Role Obligations to Alter Role Obligations

Stephanie Collins

This is the final draft of a paper who definitive version will be published in:
Alex Barber and Sean Cordell (eds), *The Ethics of Social Roles*. Oxford University Press.

1. Introduction

Many roles are situated within collective agents. These include occupational roles (e.g., 'financial advisor'), community roles (e.g., 'chair of the condominium owner assembly'), and political roles (e.g., 'citizen'). There are many situations in which the demands of such roles appear to conflict with the demands of morality.

For example, a financial advisor might face the demand to satisfy criteria for probation, where this requires putting advisees into high-risk investments, which in turn requires giving advice that is not in advisees' best interests. Or the chair of a condominium owner assembly might face the demand to lower the costs of rubbish disposal within the condominium building, where this requires disposing of rubbish in environmentally unsustainable ways. Or a citizen might face the demand to pay taxes that will fund a clearly unjust war . When such conflicts arise, role-occupants face the question: does morality trump roles, or do roles trump morality? Many forms of role-based protest—such as employees' whistle-blowing or citizens' taxation boycotts—can be seen as motivated by the conviction that morality trumps roles.

This chapter argues that the question, as just stated, often contains a false presupposition. The false presupposition is that there is a fundamental conflict between the demands of the role and the demands of morality. I suggest we need a capacious understanding of the 'role' at issue and the 'demands' the role produces. With a sufficiently capacious characterisation of the role and its demands, we can see that the demands of the role are often (though not always) concordant with the demands of morality, such that there is no conflict between the role and morality. The result is that whistle-blowing or taxation boycotts (as well as more conservative forms of role protest) can be seen as the *fulfilment* of role obligations—not the flouting of role obligations in the service of a morality that is 'external' to the role and its obligations.

The general proposal is this. When a role-versus-morality conflict appears to arise, role-occupants acquire an obligation to use (what I will call) the 'fundamental' demands of their

role to do what they can to challenge or alter (what I will call) the 'imposed' demands of their role. These 'role obligations to alter role obligations' are distinctively *role* obligations, rather than garden-variety moral obligations, insofar as the role plays an ineliminable part in justifying (and determining the content of) the obligation to alter the role obligation.

The precise demands of role-obligations-to-alter-role-obligations will vary with context. Sometimes, these obligations might require whistleblowing or civil disobedience; at other times, they might require 'working within the system' in more incremental ways. I'm neutral on those details. The more basic, and theoretical, point of the chapter is this: insofar as there appears to be a tension between organisationally-embedded role obligations and moral obligations, this tension can often be resolved via a proper conceptualisation of what the 'role' amounts to. Embracing the role's fundamental demands might require rejecting the role's imposed demands, but it (often) doesn't require rejecting morality.

I begin in Section 2 by laying out a view of what organisations are and how roles relate to them. This demonstrates the breadth of the account and the features of organisations that produce (seeming) role-versus-morality conflicts. Section 3 elaborates on how (seeming) role-versus-morality conflicts arise. To resolve the conflicts, Section 4 distinguishes between the 'fundamental,' 'imposed,' and 'derived' obligations of an organisationally-embedded role, using this distinction to explain how the seeming conflict can be resolved in a wide range of cases, such as the examples mentioned above. Section 5 concludes.

2. Organisational Roles

As I'll use the term, an 'organisation' is a collective agent that is composed of over 150 people, where those people realise a structure that coordinates their divided labour. The coordination occurs via rules and hierarchical command relations. All of this is guided by a collective decision-making procedure.[1] Organisations are a species of the genus that is 'collective

[1] I develop this definition in Collins (ms, ch. 1). This definition follows a long tradition in sociology, including Max Weber's (1968, vol. I, 223ff; 1968, vol. III, 956ff.) characterisation of bureaucracies, Rom Harre's (1979) characterisation of institutions, and Jonathan Turner's (1997) characterisation of institutions. See also Geoffrey M. Hodgson's (2007) definition of organisations in management theory. Within philosophy, similar characterisations can be found in the work of Peter French (1984, 13ff.), Raimo Tuomela (2013), and Lisa Herzog (2018).

agents.'[2] A collective agent is any entity that is composed of people who are united under a rationally-operated collective decision-making procedure. Amongst collective agents, organisations are distinctive in a few ways: organisations have a large size, while other collective agents might have as few as two members; organisations have a structure, which less organised collective agents might lack; organisations divide role-occupants' labour, whereas other collective agents might be less prescriptive on this front; and organisations have rules and hierarchical command relations, which more egalitarian collective agents might lack. As I'll now explain, these various distinctive features imply that organisations are particularly apt to produce the role-versus-morality conflicts that this chapter aims to address.

To begin to see how organisations are ripe for role-versus-morality conflicts, consider the sense in which organisations have a 'structure.' A structure is a collection of roles that stand in relations. An organisation's structure thus produces a differentiation between various roles: the roles are differentiated by the relations they stand in to one another. (In those collective agents that lack a structure, there is therefore also a lack of 'roles,' in the sense I'll use that term.) In an organisation, the roles are jobs or tasks, while the relations between the roles are usually reporting and delegation lines. The roles and their relations can be highly diverse, and more than one person can occupy each role. Both the roles and the relations are representable in an organisation chart.

