentional agency is whether the entity's behaviour is *best explained* by attributing intentional mental states such as beliefs and desires to it, where those intentional mental states make intelligible or rationalize, at least approximately, what the entity does. For example, since we cannot make sense of human behaviour in many circumstances unless we interpret people as acting based on certain intentional mental states, we have good reasons to think that those people are indeed intentional agents. Something similar might be said about primates. It would be difficult to make sense of the behaviour of chimpanzees, for instance, without taking what Daniel Dennett calls an "intentional stance" towards them: we explain their behaviour by viewing them as goal-directed, intentional agents.[27]

That said, my understanding of agency is not just "interpretivist", taking agency to be in the eye of the beholder, but fully realist. For me, interpretability as an agent is *evidence* for intentional agency, not the defining criterion.[28] To *be* an agent is to be an entity that acts in a goal-directed way based on intentional mental states. Still, if an entity's behaviour is *best explained* by viewing it as an agent, this is good (albeit defeasible) evidence that the entity *is* an agent.

3.2. Alternative possibilities

Let me next turn to *alternative possibilities*. I assume that, for an agent to have alternative possibilities, he, she, or it must sometimes be faced with what decision theorists call "choice

[26] The most significant difference between my preferred variant of libertarianism and better-known variants (such as those cited in footnote 8) lies in the interpretation of "alternative possibilities". As will become clear, I interpret that notion in a way that differs (though in a principled manner) both from traditional compatibilist and traditional libertarian interpretations.
[27] See Dennett (1987).
[28] That's why I speak of a "*test* for intentional agency" when I invoke Dennett's intentional stance.

nodes": points at which the agent can choose between two or more distinct courses of action (or simply "options"). Crucially, each such course of action must be *possible for the agent*. By this, I mean that the agent's trajectory up to any choice node must permit two or more distinct possible continuations, corresponding to the different available options. The decision tree in Figure 1 illustrates this idea. The dots represent choice nodes that the agent may be faced with, and the branches, going from bottom to top, represent possible trajectories the agent could take.

Figure 1: A decision tree

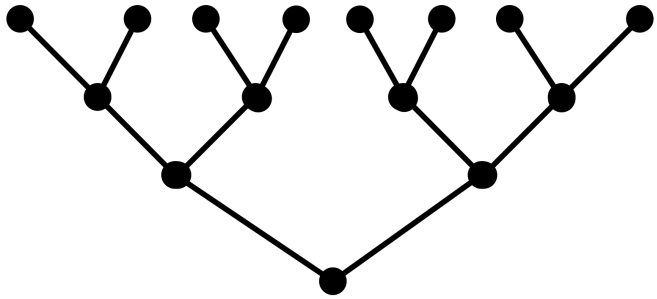

For instance, when I have a choice between coffee and tea, a future in which I choose coffee and one in which I choose tea must each be available to me. If it were already settled that I was going to choose coffee and the choice of tea was simply impossible, I wouldn't be faced with a genuine choice at all.

By insisting that different choices must be possible for me as an agent, I am committing myself to a kind of libertarian understanding of alternative possibilities, not merely a traditional compatibilist one. Unlike those compatibilists who seek to offer a determinism-friendly reinterpretation of alternative possibilities, I don't think it is enough for the alternative possibilities to be available in a purely counterfactual sense. The compatibilists in question say that even if it was completely predetermined that I would have coffee rather than tea, I could have chosen tea *if conditions had been different*. Along these lines, G. E. Moore thought that a sufficient condition for "alternative possibilities" was that I could have chosen something else *if I had wanted to or intended to or if conditions had been different*. This condition is compatible with my actual choice having been predetermined.[29] However, I broadly agree with Susan Hurley, who wrote:

> "The ability to do otherwise entails the *outright possibility* of acting otherwise: it entails that there is a causal possibility of acting otherwise, holding all else constant. A counterfactually conditioned disposition to act otherwise [i.e., the sort of thing G. E. Moore had in mind] is not the same thing as an outright possibility of acting otherwise."[30]

Despite my insistence that "alternative possibilities" means alternative *possibilities*, my view differs from that of mainstream libertarians in one respect: I require an agent's options to be

[29] See Moore (1912).
[30] See Hurley (1999, pp. 205–6).

possible in a distinctively *agential* sense, not the physical sense that Newtonian physics, quantum mechanics, or another physical theory might refer to. Therefore, I would replace Hurley's reference to a "causal" possibility of acting otherwise with a reference to an "agential" possibility. Many libertarians employ essentially the same notion of possibility regardless of whether they speak about possible physical trajectories under the laws of nature or about what an agent can and cannot do. They might call this notion "nomological possibility", "physical possibility", or even just "actual possibility". But independently of the terminology, they tend to assume that the sense of possibility in which it is possible for an agent with free will to do one thing or another is not essentially different from the one in which, given physical indeterminism, it is possible for a photon to take one path or another.[31] I consider this assumption a mistake. We routinely use different notions of possibility for different levels of analysis in the natural and social sciences, such as physical, chemical, biological, psychological, and even socio-economic notions. These can be understood as the notions of possibility supported by our best theories in these domains. Many claims about possibilities that we make in the special sciences (say, about what could happen to the weather, the climate, the economy, the next election, and so on) make sense only if we understand "possibility" in a distinctively macroscopic sense, as supported by our best scientific theories within the relevant domains. Similarly, I suggest that when we analyze what agents can and cannot do, we should employ the notion of possibility supported by our best theories of agency, which I call "agential possibility".

