Article

David Strohmaier*

Organisations as Computing Systems

https://doi.org/10.1515/jso-2020-0052
Published online March 1, 2021

Abstract: Organisations are computing systems. The university's sports centre is a computing system for managing sports teams and facilities. The tenure committee is a computing system for assigning tenure status. Despite an increasing number of publications in group ontology, the computational nature of organisations has not been recognised. The present paper is the first in this debate to propose a theory of organisations as groups structured for computing. I begin by describing the current situation in group ontology and by spelling out the thesis in more detail. I then present the example of a sports centre to illustrate why one might intuitively think of organisations as computing systems. To substantiate the thesis, I introduce Piccinini's restrictive analysis of physical computation. As I show, organisations meet all criteria for being computing systems. Organisations are structured groups with the function of manipulating medium-independent vehicles according to rules. Furthermore, I argue for the modal claim that this is a necessary feature of organisations. Having sketched the computational account of organisations, I compare it to other proposals in the literature.

Keywords: social ontology, group agency, philosophy of the social sciences

1 Introduction

Organisations are computing systems. This thesis covers diverse organisations, such as sports centres, committees and governmental authorities. A university's sports centre is a computing system for managing sports teams and facilities. A tenure committee is a computing system for assigning tenure status. The national passport centre is a computing system for processing passport applications. Despite an increasing number of publications in group ontology (Effingham 2010; Epstein 2015, 2019; Hawley 2017; Ritchie 2013, 2015, 2020; Sheehy 2003, 2006a, 2006b;

***Corresponding author: David Strohmaier**, Department of Computer Science and Technology, Cambridge University, Cambridge, Cambs., United Kingdom of Great Britain and Northern Ireland, E-mail: ds858@cam.ac.uk. https://orcid.org/0000-0002-1430-8212

Uzquiano 2004, 2018), the computational nature of organisations has not been recognised. The present paper is the first in this debate to propose a theory of organisations as groups structured for computing.

I begin by describing the current situation in group ontology and by spelling out the thesis in more detail. I then present the example of a sports centre to illustrate why one might intuitively think of organisations as computing systems. To substantiate the thesis, I introduce Piccinini's restrictive analysis of physical computation. I show that organisations meet all the criteria for being computing systems. Organisations are structured groups with the function of manipulating medium-independent vehicles according to rules. Furthermore, I argue for the modal claim that this is a necessary feature of organisations. Having sketched out the computational account of organisations, I compare it to two other proposals in the literature: theories of extended cognition and theories of group agency. Both treat groups as computing systems under some interpretation, but neither entails the specific theory of organisations I propose.

2 The Place of Organisations in Group Ontology

Over the last two decades, group ontology has turned into a blossoming field. There have been extended debates about the metaphysical analysis of groups, the two main approaches being theories of constitution and neo-Aristotelian mereology.[1] According to theories of constitution, groups are constituted by their members (Epstein 2015, 2019; Uzquiano 2004). The accounts using neo-Aristotelian mereology instead analyse groups as structured wholes, with members as parts of he group (Hawley 2017; Ritchie 2013, 2015, 2020; Sheehy 2003, 2006a, 2006b; Strohmaier 2018). While the constitutive proposal turns groups into a *sui generis* kind, the mereological account offers a reduction. Groups are wholes composed of parts in a certain structure, just like chairs, cars and laptops. Although both proposals are compatible with the purpose of the present paper, I will assume the mereological account.

Recently, group ontology has started to sort groups into different kinds. Ritchie (2013, 2015, 2020) proposes a very broad distinction between structured groups, such as committees, and feature groups, such as genders. Epstein (2019) questions whether any such broad typology of groups is possible, and instead has proposed a scheme for describing kinds of groups more finely. Rather than distinguishing structured and feature groups, he individuates kinds of groups on the level of social classes or even *Tufts University faculty committee*. At least one of

1 Other options are a set-theoretic analysis (Effingham 2010) and an analysis of groups as pluralities (Uzquiano 2018). Uzquiano's account of groups as pluralities is in fact very closely related to neo-Aristotelian mereology.

Epstein's motivations for abandoning Ritchie's coarse-grained distinction is that he considers it unlikely that this level of grain enables interesting generalisations. He argues that some generalisations suggested by Ritchie, such as the claim that structured groups always have unified functions, cannot be maintained on this level of grain. By moving to a finer level, Epstein hopes to account for the heterogeneity of the social and enable successful generalisations.

Since I offer a partial analysis of organisations as a kind of group, I contribute to the effort of creating a typology of groups. For at least two reasons, organisations – including committees, bureaucratic centres and governmental bodies – are of special interest for this effort. First, if my proposal succeeds, it describes an important kind of group on a coarser level than most of Epstein's work. The kind *Tufts University faculty committee* is a sub-kind of organisations. Organisations are a frequent and cross-cultural kind of group, for which I will establish one general and explanatorily powerful theory. While the proposal does not lay claim to providing a complete typology of the social, it establishes a further example of broad kind of groups in addition to social classes.