I follow Katherine Ritchie (2013) in identifying the organisation with a realisation of a structure. A structure becomes realised—and, thus, an organisation exists—when a sufficient number of the roles in the structure's chart are occupied. Exactly what number of (and which) roles must be occupied, in order for the organisation to exist, is a matter determined by the structure itself.

By viewing organisations as realisations of structures, we can integrate them into a naturalistic metaphysics. That is, by viewing organisations as realised structures and by including those realised structures in our ontology, we do not posit anything 'spooky' or 'mysterious.' Yet including organisations in our ontology allows us to attribute purposes or goals to them. Those purposes and goals are important for role obligations, since, as we'll see, organisations' purposes and goals are the grounds of role-occupants' role obligations. An

[2] The latter have been theorised by, e.g., Carol Rovane (1998), Christian List and Philip Pettit (2011), Deborah Tollefsen (2015), Kendy Hess (2018a, 2018b), Frank Hindriks (2018), Collins (2019), and others.

organisation's purposes and goals are dictated by its structure, including when that structure is used to produce new purposes and goals. And an organisation's purposes and goals produce a kind of internal normativity for organisations—which, in turn, produces (seeming) role-versus-morality conflicts.

I said above that roles are jobs or tasks, which are related by delegation and reporting lines. For the argument that follows, it'll be important to appreciate the breadth and diversity of norms that are entailed by any given 'job' or 'task.' I'll assume that the demands of a job or task include all norms that are operative within the organisation, as well as the more specific requirements of the particular role. Following Amie Thomasson (2019), the demands of an organisation can broadly be divided into 'internal,' 'structuring,' and 'external' demands. As Thomasson puts it, internal demands regard "how members of the group are to behave, regard themselves and other group members, and so on (what they are to wear, how they are to eat and prepare food, what they are (and are not) to do…)" (2019, 4838). Structuring demands "place different group members at different nodes, with different norms regarding those who occupy different nodes (members of the congregation, the priest, the organist, the bishop; the president, secretary, treasurer, member-at-large, new initiate; privates, corporals, lieutenants, colonels…)" (2019, 4838). Finally, external demands are "norms regarding how members of that group are to be treated, regarded, behaved towards by those who are not members of the group. … one way social groups may be constituted is by their members having shared external norms of treatment (privileging, subordinating, or just different), based on any of many kinds of different 'markers'…" (2019, 4839) For example, an external demand on members of a bank might come from existing and potential customers: all members of a bank might face the external demand from non-members to do what they can to satisfy (potential) customers' banking needs.

Thomasson is thinking of social groups, such as genders, races, and classes—not organisations in the present sense. But these categories of demands or norms also apply to role-occupants in organisations. I will use a broad concept of a role's 'demands,' which subsumes all the internal, structuring, and external norms associated with that role—not just the structuring norms, which are perhaps what would first spring to mind. As I hope is clear, the internal norms are perhaps the most ready source of role-versus-morality conflict. Internal norms in an organisation might include norms of competitiveness, selfishness, ambition, cost-cutting, and so on. Such internal norms can often be rationalised by their usefulness in achieving the organisation's goals (such as selling bank loans)—even if the internal norms, and the goals they serve, conflict with morality.

So much for the notion of organisations as constituted by role-occupants in a structure. When characterising organisations, I mentioned other ways in which organisations are distinctive amongst collective agents: organisations have over 150 role-occupants; organisations divide those role-occupants' labour; and organisations have rules and command relations. These three features compound to create a high propensity to role-versus-morality conflicts.

First, an organisation's size—over 150 role-occupants—means that role occupants cannot all be known to each other personally. (150 is 'Dunbar's number': the maximum number of humans with which any human can have meaningful contact.) Due to simple human cognitive limitations, many role-occupants will therefore tend to interact with one another in ways shaped by each other's role. Many role-occupants in the organisation will not appear to one another as unique individuals worthy of particularised engagement, but rather as faceless and nameless occupants of roles or role-types. For an example familiar to those in university settings, we can think of the specific role 'lecturer in philosophy' or the general role 'academic.' Both of these roles mediate one's relationship with many of one's non-immediate colleagues. Such mediation often determines the demands, expectations, and hopes that we place upon one another. The demands or expectations that come from others within our organisation are unlikely to be tailored to our specific circumstances, but rather derive from our role.