Agential possibility is the notion of possibility in terms of which the decision trees that we find implicitly or even explicitly in intentional explanations in the social sciences (and which decision and game theorists model formally) can be interpreted most literally. According to such decision trees (as in Figure 1), an agent has different possible options at each choice node that he or she faces, and these – at least when interpreted literally – correspond to possibilities among which this agent can genuinely choose, albeit *possibilities at the level of agency*. This reference to a distinctive notion of *agential possibility* is consistent with John Maier's claim that there is a *sui generis* form of "agentive" modality, which he associates with option availability.[32]

3.3. Causal control

Thirdly, I want to comment on my understanding of *causal control*. I suggest that an agent has causal control over an action if and only if the agent's action-guiding intentional state (rather

[31] In particular, the assumption seems to be that any choice node at which an agent has different possible options among which he or she can choose must correspond to a branching point in the physical trajectory of the universe, where different future state evolutions are physically possible. Kane (1999), for instance, suggests that quantum indeterminacies in the brain are relevant to free will, albeit in the context of a sophisticated account of "self-forming actions". An agent's alternative possibilities are then, in effect, macroscopically amplified alternative possibilities from physics. The application of physical-level notions of possibility to the analysis of agency is arguably also the reason why libertarians typically think that free will is incompatible with physical determinism, thereby endorsing van Inwagen's (1975) consequence argument.

[32] See Maier (2015).

than some physical state of the brain, body, and/or environment) is the difference-making cause of that action.[33] Difference-making causation, in turn, requires the truth of two conditionals:

> **The negative conditional:** If the cause property (here the relevant intentional mental state) were absent, then – other things being equal – the effect property (here the performance of the action) would be absent too.
>
> **The positive conditional:** If the cause property were present in nearby counterfactual circumstances, then – enough other things being equal – the effect property would still be present.

This difference-making understanding of mental causation contrasts with the sort of "production" understanding on which some influential critiques of mental causation are based, such as Jaegwon Kim's "causal exclusion argument" and arguably also the famous neuroscientific experiments by Benjamin Libet and others.[34] According to those critiques – some conceptual, others neuroscientific – anything we do as human agents is produced by certain sub-conscious physical processes in the brain and body, and we should regard these as the real causes of our actions; any intentional mental states that are distinct from those physical processes are just epiphenomena.

In response to this critique, I want to note two points. First, if we look at how scientists understand causation in disciplines ranging from the life sciences to economics – and here, I mean causation in general – this supports a difference-making understanding according to which *causes* are systematic difference-making factors for their effects that withstand careful experimental and statistical controls and that could be used for inventions in a system.[35] And second, the most systematic difference-making factors behind human actions can often be found at the intentional, mental level, not at the neural or physical level, even though such higher-level difference-making regularities are implemented by lower-level physical mechanisms. This view has been defended by a number of scholars, who have argued that our actions co-vary more systematically with our intentional mental states than with their precise neural realizers.[36] Those neural realizers, while central for the physical *implementation* of human agency, are too fine-grained to serve as difference-making causes for the resulting actions.

To sum up, then, I suggest that free will requires intentional agency, alternative possibilities between which an agent can choose, and difference-making causal control over the resulting actions. Before asking whether group agents can have free will in this recognizably libertarian sense, I must clarify two more things.

---

[33] This analysis is based on List and Menzies (2009).
[34] See Kim (1998), Libet et al. (1983), and, for discussion, List and Menzies (2017) as well as List (2019, ch. 5).
[35] On this difference-making and interventionist understanding of causation in the sciences, see, among others, Pearl (2000) and Woodward (2003).
[36] See, for example, Woodward (2008), List and Menzies (2009, 2017), Raatikainen (2010), and Roskies (2012).

3.4. Two clarifications

First, free will alone is not the same as fitness to be held morally responsible. While free will is necessary for fitness to be held morally responsible (we wouldn't attribute responsibility to any entity that completely lacks free will), it is not by itself sufficient.[37] Even non-human animals such as primates can qualify as having a basic form of free will, insofar as, arguably, (i) they are goal-directed, intentional agents, albeit without all of the human cognitive capacities; (ii) they face choices between alternative possibilities; and (iii) when they make those choices, they exercise difference-making mental causation. Nevertheless, because they lack the rich capacity for normative cognition that humans have, we wouldn't hold them morally responsible for what they do. This suggests that fitness to be held morally responsible requires the capacity for normative cognition in addition to free will.[38]

Secondly, free will, even in the present libertarian sense, is not the same as what one might call "ultimate control": control not just over the action itself but also over its entire causal pre-history. When I make a choice, even a free choice on my account, I still make that choice in a particular context: I am in a particular mental state, I have certain beliefs and goals, and so on. Perhaps some of this "backstory" was itself the result of free choices I made in the past. But evidently, if we trace back my history further, there comes a point at which I was no longer in control over everything that made me the person I have become. If a choice didn't count as free unless the agent had ultimate control over it, then free will would be an empty concept, whose conditions can never be satisfied. However, free will is a diagnostically helpful concept only if it allows us to distinguish between things agents do freely and for which they can potentially bear responsibility and things they don't do freely and for which the question of responsibility typically doesn't arise. If we were to define free will as ultimate control, the concept would lose its diagnostic usefulness.[39]

I suggest that by defining free will in the carefully qualified libertarian sense I have described we are striking a balance between making free will too "watered down" (as would arguably be the case if we did not require alternative possibilities) and making it too demanding (as would be the case if we insisted on ultimate control).