Second, organisations are particularly conspicuous groups. Organisation studies is an interdisciplinary field in its own right with considerable practical influence, ranging from managerial approaches to theories of corruption.[2] Furthermore, organisations have been frequently used for illustration in group ontology debates and they provide an intriguing contrast to race and gender groups, which have received much of the remaining attention. Committees are ubiquitous in the literature, not just in Epstein's paper. Given the extensive role organisations play, providing a theory of this type of group is a worthy effort.

Despite their frequent use as examples, no analysis of organisations has prevailed in the literature. As a first pass, one might be tempted to identify them with structured groups. The philosophy department is an organisation and it has a structure. But although organisations are groups with a structure, the identification would be a mistake because not all structured groups are organisations. While schools and sports centres are organisations, groups of friends and families are not. Nonetheless, both a group of friends and a family are structured groups. Just like the school and the sports centre, a group of friends has structuring relations holding between the members, such as two members standing in the relation of declared best friends forever or one member relating to the others as advice provider.[3] Likewise,

---

2 For sources on the latter, see Vaughan 1999 and Ashforth and Anand 2003.

3 One might raise the question what makes a relation a relevantly structuring relation. Why is the "being of the same gender" relation not relevant and does not render gender groups a structured group? The onus lies here on those defending structured groups as a kind of group, as Epstein (2019) points out. The issue does not affect my proposal.

the parents in the family bear a specific relation to the children. Organisations are a proper subset of structured groups. What, then, makes organisations special, if anything?

I propose a computational theory of organisations with a modal aspect. Organisations are groups which necessarily have their structure for computing. There is no organisation and there cannot be an organisation which does not have the function to compute, where I use "function" in a broadly teleological sense. An organisation is supposed to be arranged to fulfil its computational purpose, just as a laptop is. Both are mereological objects designed for computing, and created with appropriate structures if things go right.[4]

The claim can be formally expressed as $\forall x(\text{organisation}(x) \rightarrow \Box(\text{group}(x) \& \text{structured-for-computing}(x)))$. As can be seen, I endorse only a one-sided conditional, rendering the computational nature a necessary but insufficient feature for being an organisation. This humility is in place, because social groups are a challenging field for metaphysicians which rarely allows for simple and precise analyses. With an eye to the social sciences (e.g. Preisendörfer 2011), I expect there to be further necessary features, candidates for which are depersonalization of group roles, a hierarchy, or shared resources.

While my computational proposal has no defenders in the contemporary group ontology debate, it has predecessors in the social sciences and philosophy.[5] I will treat the social sciences here and discuss the more subtle differences to philosophical contributions after establishing my own analysis.

Organisations have long been described as goal-directed and rule-governed systems in the social sciences (Scott and Davis 2006; Weber 2013 [1921–1922]) and the image of organisations as machines has exerted great influence (Morgan 2006). Probably the most overtly computational approach to organisations was put forward by the Carnegie School of organisational behaviour. Unsurprisingly given Herbert Simon's involvement, it drew explicitly on computer science, for example by employing the notion of a programme (March and Simon 1958) to describe organisational behaviour.

More recently, authors in the computational social sciences (Cioffi-Revilla 2017) have described organisations as complex adaptive systems from a

4 This structure, however, can change and malfunction. Likewise, the structure of a laptop can change, for example if someone installs more RAM, and malfunction, for example if the RAM was inserted incorrectly. The structured-for-computing predicate entails that the group has some structure and that the structure has the function to compute, not that the group does so correctly. The group is no different from a laptop in this regard.

5 I thank Fabio Rojas for pointers regarding the social science literature.

computational perspective. Practitioners in this field use concepts and methods from computer science to understand social organisations (e.g. Miller and Page 2007, chapter 11). These efforts suggest that there is explanatory value to a computational theory of organisations.

The proposed account offers three benefits in comparison to previous computational theory of organisations in the social sciences. First, it improves the foundations of conceiving of organisations as computing system by employing a recent theory of physical computation. Current contributions to computational science often employ a very broad notion of information processing (e.g. Cioffi-Revilla 2017), that makes it hard to discern when something is described as a computer and when only modelled by one. By introducing recent theories of physical computation, I am able to strengthen a position that introduces computational explanations to an influential field of the social sciences.

Second, the present paper is the first to describe organisations as a specific type of *group* distinguished from other types, such as families and gender groups, by their necessarily computational nature. The proposal connects research from the social sciences to the philosophical group ontology debate about the nature and typology of groups.