Second, organisations' *division of labour* further entrenches the salience of roles, when it comes to an individual's own reasoning about what she should do in the organisational context. The division of labour (combined with organisations' size) implies that role-occupants don't know all the details of what all other role-occupants are doing vis-à-vis the organisation's goals and plans. (This differs from not knowing them personally, which was the point made immediately above.) This epistemic blindspot makes it difficult—if not impossible—to work out what one should do: if I don't know what others in my organisation are doing, then I can't take it upon myself to confidently work out exactly what I should do in response to what they are doing. In the face of limited information about others' precise tasks, it may be perfectly reasonable for individuals to trust 'the system'—at least if they have some independent reason to believe that the system is overall abiding by moral principles. Trusting the system means performing one's role while trusting that others are performing complementary roles and trusting that the roles together abide by morality's demands. Such trust can be reasonable even when morality's demands are in fact being violated (either by other role-occupants, or by the system as a whole).

Third, organisations have *rules and command relations*. That is, not only are people structured into a role-based system, and not only are there are a lot of those people, and not only do those people not know what one another is doing—additionally, all of that happens in the presence of hierarchy. I mentioned this when characterising an organisation's structure, when I noted that the relations between roles will often be 'reporting and delegation' lines. Hierarchies—that is, rules and command relations—have the function of turning expectations or hopes into instructions or demands. From the perspective of any given role-bearer, others within the organisation have the role of commanding, instructing, admonishing, reprimanding, and sanctioning that role-bearer (except, perhaps, the person at the top—though even here, there are usually shareholders or other top-dogs that can remove an errant director). The result is that the 'role' side of role-versus-morality conflicts are explicit demands or instructions from particular other role bearers, where those demands carry sanctions. This increases the psychological force of role-based demands.

These special features of organisational roles are not confined to the commercial world, or even to organisations in which role-bearers are formally employed for pay. As Section 1's examples demonstrated, organisational roles (in my sense of that term) cross-cut Michael Hardimon's (1994, 353) distinction between roles in civil society (which he views as essentially voluntary) and roles in the family and state (which he views as essentially non-voluntary). As I have characterised organisations, some organisational role-occupancy is voluntary, while some is non-voluntary; of those that are voluntary, some are formally contracted, while others are not. As should now be clear, states are organisations, on my characterisation (and democratic states include 'citizens' as role-bearers). Likewise, many entities in 'civil society' are organisations—such as clubs, churches, associations, committees, charities, and so on. Organisations are ubiquitous—and so, therefore, are seeming conflicts between organisational-role obligations and moral obligations. I now turn to consider those conflicts

3. Role-Versus-Morality Conflicts

In the previous section, I noted that it will often be reasonable for organisational role-occupants to perform their role as instructed. The reasonableness of doing this is generated by two broad features of organisational roles: (1) the difficulty (sometimes, impossibility) of knowing exactly what others in the organisation are doing (produced by organisations' large size and division of labour), and (2) the reasonableness of believing that the organisation overall

functions in ways that are morally benign. Of course, features (1) and (2) do not always hold. When they do not, role-occupants may face a conflict between the role's requirements and independent moral norms. How might such conflicts arise?

First, consider a case in which feature (1) fails to hold: a case in which, despite an organisation's size and divided labour, a role-occupant comes to know the details of how others are performing their roles—and, specifically, comes to know that these other role-occupants are not acting in accordance with the organisation's overall benign purpose. For a concrete example, consider the revelation in 2018 that Australia's largest banks had engaged in widespread and pervasive dishonest practices. These practices included charging fees when no service had been provided (including to deceased customers), lying to customers, forging customers' signatures, impersonating customers, falsely witnessing documents, transferring customers' funds to advisors' personal bank accounts, underpaying interest on term deposits, and other such problematic behaviours (Royal Commission 2018; Royal Commission 2019).

These practices were embedded in problematic organisational cultures—yet they nonetheless required actions on the part of particular role-occupants. If one role-occupant knows that another is engaged in such practices, then the moral question is raised of whether the first role-occupant can, or should, continue in their role as normal. This is not just a question of performing actions outside the organisation, such as publicly 'blowing the whistle.' It's also a question of acting within the organisation, for example using their role to act within the organisation to challenge the problematic conduct. Both kinds of reaction are potentially demandable by morality, yet both potentially conflict with demands (explicit or implicit) that issue from the organisation's rules and command relations. Thus we have a (seeming) role-versus-morality conflict.

A second way in which role-versus-morality conflicts might arise is via a failure of feature (2): a case in which it's not reasonable for a role-occupant to believe that the organisation overall functions in ways that are morally benign. Unlike a failure of feature (1)—which was a failure of other role occupants—a failure of feature (2) is failure at the holistic level of the organisation itself. Such failures are not always identifiable with or reducible to failures at the individual level (French 1984; Copp 2007; Pettit 2007). Again, the Australian banking scandal demonstrates this possibility. The Australian government tasked a Royal Commission with investigating potential misconduct in the banking sector and providing recommendations for how the government should respond. In its report, the Royal Commission attacked the "culture" of "dishonesty and greed" in Australia's largest banking and finance corporations (Royal Commission 2018, 73; 2019, 138). The Royal Commission resisted the

idea that this misconduct was a matter of "a few bad apples," instead insisting that the banks themselves (as distinct from individual role-occupants) receive some share of the responsibility (2018, 87-88).