4. Free will in group agents

We can now return to our main question: do group agents have free will? My three requirements for free will give us a template for approaching it. For any given corporate entity, we must ask:

---

[37] On this notion, see Pettit (2007).

[38] Similarly, Hess (2018) takes "normative competence" to be a necessary condition for moral agency, citing in turn Wallace (1994), who proposes normative competence as a requirement for moral agency (though on Wallace's account, unlike the present one, free will is not required for moral agency), and Hindriks (2014), who discusses the evaluative understanding of corporate agents as a condition for their moral agency.

[39] While my aim here is not literature exegesis, perhaps Haji's (2006) critique of the idea that collectives can be responsible for what they do stems, in part, from a relatively demanding understanding of control. Haji speaks of "ultimate responsibility" and "ultimate origination", though Haji recognizes that this notion can be spelt out in different ways. Be that as it may, I will suggest below that group agents can exercise a form of causal control over their actions, even if it's not ultimate control. I will briefly come back to this issue in the concluding section.

is it an intentional agent; does it have alternative possibilities between which it can choose; and does it have causal control over its actions?

Obviously, there is no blanket answer that applies to all corporate entities. One must look at each entity on a case-by-case basis. But I want to sketch an answer strategy that applies to "textbook group agents": paradigmatic examples such as corporations and firms as understood by the theory of the firm and states as understood by realist theories of international relations, collegial courts, and other purposive organizations.

4.1. Group agents are really intentional agents

As a first step, we must ask whether such entities truly possess intentional agency. Now, in effect, we have already presupposed a positive answer to this, by restricting our consideration to *paradigmatic* group agents. But one might still wonder whether the ascription of agency to them is merely a useful construct. Several authors, including Peter French, Kendy Hess, Deborah Tollefsen, Philip Pettit, and myself, have argued for realism about group agency.[40] For present purposes, I want to briefly (re)state an indispensability argument which supports a realist interpretation of group agency.[41] I don't have the space to fully defend the argument here, but I am stating it to give readers a sense of my preferred argumentative approach. The argument has two premises:

> **Premise 1:** The ascription of intentional agency to the collective entities in question (such as corporations and states) is explanatorily indispensable in the relevant social sciences (which seek to explain those entities' behaviour).
>
> **Premise 2**: If the ascription of some property to an entity is explanatorily indispensable in the relevant sciences, then we are provisionally warranted in assuming that the entity really has that property.

I assume that explanatory practices in the social sciences support the first premise – recall the examples of the theory of the firm and realist theories of international relations – and the second premise is familiar from arguments for realism in the philosophy of science. It is a version of the principle that explanatory indispensability in science is a guide to what is real, sometimes also called the "naturalistic ontological attitude".[42] Scientists often argue for realism about unobservable particles, fields, and forces in physics based on this principle. It is this sort of principle that leads people to think that electromagnetic fields and gravity are real, although we cannot directly observe them. We can only observe certain empirical regularities that are *best explained* by postulating electromagnetic fields or gravity, and this, in turn, makes us confident that there really are electromagnetic fields and gravity.

---

[40] See, e.g., French (1984), Pettit (2001, ch. 5, 2003), Tollefsen (2002, 2015), List and Pettit (2006, 2011), and Hess (2013).
[41] An earlier version of this argument can be found in List (2018); the argument is also implicit in Tollefsen (2002) and List and Pettit (2011).
[42] See, e.g., Quine (1977) and Fine (1984).

If we combine the two premises, we get the following conclusion:

> **Conclusion**: We are provisionally warranted in assuming that the collective entities in question really are intentional agents.

Of course, one might contest whether the ascription of agency to the collectives in question is genuinely indispensable or whether an individualist re-explanation might be theoretically feasible. Lars Moen, for instance, has argued that unless we focus on the individual level, we cannot adequately explain strategic behaviour on the part of the group members.[43] (But note that Premise 1 is compatible with the existence of *some* explanatory purposes for which we must unpack the "black box" of the group agent.) Similarly, one might contest whether explanatory indispensability should really guide our ontological commitments. Scientific instrumentalists, as opposed to realists, would raise this worry.[44] However, for the purposes of this paper, I set these questions aside and note that any skepticism about corporate free will that stems from skepticism about whether the groups in question are genuine agents is not really skepticism about corporate *free will* but merely a consequence of prior skepticism about group *agency*; and there is already a large literature on this topic.[45]