Third, the modal aspect distinguishes the proposal from previous descriptions of organisations. An organisation is a group which necessarily has a structure for computing. The proposal requires that the group's structure be for computing across all times and possible worlds. The motivation for this modal move is to explain an intuitive difference between organisations and groups which contingently engage in computations, as I explore in depth later.

Organisations share the modal character with natural kinds such as water.[6] As in the case of water being H_2O, we were not aware of the correct modal analysis at the time of inventing the term "organisation", but in retrospect can see what our term has latched onto. That organisations are groups structured as computing systems is an aposteriori necessary truth (cf. Kripke 1980). In the following, I will discuss the example of an organisation in depth, then apply an analysis of computing systems to organisations, and defend the modal claim. Before concluding, I will compare my proposal with others made in philosophy and the cognitive sciences.

6 I leave here open whether organisations form themselves a natural kind. On the question whether social kinds can be natural kinds, see Khalidi 2013, 2015 and Mason 2016.

3 The Example of the Sports Centre

The sports centre of a university is an innocuous example of an organisation. Like other organisations, it has a function, in its case registering sports teams.[7] To fulfil its purpose the centre will follow a process according to rules. A closer look at this process reveals the nature of organisations.

In the days of pure paper bureaucracy, the following steps established either an intramural basketball or water polo team at the university. A student needs to request a registration form from the sports centre, which has to be returned within a week. If the week passes without the form being returned, the registration process is aborted. If the form is returned in time, the sports centre then expects the student to attend an introduction to the facilities. There is one court and one pool and depending on the sport of the team, basketball or water polo, the student has to visit one or the other facility. At the introduction event, the applicant receives a signed sheet which they have to bring back to the sports centre. Upon receipt, the centre will finally approve the application and the team is thereby registered.

Moving to a higher level of abstraction, the process can be described as the centre taking a number of states in response to various inputs. There is a start state (q_O), from which the centre is moved by requesting a form to the state (q_1) of expecting the completed form or the passage of a week. If a week passes without the completed form being received, the centre returns to state q_O. Otherwise, it advances either to state (q_2) of expecting the sheet certifying that the introduction to the pool has been attended or the state (q_3) of expecting the sheet for the court. Finally, if those sheets are returned, it reaches the state (q_4) of having registered the team.

This type of description should be familiar to those having studied computational systems. The sports centre implements a finite state machine for registering a team. There is an input alphabet $\mathbf{\Sigma}$, namely the various forms and pieces of paper; a set of states S in which the centre can find itself, including the start state q_O; a transition function δ from state to state based on input, i.e. δ: $\mathbf{\Sigma} \times S \rightarrow S$; and there are two final states in my description, the initial state (q_O) and the acceptance state (q_4) after all documents have been received in order.

The finite state machine allows for a graphical depiction according to the conventions of computer science (see Figure 1).[8]

7 Presumably, the sports centre also has other functions. The multiplicity of organisational functions will be addressed in due course. Although substantially changed, perhaps beyond recognition, and used for an entirely different purpose, the example of the registration process is inspired by one found in Epstein (2015: 184).

8 I left out the input to keep the image clean.

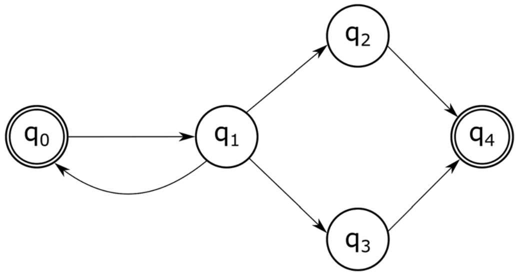

Figure 1: Representation of the sports centre as a finite state machine.

There are, of course, several shortcomings with the description. The sports centre in my example appears to be a finite state transducer, taking inputs and giving outputs, rather than a finite state machine which takes only inputs; finite state machines have no memory, but the best theory of the sports centre attributes some sort of memory to it; a real-world sports centre does not just take input from one person at a time, instead multiple registration processes will overlap; a real-world sports centre can get in contact with other people and institutions, for example, to check with a department whether a student is really registered there.

None of these worries touch upon the core of my position, because they only concern the specific architecture. I am not arguing that organisations like the sports centre are specifically finite state machines. That would be true for only a few organisations designed by overzealous computer scientists. I argue that organisations such as the sports centre are computational systems. They reach goal states by engaging in computational processes for which they have been designed.

4 Organisations as Physical Computing Systems

The sports centre is a motivating example of why one might think that organisations are computing systems, but turning the intuition into a substantial thesis requires an account of physical computation. The literature offers multiple theories which differ in their criteria for being a computing system. The theories range from pancomputationalism (Chalmers 1996; Putnam 1988; Searle 1992), which attributes computation to every physical object, to restrictive proposals requiring the systems to have very specific functions.