It's easy to imagine how an individual role-occupant might come to be aware of such a culture—and, thus, come to have a reasonable basis to doubt the organisation's moral credentials—even without any knowledge of which precise behaviours by which precise other role-occupants contributed to this culture. To see this, consider the example I gave in Section 1: a financial advisor might face the demand to satisfy criteria for probation, where this requires convincing their advisees to adopt high-risk investments (which are in the long-run financial interests of the bank), which in turn requires giving advice that is not in advisees' best interests. Such demands arise out of an organisation's rules and command relations, including internal norms around profit maximisation and the instrumental treatment of customers. Yet due to the organisation's large size and division of labour, such demands may not come from any identifiable role-bearer. Indeed, it's possible that such problematic organisation-level demands result from a combination of individual-level decisions that are themselves entirely benign (Pettit 2007; List and Pettit 2011; Hess 2014).

Failures of features (1) and (2) are not the only mechanisms by which the demands of roles can come to conflict with the demands of morality, though they are illustrative of the general idea. In general, a role-versus-morality conflict is any tension between, first, the demands imposed on a role-occupant in virtue of her role, where these include informal internal norms, and, second, the demands that would ideally be imposed on the role-occupant, were the whole system concordant with morality's demands, or the demands the role-occupant would abide by, were she being morally responsive to the system's moral failings. (In the final section, I will address cases where morality mandates the abolition of the organisation.)

It's important to see that role-versus-morality conflicts do not have an easy resolution. As emphasised in Section 2, it is often perfectly reasonable for a role-occupant to go along with the demands imposed on her. Those demands cannot simply be thrown out the window as soon as a failure of features (1) or (2) presents itself. After all, there are good moral reasons why we normally toe the organisational line: we are cognitively limited individuals who usually cannot know what others in our organisation are doing, and we often have good reason to believe our organisations are acting for the good. These facts create morally powerful presumptions of

obedience, which it would often be hubristic to override.[3] As Cristina Bicchieri and Ryan Muldoon put it when discussing social (rather than organisational) norms:

> Most of the time, we are not aware of our expectation [that others expect us to comply with a norm], and compliance may look like a habit, thoughtless and automatic, or it may be driven by feelings of anxiety at the thought of what would happen if one transgresses the norm. Upholding a norm is not a matter of conscious cost/benefit calculations; rather, people tend to repeat patterns of behavior that they have learned and, on average, work well in a variety of situations. (Bicchieri and Muldoon 2014)

Thoughtless and automatic habits are difficult to override, and it's not obvious that we have obligations to do so if they have served us well in the past.

Moreover, even if one could override the presumption that favours organisational obedience, it's far from obvious that an individual should 'do the right thing' in circumstances where others (whether individual or organisational) are doing the wrong thing. Such high-mindedness often makes one a target of resentment, discipline, and accusations of virtue signalling—exactly the opposite of what will change the problematic conduct or culture within the organisation. We are often justified in trusting the system, and the legitimacy of this trust produces presumptively decisive reasons for abidance across a variety of different situations, even in cases where wrongful behaviour is being asked of us (as long as wrongful behaviour is not being asked of us too often). Role-versus-morality conflicts produce a normative (as opposed to merely psychological) tension.

4. Dissolving the Conflict: Fundamental, Imposed, and Derivative Demands

I propose we begin resolving the tension by asking not about any individual role-occupant, but rather about the organisation as a whole. Specifically, we begin by asking: what is the purpose of this organisation? The answer to this question is not determined by the goals that are explicitly endorsed by leaders, executives, directors, managers, or shareholders. Instead, the

[3] Here there are analogies to arguments that there are good reasons to obey the law, even in cases where the law is wrong (e.g. Raz 1986). This analogy becomes even more apt once one sees that citizenship is a role, the demands of which are largely obligations to obey laws. I lack space here to fully defend an analogy between the obligation to obey the state's laws and the obligation to obey an organisation's demands; my suggestion is just that the justification for these obligations take roughly the same shape.

answer is determined by looking at what's essential to this organisation: what purpose is such that the organisation currently has that purpose and the organisation would change its personal identity if it lost that purpose? That is the organisation's fundamental purpose.

To see how this works, consider a legislature. Suppose the fundamental purpose of that legislature is to write laws that reflect the general will of constituents. This is just to say that: if this particular legislature stopped doing this, then there would be a rupture in the identity of the organisation; we'd now be dealing with a new and different entity. Based on this fundamental purpose of the legislature, we can deduce the fundamental purpose of each individual role-occupant in that legislature: very roughly, the fundamental purpose of each individual legislator is to contribute to writing laws, in such a way that those laws reflect the interests or will of that legislator's particular constituents (supposing a system with district-based representation). In this way, the fundamental purpose of the organisation (as determined by its structure) determines the fundamental demands on each individual role-occupant.[4]

But a role's fundamental demands are not directly action-guiding. The action-guiding demands of any role are more specific than the fundamental demands. The action-guiding demands have two faces.