4.2. Group agents have alternative possibilities

So, let me move on to the question of whether the group agents in question have alternative possibilities. Given my earlier clarification, this question becomes whether those corporate entities are sometimes faced with "choice nodes": points at which they can choose between two or more distinct courses of action, and where the group agent's prior trajectory permits distinct possible continuations, corresponding to those different available courses of action. It seems that the answer is clearly "yes". A state might face a choice between conducting a military operation and refraining from doing so. Up to the decision in question, each option may be a genuine possibility that is not ruled out by any prior constraint. Similarly, a firm might face a choice between investing in a new market and not doing so. The firm might enter the deliberative process with an open mind and not be subject to any hard constraint that would rule out one of these courses of action from the beginning. And a court might face a choice between convicting a defendant and acquitting him or her, again (and even by law) approaching the deliberation in an open-minded manner. In each case, our best social-scientific theories will describe this as a genuine choice that the group agent faces, and a natural way to capture this is to say that the different courses of action are *agentially possible* for the corporate entity. Indeed, if we are realists about our best theories of group agency and thus regard those theories'

---

[43] See Moen (2023). I accept that there are some explanatory purposes for which an individual level of analysis is more appropriate than the corporate one. Relatedly, see the discussions of individual incentives and individual control in the context of group agency in List and Pettit (2011, chs. 5 and 6). What I deny is that we can eliminate references to group agency for all explanatory purposes in the social sciences.

[44] An influential "instrumentalist" approach to the philosophy of science is the constructive empiricism of van Fraassen (1980).

[45] As noted, I am adopting a fully realist rather than interpretivist account of agency (as in List and Pettit 2011), and so I must emphasize that I am treating the explanatory indispensability of ascribing agency to certain collectives by itself merely as *evidence* for the presence of group agency. It is only in conjunction with the principle expressed by Premise 2 that I arrive at my realist conclusion about group agency. For a more explicitly interpretivist account of group agency, see Tollefsen (2015).

postulates as at least approximately correct, then we should also be realists about the choice nodes those theories attribute to the group agents in question.

More generally, the ascription of agency to some entity seems of little explanatory use unless it goes along with the ascription of the capacity for choice between alternative possibilities. Intentional explanations, which explain an entity's behaviour by attributing agency to it, are often contrastive: they explain – based on the agent's beliefs, desires, and intentions – why an agent chooses one thing *rather than* another.[46] Unless we think that agents face real choices between alternative possibilities, it is not clear to me that we can understand such explanations as literally true and as being more than just useful narratives. Therefore, I am inclined to think that realism about agency also supports realism about alternative possibilities.[47] And this applies to any entity whose agency we consider real – individual or collective.

One might still wonder whether the proposed realism about alternative possibilities at the corporate level could be justified if it turned out that the individual-level events underlying the group agent's behaviour were deterministic at some micro-level. While I am unclear why one should think that those individual-level events are deterministic – after all, the members of a group agent themselves face choices between alternative possibilities – I want to emphasize that alternative possibilities at the corporate level are in principle compatible with determinism at the individual level. The notion of *agential possibility at the corporate level* is best understood as a macroscopic modal notion, which can be defined as possibility relative to the agential state of the group agent, not as possibility relative to the microphysical or even detailed individual-level state of its members. And it can be argued that the possible *macro*-trajectories of a system can behave indeterministically even if the underlying *micro*-trajectories on which these supervene are deterministic.[48] A simple illustration is given in Figure 2.[49]

Figure 2: Lower-level and higher-level trajectories

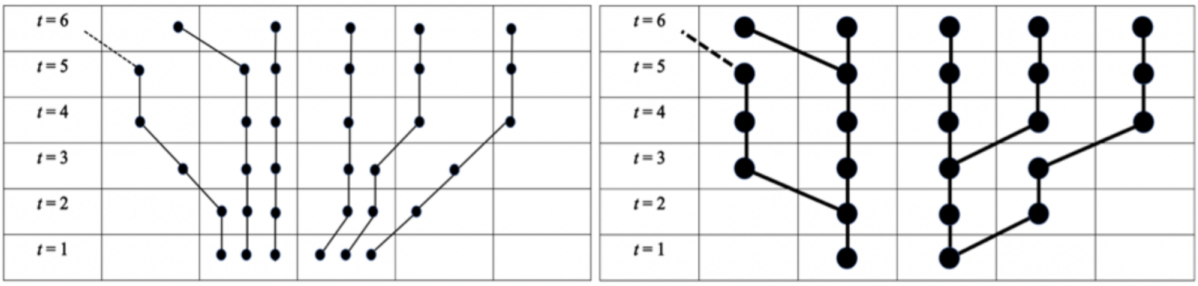

The left-hand side of the figure shows the possible lower-level trajectories of a simple toy system, with the dots representing the states of the system at a particular time and the lines

[46] This is consistent with the claim that agency always involves "two-way powers" (Steward 2012, Alvarez 2013). Roughly speaking, anyone who performs an action could also refrain from doing so or act differently. Intentional explanations must then explain why an agent exercised such a two-way power in one way *rather than* another.
[47] Relatedly, Steward (2012) has argued that agency itself requires alternatives possibilities.
[48] Versions of this result can be found in Butterfield (2012), Yoshimi (2012), List (2014), List and Pivato (2015), and with a different interpretation (observational indistinguishability of deterministic and indeterministic descriptions), Werndl (2009).
[49] This figure is reproduced from List (2019).