My proposed theory of organisations is in tension with pancomputationalism. If everything was a computer, then my claim that organisations are computating systems would become trivial and of little explanatory value. Since establishing a particular analysis of computing systems is beyond the scope of the present paper, my strategy will be to revise common arguments why pancomputationalism fails and then take the most developed restrictive analysis of computing systems and discuss whether organisations fit it.

4.1 Pancomputationalism

Pancomputationalism comes in an unlimited and a limited form. According to unlimited pancomputationalism, not only is every system a computer, but every system computes everything. Famously, Hilary Putnam (1988) tried to undermine his own previous functionalisms about the mind by showing that the states of a physical system can be mapped to all possible computational states. On the assumption that the availability of such mappings establishes that the objects are equivalent as computing systems, all physical objects engage in all computational processes.

As Piccinini (2015: 53–54) has forcefully argued, unlimited pancomputationalism can neither account for our everyday understanding of computation, nor for the practices of computer scientists. The common user is not willing to replace their laptop, or for that matter its CPU, with a rock. In everyday conversation, only certain objects are considered computing devices and they are considered to run specific processes.

Similarly, computer scientists study very particular artefacts. Most of them investigate digital computers, a smaller number of them analogue and quantum computers, but none of them study the computational processes of rocks. Unlimited pancomputationalism does neither track what ordinary language users, nor the relevant scientists mean by “computational system” and I will, thus, leave it aside.

Limited forms of pancomputationalism (e.g. Chalmer 1996: 331) fare somewhat better in capturing the meaning of “computing”. While they also treat every physical object as a computing system, each system is only assumed to perform a limited set of computations. Consequently, there is a difference in the computational processes in which they engage. With this change the unwillingness of laptop users to switch to rocks is explained. Whatever rocks compute, it is not what a laptop computes. Similarly, computer scientists might not study all kinds of systems, because only a few engage in interesting computations. By counting planet systems and the weather as computers, limited pancomputationalism still stretches the notion of computation beyond the limits of common sense, but some might find the extent to which it does so palatable.

Limited computationalism, however, introduces a distinction between computation systems which engage in explanatorily interesting computations and those that do not. This distinction would allow to recapture the main thrust of my argument. Instead of ascribing some unspecific computations to organisations, the claim is that organisations necessarily have their structure for fulfilling a function by way of interesting computations. The processes through which organisations achieve their goals are computational processes that are relevantly taken as such. In the example of the sports centre, the registering of a team is a computational process.

That being said, even limited pancomputationalism is a rather extreme position, given our ordinary intuitions about computing systems and the scientific notion employed in computer science. In the following, I will argue on the basis of a more restricted analysis.

4.2 Piccinini's Restrictive Analysis

In the domain of current analyses, Gualtiero Piccinini's proposal is amongst the most restrictive and most developed.[9] According to him, "[a] physical computing system is a mechanism whose teleological function is performing a physical computation" (Piccinini 2015: 10; see also Coelho Mollo 2018, 2019). A laptop is for computing and the same is true for an organisation. Of course, that moves the question only a step further, raising the questions what distinguishes computation from other processes.

Piccinini understands physical computation as "the manipulation [...] of a medium-independent vehicle according to a rule" (ibid.). A vehicle is a physical variable, which in the case of digital computation might for example consist in a string of digits.[10] The states of a Turing machine's tape, for example showing a 1 or a 0, are vehicles; the tape itself is a medium.

The analysis establishes three restrictions. Organisations must have (1) teleological functions for (2) rule-governed manipulation of (3) medium independent-

9 Two other prominent restrictive accounts (Fodor 1975, 1981) analyse computation as involving the manipulation of meaningful symbols according to their semantic or syntactic properties. Computational processes would then either be individuated by semantic or syntactic properties. Piccinini (2015: 26–50) argues against both accounts, but it would not undermine my proposal if either of them were correct. Organisations would fulfil these other conditions as well, since they operate on natural language or closely related types of representation.

10 The term "digit" is here not intended in the sense of "numeral", but rather it refers to "a portion of a vehicle that can take one out of finitely many states during a finite time interval" (Piccinini 2015: 121). A ticked or unticked box on a form would be such a digit. In the case of analogue computing, there are an infinite number of states.

In this section, Piccinini also notes that his use of "vehicle" is in fact ambiguous. "Vehicle" can denote either a state that can take different values or such a value. The context disambiguates the term sufficiently.

vehicles. I will go through each of these three requirements in turn, discussing the case of the sports centre as well as organisations in general.