The first face is composed of the demands that an individual confronts as a result of the instructions, commands, hopes, expectations, and so on that others place on her in virtue of her role. These 'others' include role-occupants within the organisation, the organisation itself, and even outsiders who interact with the role-occupant qua role-occupant. Call these the *imposed* demands, since they are imposed by others on a role-occupant, in virtue of her role. In the case of a legislator, one example of an imposed demand is the demand to obey the party's whip in a particular context. The imposed demands constitute the 'role' side of role-versus-morality conflicts.

Yet the imposed demands are only one face of a role's action-guidingness. The second face is composed of the demands that an individual faces in virtue of how her role's fundamental purpose interacts with the context. Call these the *derived* demands, since they are derived from the role's fundamental purpose (in interaction with a particular context). The

[4] One might doubt that any legislature has such a democratic fundamental purpose: after all, isn't the UK Parliament now identical to the English Parliament before it was democratic? This is a question I cannot resolve here; I use the example of a democratic fundamental purpose simply to get across the idea. Plausibly, the precise fundamental purpose of each organisation will be a highly contestable matter.

derived demands might not be imposed on the role-occupant by anyone, not even by herself. No one may even have considered what the derived demands might be. The derived demands are, in this way, demands that can depart from any (explicit or implicit) demands that anyone imposes on the role occupant. In the case of a legislator, one example of a derived demand is the demand to disobey the party's whip in a context where the legislator knows the whip is enforcing a view that is both wrong and contrary to the party's manifesto. The derived demands are action-guiding applications of the fundamental demands in a context.

In light of these distinctions, my proposal for resolving the role-versus-morality tension is simple. When the imposed demands depart from the derived demands (which is to say, when the imposed demands are poor applications of the fundamental demands), then organisational role-occupants have obligations—grounded in the fundamental demands of the role—to use the fundamental demands to challenge the imposed demands. These obligations are themselves derived demands: they are demands that derive from the fundamental demands, in a context in which the role-occupant faces imposed demands go against the proper application of the fundamental demands. Specifically, they are demands to act *because* of the fundamental demands, while *challenging* the imposed demands. By looking to the fundamental demands, we can reconcile the action-guiding demands with morality. We simply need find the *derived* action-guiding demands, rather than focusing only on the *imposed* action-guiding demands.

To see how this plays out, consider the examples from Section 1. The easiest of these is the democratic citizen, whose imposed demands require that she pay taxes that fund a clearly unjust war. Meanwhile, her fundamental demand is to do what she can to contribute to the just governance of her state—by voting, petitioning, protesting, and so on. What about her derived demands? If we were in a context where her state was pursuing justice on all fronts, then her derived demand might simply be to pay her taxes and otherwise obey the law. However, given a context in which her state is pursuing a clearly unjust war, her derived demand is to protest, petition, and so on against the war. This is her *fulfilling* her (fundamental) role obligation; it isn't her resisting her (fundamental) role obligation. Her derived role obligation is to use her fundamental role obligation to challenge her imposed role obligation. Now, she may not be required (by her role's fundamental demands) to withhold taxation: exactly what "challenge" she must make to the imposed demand will vary, depending on the values at stake.

A trickier example is the chair of a condominium owner assembly. This person has the (fundamental) role of organising, chairing, and contributing to meetings of all owners of apartments within a particular building, so that the owners can together deliberate about, and decide on, various matters of building governance. Let's suppose the assembly is debating

whether to dispose of common rubbish in a cheap yet environmentally unsound way, or in a more expensive yet environmentally sound way. Most owners are in favour of the cheap option. It's the chair's turn to speak in the debate. What should she say? This example differs from the citizen example, insofar as the chair's fundamental role makes no mention of moral values, such as justice or environmentalism: her fundamental role does not (it seems) *require* her to contribute to the building being run in an environmentally sustainable way. We might think her fundamental role demand is to contribute to meetings in a way that is *likely to lead to consensus amongst owners*, or in a way that is *likely to make most owners satisfied*. These fundamental role demands don't seem to produce morality-respecting derived demands. However, crucially, when assessing the fundamental purpose of an organisation, we will often need to look at the human purpose for which the organisation was set up.[5] The organisation might have been set up to govern a territory, as in the case of states, or to provide satisfactory financial products and services, as in the case of banks, or to responsibly govern an apartment building, in the case of condominium owner assemblies. Such goals are often an important part of an organisation's fundamental purpose. Once such goals are included in the organisation's fundamental purpose, the chair's derived role obligation will largely be determined by what it means for the assembly to *responsibly* govern an apartment building—where this produces the chair's fundamental and derived role obligations. In a particular context, the assembly's fundamental purpose may well give the chair the derived obligation to speak up in favour of the environmental option for rubbish disposal, on pain of frustrating the assembly's fundamental purpose of governing responsibly.