representing the trajectories from time $t = 1$ to time $t = 6$. The right-hand side shows the corresponding higher-level trajectories, under the assumption that any distinct lower-level states (thin dots) on the left-hand side give rise to the same higher-level state (thick dot) on the right-hand side: a case of "multiple realizability" of any higher-level state at the lower level. Individual mental states are multiply realizable at the physical level of the brain, for instance; and similarly, corporate intentions in a group agent are multiply realizable at the level of the group's members. As can easily be seen, the higher-level trajectories are indeterministic in this example (they exhibit branching), while the lower-level trajectories are deterministic (there is no branching).

Elsewhere, I have used this observation to argue that alternative possibilities in an individual human agent – understood in coarse-grained terms, as a macro-system – can go along with determinism in the underlying microphysical system.[50] Structurally, the point carries over to the corporate context: indeterminism at the macroscopic level of a complex group agent could go along with determinism at the microscopic level of its individual members.[51]

Nonetheless, I want to stress once again that we have no good reason to assume that the individualistic building blocks of a group agent should be thought of as deterministic; individuals have alternative possibilities themselves. In general, macro-indeterminism, such as at the group level, could also stem from the amplification of micro-indeterminism, such as at the individual level, and in the case of group agents this seems entirely plausible.

4.3. Group agents have causal control over their actions

Let me finally turn to the question of whether group agents have causal control over their actions. Philip Pettit and I have given an affirmative answer to this question in our book.[52] Our strategy was to invoke the Jackson-Pettit "programme model" of causation, according to which one must distinguish between causation as "programming" (a high-level notion) and causation as "implementation" (a low-level notion).[53] We suggested that group agents exercise causal control over their actions in the programming sense, even if the implementation of what they do takes place at the individual level. This analysis was meant to offer a response to the social analogue of the "causal exclusion argument" in the philosophy of mind, already briefly mentioned earlier. The original exclusion argument, well-known from the work of Jaegwon

[50] See List (2014, 2019).
[51] There is one complication. When we analyze the mind-brain relationship in a human individual, it is reasonable to distinguish between two levels that should not be conflated: a lower level, corresponding to the brain, and a higher one, corresponding to the mind. By contrast, it is less clear that we should think of individual and group agency as corresponding to two entirely separate levels. After all, we are dealing with intentional agents at each level. Suppose, for instance, we wish to investigate the interaction between an individual and a group agent, as in a lawsuit in which an individual sues a corporation. When we analyze such mixed individual-corporate interactions and take an intentional stance towards the corporate entity, I suggest we should "black box" the inner workings of the group agent and use the appropriate course-grained notion of agential possibility to analyze what it can and cannot do *qua corporate entity*. On the other hand, when we analyze the group agent's internal dynamics, we must open up that black box, but we may then temporarily suspend the intentional stance towards the corporate entity as a whole. I am grateful to Johannes Kleiner and Toby Solomon for prompting me to address this point.
[52] See List and Pettit (2011, ch. 7).
[53] See parts 1 and 2 of Jackson, Pettit, and Smith (2004).

Kim, concerns the relation between the mind and the brain and asserts that (high-level) mental states, as distinct from physical brain states, cannot be causally efficacious because all the causal work takes place at the level of the physical brain states on which those mental states supervene.[54] The social analogue of this argument concerns the relation between group agents and their members and asserts that certain high-level states of a group agent – such as "corporate intentions" – cannot be causally efficacious because all the causal work takes place at the level of the individual members. It is the individuals who bring about a group agent's decisions and enact the resulting actions. Pettit and I noted that we can bypass this problem if we interpret corporate-level causation as "programming" and individual-level causation as "implementation", which are two distinct modes of causation.

While I continue to think that this analysis works if one adopts the programme model of causation, I have meanwhile amended my view on the causal control of group agents. I now prefer to argue for a group agent's causal control, or "corporate mental causation", by invoking the difference-making account of mental causation mentioned earlier. The key question, on this account, is whether the difference-making causes of a group agent's actions are to be found at the individual level or at the corporate level. In particular, we must respond to the challenge that the mere fact that what a group agent does supervenes on what the members do implies a "crowding out" of difference-making causation at the corporate level, as a Jaegwon-Kim-style causal exclusion argument would suggest.

Given the non-reducibility thesis outlined earlier, I suggest that if a group agent forms the intention to take an action and then acts on its basis, it is typically the corporate intention that qualifies as the difference-making cause of the action, while the underlying profile of individual-level attitudes need not. In cases of genuine corporate actions, the following two conditionals seem to be true:

> **The negative conditional:** If the corporate intention were absent, then – other things being equal – the corporate action in question would not take place.
>
> **The positive conditional:** If the corporate intention were present in nearby counterfactual circumstances, then – enough other things being equal – the corporate action would still take place.

By contrast, no such conditionals would need to hold if we were to look for the causes of the corporate action at the level of the group members' individual attitudes. There are at least three reasons for this, which I will sketch only briefly.