4.3 Teleological Functions

Organisations such as the sports centre from our example have teleological functions, such as registering sports teams. In this respect organisations do not differ from other artefacts. Just like a laptop is designed for computation (cf. Piccinini 2015: 111–112; Coelho Mollo 2019: 445–446), so an organisation is designed for its function.[11] One might question, however, whether they have the right kind of function. One might suggest that the purpose of the centre is not literally to compute, but to register sports teams, while Piccinini requires that the function to be "performing a physical computation" (ibid.). Thus, one might question whether the function has the correct target.

It is difficult to deny that many organisations are not pure computing systems that are systems designed exclusively for the purpose of computing. The criticism is nonetheless mistaken, because it assumes that a computing system has only one unique function. Computing systems need to have the function to compute, but nothing stops them from having further functions. Many paradigmatic computing systems have functions in addition to pure computation. Piccinini's allows such further functions, carefully stating that "[o]ne of" (2015: 121) the system's functions has to be to compute. These functions can depend on the computing, as they do in the case of organisations like the sports centre, but they can also be unrelated. A server that is also a doorstopper remains a computing system, despite also having the function of keeping the door in place.

A smartphone is a computing device even if its purpose is to communicate and not just compute the first 5000 digits of pi by manipulating medium-independent vehicles. The function of the smartphone is to serve communication by computing.

Robots offer another example of computing systems that have functions beyond computing.[12] The purpose of a Roomba is not just to compute a representation of the

---

11 While I focus on organisations which are artefacts as the core case, I do not rule out that there might be extraordinary cases in which organisations have functions not grounded in design, but in virtue of the same phenomena that ground the teleological functions of the human brain. If one believes that evolutionary selective processes or dispositions to contribute to goals bestow teleological functions (cf. Coelho Mollo 2019: 439–440), then I do not rule out that in some possible world an organisation could have its function in virtue of these phenomena rather than design. I commit, however, to the claim that all organisations will have appropriate teleological functions.

12 A more controversial further example would be human agents. I also consider these to be computing systems, although they have many other functions. Human agents, however, differ from robots and organisations by not being artefacts.

house, but to move through it and clean it. Nonetheless, one function of the Roomba is to compute where to move. It is a computing system with additional functions regarding the world. Likewise, I propose that the function of the sports centre is to serve the registration of teams by computing.

The opposite worry would be that some organisations lack functions altogether. Devising such a counterexample, one has to keep in mind that many organisations which seem to *lack* a function are just *mal*functioning. Consider a bureaucratic organisation which has been founded for the purpose of fighting poverty, but by now does little else than keep its members employed. Although it does not fulfil its function, it surely is still an organisation. This example does not undermine the claim that organisations are computing systems any more than the failure of a misconfigured laptop to start undermines the claim that laptops are computing systems. A malfunctioning computer is still a computer.

A more challenging example of a group seemingly lacking a function, can be found in John Updike's short story "Minutes of the Last Meeting", originally discussed by Margaret Gilbert (1992: 172). This text is written as the record of a meeting in which the chairman of a committee asks for the purpose of the organisation, only to be frustrated:

> "The Secretary pointed out that the bylaws perfectly clearly specify the purpose of the Committee and read excerpts spelling out that 'no political candidates or partisan causes should be publicly espoused,' 'no stocks or bonds were to be held with the objective of financial profit or gain,' and 'no gambling or licentious assignation would be permitted on any premises leased or owned entirely or in part by the said Committee.'
>
> Mr. Langbehn asked to see the bylaws.
>
> The Secretary graciously complied.
>
> Mr. Langbehn claimed after examination that this was a standard form purchasable in any office-supplies or stationery store.
>
> Mrs. Hepple said she didn't see that it made any difference, that here we all are and that is the main point." (Updike 2003: 637–638)

Despite the lack of a purpose, the meeting continues and the organisation persists. Here, one might suggest, is a literary example of an organisation without any function whatsoever.

Quite to the contrary, the story offers a purer example of a computing device than the sports centre. As noted, the centre has a more specific function than just to compute, i.e. it is supposed to register teams. The committee of Updike's story lacks any such *specific* purpose, but its function is still to compute. It resembles a general-purpose computing device, say a CPU that could be integrated in many specialised systems.[13] The committee has rules which are continuously cited and applied, even though the larger goals are not settled.

One last example would be a phantom company created for tax purposes consisting of nothing more than a letter box. It processes nothing and was never supposed to process anything. If it has a function, it is to mislead tax authorities or exploit a legal loophole, not to compute anything or accomplish anything by computing.[14] At this point, I have to draw the line and claim that there is no organisation, but only the appearance of one. In the same way a corporation can be a legal person without being a person in the core sense of the term, a corporation can be treated as an organisation for legal purposes without being one. All organisations have functions as required for being computing systems.