Another tricky example is a financial advisor in a bank. The advisor's fundamental obligation is to provide sound financial advice to customers. This fundamental obligation is produced by the organisation's fundamental purpose (providing satisfactory financial products and services). The advisor's imposed demand is to convince customers to make high-risk

[5] Often, but not always: sometimes, an organisation's fundamental purpose cannot be ascertained by asking what it was 'set up' to do, for example since the organisation arose in an unintended way out of earlier organisations. It is difficult to say much in general about what determines an organisation's personal identity and, therefore, its fundamental purpose. This difficulty is one attraction of tying an organisation's fundamental purpose to its personal identity: organisations' fundamental purposes are plausibly as varied, and varied for the same reasons and in the same ways, as the determinants of their personal identities. On the latter, see Rust 2019.

investments. In an ideal world, the advisor's derived obligationwould be to give advice that will satisfy the financial goals of the particular customer. This derived obligation is ultimately grounded in the organisation's fundamental purpose of providing *satisfactory* financial products and services. However, in a context with morally problematic imposed demands, the advisor's derived demand is to use her fundamental role to challenge the imposed demands. This example differs from the citizenship and condominium examples, insofar as the advisor might have few avenues for using her fundamental demand to challenge the imposed demand. Yet even here, there are informal avenues for dissent. These include raising the issue with superiors or equals in the structure, or (in more extreme cases) publicly blowing the whistle on the incentive structures that reward giving unhelpful advice. Exactly which actions are required will vary with the values at stake, just as in the other examples.

One objection presents itself immediately. What if the fundamental purpose of an organisation is morally dubious, and therefore produces morally dubious fundamental and derived role demands? Consider venture capitalist organisations. One might be tempted to say that the fundamental purpose of these organisations *just is* maximising profit: this is what they aim to do and what they were set up for. They don't have more substantive or morally respectable goals (such as 'provide satisfactory financial products and services'). In these cases, the fundamental and derived role obligations will trace back to the organisation's fundamental purpose of maximising profit. In some contexts, that purpose will produce a derived demand to act corruptly. And that derived demand can be defeated (one might think) only by a morality that is wholly 'external' to the organisation. The role-versus-morality tension is not dissolved.

An initial reply runs as follows. The fundamental purpose 'making profit' won't produce the dilemma as outlined in Section 3. That dilemma relied on role-bearers having *good reasons* to obey their imposed obligations, as a matter of a general presumption or habit. Only when a role bearer has good reasons for a general presumption in favour of the imposed obligations, does the role bearer face a genuinely normative role-versus-morality dilemma. And perhaps purely profit-driven organisations simply cannot produce good reasons for a general presumption in favour of abiding by imposed obligations.

The initial reply is too quick. The moral legitimacy of the purpose 'make profit' was famously endorsed by Milton Friedman (1970) from within a utilitarian perspective. There is reasonable disagreement over the cogency of Friedman's argument. Thus, it is not unreasonable for at least some role-bearers, in at least some profit-driven organisations, to presumptively abide by their imposed obligations. And such a presumptionmight further be

justified via other moral theories. For example, by adapting the Aristotelean notion of *eudaimonia*, one might think that the truly flourishing venture capitalist corporation just is the one with the highest profits. If profit maximisation is the fundamental purpose of a venture capitalist organisation, and if that purpose might be (reasonably thought to be) morally benign, then a role-bearer might reasonably hold a general presumption in favour of performing their imposed obligations. But that benign organisational purpose might, in certain contexts, nonetheless produce role obligations (both fundamental and derived) that are clearly immoral—because maximising profits will sometimes require role-bearers to act in ways that are clearly immoral.

The worry, then, is that some role-bearers will face a genuinely normative role-versus-morality tension—but their organisation's fundamental purpose will not contain the resources to support the 'morality' side of that tension. Instead, the fundamental demand of any given role-occupant will be 'do what you can to most efficiently contribute to the organisation's making profit.' Meanwhile, the derived demand of any given role-occupant will require the role-occupant to do whatever profit maximisation demands from them in a particular context. In such organisations, the objector might continue, the imposed demands will often perfectly track the derived demands—where the latter are properly understood as demanding whatever *profit maximisation* demands in a context. This is because the imposed demands—as at least when they come from directors, managers, and shareholders—will be in the service of profit. But then (the objector concludes) the idea of an organisation's 'fundamental purpose' won't have captured, or responded to, the role-versus-morality tension. Instead, we will have vindicated the derived role obligations of lying to customers, forging their signatures, twisting their arms, and so on.