First and most generally, since corporate attitudes and actions are typically multiply realizable at the individual level – insofar as different profiles of individual attitudes and actions may give rise to the same corporate attitudes and actions – the corporate actions need not co-vary systematically with changes in underlying individual attitudes and actions, while they will co-vary systematically with changes in the relevant action-guiding attitudes at the group level.

[54] See, e.g., Kim (1998).

Corporate-level attitudes and actions are usually robust to at least some changes in the underlying individual-level attitudes and actions.[55]

Secondly and relatedly, none of the individuals needs to be pivotal for the corporate attitude and action. Suppose, for instance, the group's attitudes are generated through voting. Unless the pattern of votes produced a knife-edge result (such as a majority of one), changes in individual votes will typically be insufficient to change the collective result. So, the negative conditional for difference-making may not be satisfied by any of the individual attitudes. Especially in large groups, it is unlikely that any single individual is pivotal.

And third, the relevant action-guiding corporate attitude need not even be shared by any of the individual members, as illustrated by the example of the safety device. In this example, it would be hard to identify any individual's action-guiding intention – held *qua* individual – that could qualify as the difference-making cause of the installation of the safety device. By contrast, the group's intentional attitude on proposition *S* satisfies both the positive and the negative conditionals for difference-making. As noted, such a disconnect between individual and corporate attitudes can arise very generally, because there may be only a "holistic" pattern of dependence between individual and corporate attitudes. The corporate intention may also be formed in a path-dependent way. A series of prior decisions may commit the group agent to a particular action, even though an unconstrained vote at the present time would reveal very little support for it among the members.

Finally, one should note that even in those special cases in which the precise profile of individual-level attitudes turns out to be a difference-making cause of the corporate action, for example in a knife-edge situation where even the slightest change in individual attitudes would systematically change the resulting corporate-level action, this still need not undermine the truth of the positive and negative conditionals for *corporate-level* mental causation. It will remain true that if the corporate attitude had been the same (in nearby counterfactual circumstances), the corporate action would have happened, and if the corporate attitude had been different, it would not. In a multi-level system – cell-organism, brain-mind, individual-collective, and so on – there may be special cases in which there is difference-making causation at both the lower and the higher levels.[56] What matters for corporate-level mental causation is that the corporate-level intention *is* a difference-making cause of the corporate action, not that the profile of individual-level attitudes is not (though, as I have emphasized, in typical cases, it isn't).

Putting all this together suggests that at least *paradigmatic* group agents do indeed possess intentional agency, alternative possibilities between which they can choose, and difference-making causal control over their actions. On my account, they then qualify as having free will.

---

[55] The point that certain social-level regularities may be "robust to changes in their individual-level realization" is also emphasized in List and Spiekermann (2013).

[56] In List and Menzies (2009), this is called a "compatibility result" concerning low-level and high-level difference-making causation.

5. Concluding remarks

My analysis suggests that once we are realists about group agency, the step to corporate free will, even with alternative posibilities, is actually quite straightforward. This should not be too surprising on reflection. If we weren't already implicitly relying on certain assumptions about corporate free will, we would have a hard time defending the claims about corporate responsibility that many people are perfectly willing to make. Kendy Hess established this conclusion already based on (what I interpret as) a compatibilist understanding of free will. Here, I have argued that a similar conclusion can still be reached even if we understand free will in a recognizably libertarian manner.

Returning to Haji's question of whether group agents could qualify as the "ultimate originators" of their actions, my answer is that it depends on what we mean by "ultimate origination". If it requires that group agents exercise what I have called "ultimate control" over their actions, that is, control not just over the action itself but also over its entire causal pre-history, then the answer is clearly "no", but this is unsurprising insofar as individual human beings do not have such ultimate control over their actions either. As I have pointed out, the notion is too demanding to be diagnostically useful. If, on the other hand, the "origination" question is whether group agents are real intentional agents to whom actions can be ascribed, whether they make genuine choices between alternative possibilities, and whether they exercise causal control ("corporate-level mental causation") over the resulting actions, then the answer is certainly "yes". And arguably, it is this kind of control that is necessary for fitness to be held responsible, not the unrealistically demanding form of "ultimate control" that no-one has.

I want to note, in conclusion, that my analysis works only if, as I assume, the issue of free will is separate from the issue of phenomenal consciousness. I assume that free will is a feature of agency, and that it lends itself to a third-personal and functionalist analysis, like agency itself. Although free will and consciousness appear to co-occur in the paradigmatic case of human beings, I see no conceptual problem with attributing free will to an agent that lacks phenomenal consciousness. Group agents and perhaps some kinds of AI systems may be examples of such agents. As I have argued elsewhere, group agents lack phenomenal consciousness, and in the AI case, we are arguably closer to having sophisticated AI *agents* than we are to having phenomenally conscious AI systems.[57] Conversely, one could coherently imagine an entity that is conscious but lacks free will. Some free-will skeptics presumably think that human beings are like this. While I don't consider that view plausible, it still seems coherent. And so, free will and phenomenal consciousness should be thought of as separate properties. Still, someone who endorses a more demanding notion of agency, with some notion of phenomenal consciousness built in, is unlikely to agree with my analysis.