4.4 Rules

Having a function does not suffice. The second required feature is that organisations are supposed to behave according to rules. The informal dynamics of many organisations pose a difficulty for this requirement. While such organisations have functions, they might seem not be appropriately rule-governed for being computing systems. To those familiar with organisation theory, such criticism will have a familiar ring. In the social sciences similar debates have taken place, with some proposing a rule-governed functioning of organisations and others attacking these theories.

The sociologist most closely associated with a picture of organisations as bureaucratic rule-governed and hierarchically structured groups is Max Weber (2013 [1921–1922]). Today, defenders of a rational system view (see Scott and Davis 2006) carry the torch of organisations as systematically arranged and rule-governed entities. Over the decades a multitude of criticisms has been raised against such theories in the social sciences. Many of the objections, however, do not undermine the computational nature of organisations generally.

13 One might suggest that an organisation is in fact closer to a *field-programmable gate array* (FPGA), which are circuits that can be reconfigured after production. Similarly, the rules and procedures of organisations are typically subject to reconfiguration. I thank Stephanie Collins for pressing me on this issue.

14 This counterexample could be addressed in another way by arguing that the function presumed by law trumps the function intended by the person registering the company for tax purposes. In this case, the phantom company would have the appropriate function for being a computing system, albeit it was intended to malfunction consistently by its owner.

For example, critics in the tradition of symbolic interactionism (Strauss 1988; Strauss et al. 1963) have pointed out that organisations of rely on ad-hoc and informal rules, rather than the bureaucratic formalisation and explicit hierarchies that Weber stressed. But while the heterogeneity and flexibility of such rules may be unmatched by the computing devices on our desks, that does not undermine the computational nature of the system.

Transfer of the criticism from the social sciences fails, because the notion of rules in Piccinini's analysis is very thin. According to him, a rule is merely "a map from inputs (and possibly internal states) to outputs" (Piccinini 2015: 121). Having rules for requests and replies, say that one has to request a departmental reimbursement form from a secretary to initiate the reimbursement process, is all that is needed. All organisations will have some such rules. Neither do the rules of computing systems have to be deterministic in the sense that for each computational state of the machine one input will always produce the same output.[15] Probabilistic automatons have the same right to the label "computing system" as deterministic finite state machines. Neither the indeterminism, nor the implicit nature of the party committee's rules keeps it from being a computing system.

It is also not a problem that the rules are often interpreted by the individual members of a group. An organisation is a special computing system, because it is (partially) composed of members who are themselves full-fledged agents. This feature affects the rule-application as the members will associate an intention with the rule. For example, an employee of the sports centre might correct an error in a form, because they understand what has gone wrong. Such interpretations might render the groups more robust or might on occasion undermine their functioning, for example because members wilfully misconstrue rules, but they do not affect that organisations are supposed to implement input-output rules.

A different challenge is posed by the question whether it is the organisation which follows the rules or the individual members of the organisation. After all, are the rules laid out by organisations not applied to individuals? It is correct that there are norms governing in individuals in an organisation, but these are not the rules in question when analysing organisations as computing systems. The rules in question are those of manipulating medium-independent vehicles. That these are rules of the organisation rather than rules of the individual can be seen in that they persist through changes of membership. The rules of handling sports team registration forms continue to be the same no matter which individual helps implement them in the sports centre.

15 In this case, Piccinini's mapping has to be adapted, so that the input is mapped to a weighted set of potential output options.

4.5 Manipulation of Medium-Independent Vehicles

The third requirement of Piccinini's analysis is that the rules concern the manipulation of medium-independent vehicles, where medium-independence is analysed as the rules being insensitive to the physical implementation of the vehicle (cf. Piccinini 2015: 122).[16] Generally, this part of the analysis distinguishes a laptop from a combustion engine. A combustion engine also has a function and an input-output rule, but the function is not to manipulate medium-independent vehicles, it is to convert the fuel's stored energy into movement. By contrast, the laptop has the function to manipulate the portions of a physical medium, where the medium itself is immaterial for the rule. The question thus becomes whether organisations always have the function to manipulate appropriate vehicles.

The forms of the sports centre provide a clear case of medium-independent vehicles which fit the requirement. The state of a box as ticked leads the organisation to expect a specific sheet, in the end resulting in a water polo rather than a basketball team being added to the official records. The ticked box and entry in the records are states of vehicles independent of the medium. For the purposes of the centre, the form could be on pure paper, digital, or hewn in stone. According to Piccinini and Scarantino, "the rules (i.e. the input–output maps) that define a computation are sensitive only to differences between portions (i.e. spatiotemporal parts) of the vehicles along specific dimensions of variation" (Piccinini and Scarantino 2011: 8). Being ticked or not ticked is the specific dimension of variation in the case of the form and the organisation's behaviour is only specific to it, not whether the ticking happened with a pencil or a fountain pen.