In response to this line of reasoning, let's grant for the sake of argument that some organisations have 'maximising profit' as their fundamental purpose (perhaps venture capitalist organisations are the only ones of which this is true).[6] For these organisations, we can—and should—demarcate the conditions under which that fundamental purpose can truly produce the role-versus-morality dilemma.[7] Even for someone like Friedman, the legitimacy of the profit motive does not extend to corruption, lying, forgery, theft, and so on. The organisation's profit motive produces a reasonable presumption (held by role-bearers) in favour of obedience with imposed obligations, *only if* the organisation's profit motive is embedded within a broader

[6] I thank Alex Barber for this example.
[7] I thank Alex Barber and Sean Cordell for help with formulating this point.

system that prohibits the violation of certain basic rights. Because this broader system is a necessary condition for the normativity of the organisation's profit-focused fundamental purpose, we can see the role-bearer's fundamental obligations as *including* the requirements of that broader system. Thus, the role-bearer's fundamental obligations include a prohibition on corruption, lying, forgery, theft, and so on, if those actions are prohibited by the system that gives the organisation's fundamental purpose its normativity-producing legitimacy. Importantly, these prohibitions are still genuinely *role* obligations, because they operate on role bearers only via the *organisation's* fundamental purpose. The constraints apply to the role-bearer insofar as she operates within an organisation with a certain fundamental purpose, where that purpose is legitimate only within certain constraints.

There is another, compatible, method of responding to this objection. This method of response starts by noticing that maximising profit has *rationality* as a necessary condition. You can't aim to maximise profit if you're not (to some extent) rational. By 'rational,' I simply mean that an entity is capable of taking effective means to its ends, capable of ensuring that its beliefs and desires are more-or-less mutually consistent, and that it by-and-large aims to exercise those capabilities.

Once organisations' rationality is on the table, we can notice that organisations (even for-profit corporations) have the ability to use their rationality in a way that respects persons. More generally, corporations have the capacity to respect moral side-constraints on their pursuit of profit. This capacity can be gleaned from the fact that corporations at the very least pay lip service to the ideals of human rights, stakeholders interests, law-abidance, and so on. Of course, corporations might pay such lip service simply out of self-interest—as a way of, ultimately, pursuing profit. But this doesn't matter: the fact that corporations can acknowledge such considerations demonstrates that they are *capable* of attending to them. This capacity produces side-constraints on the pursuit of profit, even if those constraints are not mandated by the social system in which the organisation is embedded.

Paradigmatically, consider the Kantian account of moral agency. As Kendy Hess (2018a) has argued, corporations satisfy three different Kantian conceptions of moral agency: they can "act on universalizable principles and treat humanity as an end in itself," they can "give such principles to *themselves*, treat *their own* 'humanity' as an end itself, and act out of respect for the law," and they can "draw on empathically generated information and insights to inflect their performance." (2018a, 67; see similarly Wringe 2014) The fact that organisations satisfy the conditions of moral agency produces a requirement to respect moral side-constraints—even if the organisation's s purpose just is the completely unbridled pursuit of

profit, and even if it is somehow reasonable for role-bearers to presumptively abide by their imposed obligations.

Thus, even if an organisation has the fundamental purpose of making profit, it *also* faces the fundamental demand to respect moral side-constraints (where that demand is grounded in its capacity to do so). This demand comes from morality, if not from the social system in which the organisation acts. The organisation's fundamental demand to respect side-constraints grounds a related fundamental demand on role-occupants: specifically, role-occupants bear fundamental role obligations to perform their roles in ways that abide by side-constraints. This fundamental demand of the role-occupant (which is derived from the fundamental demand of the organisation) then produces a derived demand to challenge certain imposed demands: that is, the fundamental demand on the role occupant produces a (fundamental) role obligation to alter (imposed) role obligations, in any organisation that is capable of attending to moral considerations. Crucially, I am suggesting that all of these demands on role-occupants are genuinely *role obligations*—they are obligations internal to the organisation. This is because the moral demand to abide by hard constraints applies fundamentally to the organisation itself, and only derivatively to the role-bearers. There is no role-versus-morality conflict, since the role itself requires morality (in all organisations that are moral agents).

Now, arguably, not all organisations are moral agents: not all organisations are capable of attending to moral considerations. Some organisations might be 'amoral,' that is, instrumentally rational but incapable of any kind of moral reasoning (Hindriks 2018). In these organisations, there is no way of using the organisation's fundamental purpose or fundamental obligations to resolve role-versus-morality conflicts. This is because these organisations' fundamental purposes (and the role-occupants' fundamental demands, which derive from that purpose) may produce derived demands that conflict with morality, while lacking the resources to produce any derived demands that accord with morality. Perhaps some criminal organisations are like this.

In such organisations, my proposal for resolving role-versus-morality conflicts will not work. But this doesn't mean that role-occupants should simply obey such organisations' edicts. Instead, if an organisation is not a moral agent—and therefore is genuinely *incapable* of producing (fundamental and derived) role demands that accord with morality—then we should doubt whether role-occupants have any reason to perform their role at all. Of course, the organisation's goals produce a kind of normativity that's internal to the organisation's perspective. But that normativity doesn't produce genuine obligations for members. In these

cases, then, the role-versus-morality conflict should be resolved in favour of morality; the role should not be obeyed. This is different from the cases I focused on above, in which the organisation has a morally benign or morally conditioned fundamental purpose (or faces the fundamental demand to be moral), such that the fundamental role obligation accords with morality, such that the role-versus-morality conflict dissolves altogether.