From my perspective, however, the picture according to which group agents have free will and can be held responsible but lack phenomenal consciousness makes good sense. Philip Pettit and I, but also others, have suggested that group agents are moral agents, but not genuine moral

[57] See List (2018).

patients: they are not objects of intrinsic moral concern.[58] They don't, for instance, have the kinds of non-derivative and non-instrumentally justified rights that we human beings possess. The present analysis is consistent with this idea. Overall, I conjecture that if we are reluctant to accept the idea of corporate free will, this is likely to be a symptom of a prior unease about the idea of group agency itself.

References

Alvarez, M. (2013). Agency and Two-Way Powers. *Proceedings of the Aristotelian Society* 113(1, pt. 1): 101–121.

Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton, NJ (Princeton University Press).

Bratman, M. E. (1999). *Faces of intention: Selected essays on intention and agency*. Cambridge (Cambridge University Press).

Bratman, M. E. (2014). *Shared agency: A planning theory of acting together*. Oxford (Oxford University Press).

Briggs, R. A. (2012). The Normative Standing of Group Agents. *Episteme* 9(3): 283–291.

Butterfield, J. (2012). Laws, Causation and Dynamics at Different Levels. *Interface Focus* 2(1): 101–114.

Chisholm, R. (1966). Freedom and Action. In K. Lehrer (ed.), *Freedom and Determinism* (pp. 11–44). New York (Random House).

Dennett, D. (1984). *Elbow Room: The Varieties of Free Will Worth Wanting*. Cambridge, MA (MIT Press).

Dennett, D. (1987). *The intentional stance*. Cambridge, MA (MIT Press).

Erskine, T. (ed.) (2003). *Can Institutions Have Responsibilities? Collective Moral Agency and International Relations*. London (Palgrave).

Fine, A. (1984). The Natural Ontological Attitude. In J. Leplin (ed.), *Scientific Realism* (pp. 83–107). Berkeley (University of California Press).

Frankfurt, H. (1969). Alternate Possibilities and Moral Responsibility. *Journal of Philosophy* 66(23): 829–839.

French, P. A. (1984). *Collective and corporate responsibility*. New York (Columbia University Press).

Gilbert, M. (1989). *On social facts*. London (Routledge).

Gold, N., and R. Sugden (2007). Collective intentions and team agency. *Journal of Philosophy* 104(3): 109–137.

Guardian (2023). Revealed: Exxon made 'breathtakingly' accurate climate predictions in 1970s and 80s. *The Guardian*, 12 January 2023. <https://www.theguardian.com/business/2023/jan/12/exxon-climate-change-global-warming-research>, accessed on 10 February 2023.

[58] See List and Pettit (2011) and Hess (2013). For helpful critical discussions of the view that group agents are moral agents but not moral patients, see Briggs (2012) and Hindriks (2014). As I see it, a key to answering some of the important objections to this view raised in the literature lies in recognizing the difference between human moral agents and corporate moral agents with respect to phenomenal consciousness (List 2018, Hess 2018).

Haji, I. (2006). On the ultimate responsibility of collectives. *Midwest Studies in Philosophy* 30(1): 292–308.

Hess, K. M. (2013). "If You Tickle Us….": How Corporations Can Be Moral Agents Without Being Persons. *Journal of Value Inquiry* 47(3): 319–335.

Hess, K. M. (2014). The free will of corporations (and other collectives). *Philosophical Studies* 168(1): 241–260.

Hess, K. M. (2018). Does the Machine Need a Ghost? Corporate Agents as Nonconscious Kantian Moral Agents. *Journal of the American Philosophical Association* 4(1): 67–86.

Hess, K. M. (2020). Assembling the Elephant: Attending to the Metaphysics of Corporate Agents. In S. Bazargan-Forward and D. Tollefsen (eds.), *The Routledge Handbook of Collective Responsibility* (pp. 113–126). New York (Routledge).

Hindriks, F. (2014). How Autonomous Are Collective Agents? Corporate Rights and Normative Individualism. *Erkenntnis* 79(9): 1565–1585.

Huebner, B. (2014). *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford (Oxford University Press).

Hurley, S. (1999). Responsibility, Reason, and Irrelevant Alternatives. *Philosophy and Public Affairs* 28(3): 205–241.

Jackson, F., P. Pettit, and M. Smith (2004). *Mind, Morality, and Explanation: Selected Collaborations*. Oxford (Oxford University Press).

Kane, R. (ed.) (1999). Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism. *Journal of Philosophy* 96(5): 217–240.

Kane, R. (ed.) (2002). *The Oxford Handbook of Free Will*. Oxford (Oxford University Press).

Kim, J. (1998). *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA (MIT Press).

Libet, B., C. A. Gleason, E. W. Wright, and D. K. Pearl (1983). Time of Conscious Intention to Act in Relation to Onset of Cerebral Activity (Readiness-Potential): The Unconscious Initiation of a Freely Voluntary Act. *Brain* 106: 623–642.

List, C. (2014). Free will, determinism, and the possibility of doing otherwise. *Noûs* 48(1): 156–178.