Of course, the sports centre is a cherry-picked example to illustrate the main thesis of organisations as computing systems. Forms are widespread throughout many organisations, but they are not a necessary feature of organisations. But forms are also not a necessary feature of computations by groups. While forms are a particularly striking exemplars of medium-independent vehicles, there is no need for the vehicles to take such a strongly bureaucratic shape.

In many cases of organisational computation, the vehicles will be realised as written words and utterances exchanged between the group members.[17] The natural language strings are medium-independent vehicles. Since these strings can be manifested in text messages, handwritten notes, or whispering, their

16 Many vehicles are symbols in the sense of standing in a representational relation, e.g. word tokens that refer are vehicles, but such a representational character is no requirement for Piccinini's theory of computation.

17 For a related discussion of language as a medium of computation in organisations, see Gordon and Theiner 2015. They do not focus, however, on the organisational rules.

independence of a particular medium is clear. The organisational rules will be specific to the language string rather than the physical manifestations.

The rule in a firm might specify that the accounting department has to provide a report on finances upon request, but the request might be given verbally or via email and the report can be printed on paper or other materials. In instances like this, organisations rely on the semantic dimension of the vehicles and on the interpretation of this dimension by their members. Again, typical organisations differ from many other computing system in virtue of their strong reliance on the capacities of their members as agents. Nonetheless, the vehicles that are manipulated are medium-independent and that is what is required.

Piccinini's third requirement – that the rules concern the manipulation of medium-independent vehicles – will be fulfilled for all organisations, because they are built upon the manipulation of either natural language strings or similar representations. Boxes to tick on a form are one example of representations that an organisation might use in place of natural language strings. One might at least imagine an organisation relying exclusively on such representations and avoiding natural language strings altogether, but no organisation can exist that is not at least supposed to manipulate appropriate vehicles under some circumstances. All organisations have the function to manipulate medium-independent vehicles according to rules. The sports centre, the tenure committee, and Microsoft are all groups which necessarily have a structure for computing. All organisations are covered by the analysis.

5 The Modal Aspect

I have argued that all organisations are computing systems, but so far I have not defended the modal aspect of my proposal, i.e. that organisations are necessarily structured for computing. The reason to endorse the modal aspect is that there exists an intuitive difference between organisations and other groups that implement a computation.

Previously, I distinguished organisations from other structured groups, such as families. One might suggest, however, that under appropriate circumstances a family could contingently be a computing system. Pancomputationalism makes this claim trivially true, which is why the computational theory of organisations is committed to rejecting it, but even if one instead endorses Piccinini's account, a family can appear to implement a computing system.

Assume that a family is playing a board game which requires them to act like a finite state machine. They play one figure on the board, roll a 5, and move accordingly. The field on which they end contains a string which is the input and the rule requires them to produce a certain output based on the string. For example, the string might be "horse" and the family has to produce "esroh" by interacting with each other according to some rules implementing a simple reversal algorithm. The rule book specifies the input to output mapping. The game reaches a goal state when they land on a particular field.

During the game, the family would seem to also be a computing system, but intuitively there still exists a difference to organisations such as the sports centre. The computation appears accidental to the family in a way it is not for organisations. The family can stop playing the board game and thereby stop being a computing system while remaining a family. Manipulating vehicles according to input-output rules appears central to being an organisation in a way it is not to being a family or a group of a friends. To explain the intuitive difference between the family and organisations, I propose the modal description. Organisations are necessarily computing systems while families are not.

In response, one might question whether an organisation cannot stop to be a computer as well. Surely, when all the members of the committee call it a day and head home, the computing of the organisation stops. This response misses its target because it confuses the predicates "computing" and "being a computing system". The laptops we use are computing systems even when they are completely turned off. Organisations, I propose, are more similar to laptops than families in this regard.

A laptop stops to exist when it stops being a computing system, and the same is true for an organisation. A philosophy department can change its rules, it can take a break, but it cannot stop having the purpose of processing medium-independent vehicles at pain of ceasing to be. Assume a philosophy department was shamefully axed by its university and that the unemployed philosophers, being friends, were to continue to meet in a pub regularly, where they still discussed philosophical issues. Not only would they not form the same organisation, they would not be an organisation at all, because they lost the distinctive necessary function of an organisation. Even if they played the finite-state-board-game, they would not become an organisation again, only a group which computes.

As computing artefacts, organisations differ from other kinds of groups, such as families and groups of friends, by necessarily being for the purpose of processing medium-independent vehicles. Some kinds of groups lack a necessary purpose, because they are not an artefact kind. For example, loose peer groups are arguably not artefacts in the way organisations are and lack a necessary function. Other kinds of groups have necessary purposes as artefact groups, but the purposes differ. A sports team might have the necessary function of playing sport, but not the function of processing medium-independent vehicles.[18] But importantly, in having such a modal property organisations do not differ from other artefact kinds. A pencil always and across all possible worlds has the purpose to produce lines on paper and similar surfaces. The pen might fail to do so under inappropriate circumstances or because of a broken pencil lead, but the object would not even be a pencil if it lacked this teleological function. Likewise, conceiving of organisations as a specific artefact kind clarifies their distinguishing modal characteristic: organisations and laptops are both artefact kinds created for a similar purpose.