5. Conclusion

I have suggested that we sometimes have 'role obligations to alter role obligations': obligations to use our role's fundamental demands as a justification for acting in accordance with derived demands, where the derived demands require us *to push back against* imposed demands. My proposal can be seen as a variation of Michael O. Hardimon's claim that "[t]here are different ways of being a flight attendant, oncologist, or police officer. Part of what it is to become a good flight attendant, oncologist, or police officer is to find a way of carrying out the responsibilities of these roles which suits one's particularities." (1994, 355) I would contest only the last part of Hardimon's statement. My suggestion has been that the good flight attendant, oncologist, or police officer finds ways of carrying out the responsibilities that accord with that role's place in its organisation's pursuit of its fundamental purpose—where that purpose will almost inevitably be infused with moral requirements.

To be sure, the derived demands that I have proposed—role obligations to alter role obligations—are not the only way to push back against morally problematic imposed demands. A role-occupant might also (or instead) attempt to act upon her organisation 'from the outside,' to induce it to change its ways. Such actions are not properly construed as *role* performances, so any obligations to perform such actions are not *role* obligations. I have focused here on the ways in which our organisationally-embedded roles *themselves* can contain the normative resources for pushing back in those cases where our organisations demand the morally dubious. By attending to the fundamental purpose of our organisations—which include the conditions for its moral legitimacy and the fact of its moral agency—we can infer the fundamental demand of our role. This fundamental demand often (though not always) contains intra-organisational normative resources for challenging imposed demands that conflict with moral requirements.

References

Bicchieri, Christina and Ryan Muldoon. 2014. "Social Norms." *Stanford Encyclopedia of Philosophy*.

Collins, Stephanie. 2017. “Duties of Group Agents and Group Members.” *Journal of Social Philosophy* 48/1: 38–57.

Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford: Oxford University Press.

Collins, Stephanie. ms. *Organisations as Wrongdoers*.

Copp, David. 2007. “The Collective Moral Autonomy Thesis.” *Journal of Social Philosophy* 38(3), 369–388.

French, Peter. 1984. *Collective and Corporate Responsibility*. Columbia University Press.

Friedman, Milton. 1970. “The Moral Responsibility of Business is to Increase Its Profits.” *New York Times Magazine.*

Hardimon, Michael O. 1994. “Role Obligations.” *Journal of Philosophy* 91(7), 333-363.

Harre, Rom. 1979. *Being Social*. Oxford: Blackwell.

Herzog, Lisa. 2018. *Reclaiming the System: Moral Responsibility, Divided Labour, and the Role of Organisations in Society.* Oxford University Press.

Hess, Kendy. 2014. “The Free Will of Corporations (and other Collectives).” *Philosophical Studies* 168(1), 241-260.

Hess, Kendy M. 2018a. “Does The Machine Need a Ghost? Corporate Agents as Nonconscious Kantian Moral Agents.” *Journal of the American Philosophical Association* 4(1), 67-86.

Hess, Kendy M. 2018b. ‘The Peculiar Unity of Corporate Agents.’ Pp. 35–60 in Kendy M. Hess, Violetta Igneski, and Tracy Isaacs (eds), *Collectivity: Ontology, Ethics, and Social Justice*. Rowman and Littlefield.

Hindriks, Frank. 2018. “Collective Agency: Moral and Amoral.” *Dialectica* 72(1), 2-23.

Hodgson, Geoffrey M. 2007. ‘Institutions and Individuals: Interaction and Evolution’, *Organsiation Studies* 28, 97-116.

Koslicki, Kathrin. 2008. *The Structure of Objects*. Oxford: Oxford University Press.

List, Christian and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

Pettit, Philip. 2007. “Responsibility Incorporated.” *Ethics* 117, 141-201.

Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.

Ritchie, Katherine. 2013. “What Are Groups?” *Philosophical Studies* 166(2), 157-72.

Rovane, Carol. 1998. *Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton: Princeton University Press.

Royal Commission into Misconduct in the Banking, Superannuation and Financial Services Industry. 2018. *Interim Report*. Volume 1. Available at https://www.royalcommission.gov.au/royal-commission-misconduct-banking-superannuation-and-financial-services-industry

Royal Commission into Misconduct in the Banking, Superannuation and Financial Services Industry. 2019. *Final Report*. Volume 1. Available at https://www.royalcommission.gov.au/royal-commission-misconduct-banking-superannuation-and-financial-services-industry

Rust, Joshua. 2019. "Institutional Identity." *Journal of Social Ontology* 5(1), 13-34.

Thomasson, Amie . 2019. "The Ontology of Social Groups." *Synthese* 196, 4829-4845.

Tollefsen, Deborah. 2015. *Groups as Agents*. Polity.

Tuomela, Raimo. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.

Turner, Jonathan. 1997. *The Institutional Order*. London: Longman.

Weber, Max. 1968. *Economy and Society: An Outline of Interpretive Sociology*. Ed. Guenther Roth. New York: Bedminster Press.

Wringe, Bill. 2014. "May I Treat a Collective as a Mere Means?" *American Philosophical Quarterly* 51(3), 273-284.