List, C. (2018). What is it like to be a group agent? *Noûs* 52(2): 295–319.

List, C. (2019). *Why Free Will is Real*. Cambridge, MA (Harvard University Press).

List, C. (2022). The Naturalistic Case for Free Will. In S. Ioannidis et al. (eds.), *Levels of Reality in Science and Philosophy, Jerusalem Studies in Philosophy and History of Science*. Cham (Springer).

List, C., and P. Menzies (2009). Non-reductive physicalism and the limits of the exclusion principle. *Journal of Philosophy* 106(9): 475–502.

List, C., and P. Menzies (2017). My Brain Made Me Do It: The Exclusion Argument against Free Will, and What's Wrong with It. In H. Beebee, C. Hitchcock, and H. Price (eds.), *Making a Difference: Essays on the Philosophy of Causation* (pp. 269–285). Oxford (Oxford University Press).

List, C., and P. Pettit (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy* 18(1): 89–110.

List, C., and P. Pettit (2006). Group agency and supervenience. *Southern Journal of Philosophy* 44(S1): 85–105.

List, C., and P. Pettit (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford (Oxford University Press).

List, C., and M. Pivato (2015). Emergent Chance. *Philosophical Review* 124(1): 119–152.

List, C., and C. Puppe (2009). Judgment Aggregation: A Survey. In P. Anand, C. Puppe, and P. Pattanaik (eds.), *Oxford Handbook of Rational and Social Choice* (pp. 457– 482). Oxford (Oxford University Press).

List, C., and K. Spiekermann (2013). Methodological Individualism and Holism in Political Science: A Reconciliation. *American Political Science Review* 107(4): 629–643.

Maier, J. (2015). The Agentive Modalities. *Philosophy and Phenomenological Research* 90: 113–134.

Moen, L. (2023). Eliminating Group Agency. *Economics and Philosophy* 39(1): 43–66.

Moore, G. E. (1912). *Ethics*. Oxford (Oxford University Press).

Nahmias, E., S. Morris, T. Nadelhoffer, and J. Turner (2005). Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility. *Philosophical Psychology* 18(5): 561–584.

O'Connor, T. (2000). *Persons and Causes: The Metaphysics of Free Will*. New York (Oxford University Press).

Olson, M. (1965). *The logic of collective action: Public goods and the theory of groups*. Cambridge, MA (Harvard University Press).

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge (Cambridge University Press).

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. New York (Cambridge University Press).

Pettit, P. (2001). *A theory of freedom: From the psychology to the politics of agency*. New York (Oxford University Press).

Pettit, P. (2003). Groups with minds of their own. In F. Schmitt (Ed.), *Socializing metaphysics* (pp. 167–193). New York (Rowan and Littlefield).

Pettit, P. (2007). Responsibility incorporated. *Ethics* 117(2): 171–201.

Quine, W. V. (1977). *Ontological Relativity and Other Essays*. New York (Columbia University Press).

Raatikainen, P. (2010). Causation, Exclusion, and the Special Sciences. *Erkenntnis* 73(3): 349–363.

Roskies, A. L. (2012). Don't Panic: Self-Authorship without Obscure Metaphysics. *Philosophical Perspectives* 26: 323–342.

Rovane, C. (1997). *The bounds of agency: An essay in revisionary metaphysics*. Princeton, NJ (Princeton University Press).

Sarkissian, H., A. Chatterjee, F. de Brigard, J. Knobe, S. Nichols, and S. Sirker (2010). Is Belief in Free Will a Cultural Universal? *Mind and Language* 25(3): 346–358.

Steward, H. (2012). *A Metaphysics for Freedom*. Oxford (Oxford University Press).

Tollefsen, D. P. (2002). Collective intentionality and the social sciences. *Philosophy of the Social Sciences* 32(1): 25–50.

Tollefsen, D. P. (2015). *Groups as agents*. Cambridge (Polity Press).

Tuomela, R. (2007). *The philosophy of sociality: The shared point of view*. Oxford (Oxford University Press).

Tuomela, R. (2013). *Social ontology: Collective intentionality and group agents*. Oxford (Oxford University Press).

Van Inwagen, P. (1975). The Incompatibility of Free Will and Determinism. *Philosophical Studies* 27(3): 185–199.

Van Fraassen, B. C. (1980). *The Scientific Image*. Oxford (Oxford University Press).

Wallace, R. J. (1998). *Responsibility and the Moral Sentiments*. Cambridge, MA (Harvard University Press).

Werndl, Charlotte (2009). Are Deterministic Descriptions and Indeterministic Descriptions Observationally Equivalent? *Studies in History and Philosophy of Science Part B* 40(3): 232–242.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. New York (Oxford University Press).

Woodward, J. (2008). Mental Causation and Neural Mechanisms. In J. Hohwy and J. Kallestrup (eds.), *Being Reduced: New Essays on Reduction, Explanation, and Causation* (pp. 218– 262). Oxford (Oxford University Press).

Yoshimi, J. (2012). Supervenience, Dynamical Systems Theory, and Non-reductive Physicalism. *British Journal for the Philosophy of Science* 63(2): 373–398.