One might object that the purpose of organisations as computing systems is not transparent to us, as is the case with artefact kinds such as laptops and pencils. After all, with artefact kinds we usually know what their teleological purpose is in virtue of the purpose's connection with the intentions of the kinds creators – a difference to the case of water being H_2O. But this transparency claim can be questioned. For example, it is not obvious that the purpose of coins as a generic medium of exchange was transparent at the creation of this artefact kind. Coins might have only been intended as storage medium or exclusively for state payments, acquiring their general purpose only later without transparent creation intentions.[19] But even if being an artefact kind implies some transparency of the kinds' function, we might still fail to realise that the function we have given to an artefact amounts to computing. We might happily describe organisations in the terms of the sufficient conditions for being computing systems without drawing the conclusion that they are computing systems. To draw this connection, we require a theory of computation, such as the mechanistic account discussed in the present paper.

The purpose of organisations as artefacts is to compute and can be further specified in a multitude of ways, from computing the registration of a sports team to a multinational cooperation computing its market strategy. Organisations exhibit an enormous variety, and they undergo change to structure and purpose, but all of them are always and across all possible worlds computing systems.

The modal aspect explains the intuitive difference between organisations and other groups such as the family playing the finite-state-board-game and a sports team. It also creates a difference to previous proposals that have described groups as computing systems. I turn now to the existing literature to establish this difference.

---

18 Qua-objects pose a subtle challenge. For example, one might wonder whether the family-qua-finite-state-board-game-team does not have the necessary purpose of computing. This challenge, however, relies on the status of qua-objects; a topic which goes beyond the confines of the present paper. Furthermore, even if the qua-family would have a necessary purpose like an organisation, I have not claimed to identify all distinguishing features of organisations, just one that separates them from many other paradigmatic group kinds.

19 For some more discussion of such artefact kinds and their functions, see Strohmaier 2020b.

6 Comparison with Previous Proposals

Brian Epstein (2019) introduces his paper on group ontology by noting that groups and their features are important innovations in human society that are often taken for granted. Organisations, I have argued, are a special kind of group and they are a special innovation. They are necessary computing systems and as such an impressive contribution to human society that has hitherto been neglected in group ontology. That is not to say, that there have not been similar ideas in the philosophical literature before.

While the description of organisations as groups necessarily structured to be computing systems has not received attention in recent group ontology, the literature on social cognition and social ontology includes related proposals. I will in the following consider two debates, the debate on the extended mind and the debate on group agency, that are related but do not cover quite the same ground, before I turn to more closely related proposals that explicitly concern group computation.

6.1 Extended Cognition

According to theories of extended cognition, cognitive processes may extend beyond the skull and also involve tools and social context, e.g. a notebook becomes an extension of memory (see Clark and Chalmers 1998; Gallagher 2018). Similarly, other members of an organisation could serve to extend our cognition beyond the cranial boundaries. Attribution of extended cognition has enjoyed increasing popularity in the social ontology and social cognition literature (e.g. Gallagher and Crisafi 2009; Galloti and Huebner 2017; Tollefsen 2006, 2015).

The theory of extended cognition has also been applied to social groups. Assuming a computational analysis of cognition, a group engaging in extended cognition computes. Hutchins (1995) early and influential study of the navigation on a ship serves as a case in point, as he endorses an explicitly computational account of the social process on a ship's bridge. Thus, the question arises whether my proposal contributes something beyond the application of extended cognition. In fact, it does so in at least three ways.

First, even if all organisations were instances of extended cognition, the present proposal is the first to describe the ontology of organisations in terms of computation. The thesis of extended cognition has served mainly as a way to challenge pre-conceptions in the philosophy of mind and cognitive science, and has been put to other uses only on the side.[20] It has not found much attention within the group ontology debate. That all organisations are necessarily computational systems is a substantial thesis that should surprise even those who have endorsed theories of extended cognition, since it modally links a particular kind of group with computation.

Second, many theories of extended cognition rely on less restrictive accounts of computation. Hutchins, for example, described computation as "the propagation of representational state across a variety of media" (Hutchins 1995: xvi) – a less demanding and underdeveloped analysis than the one offered by Piccinini. By endorsing the mechanistic account of computation, I apply a more restrictive and developed account of computation to the study of organisations.

Third, I argue for an application of the computational approach to all organisations, but not that they have to be extensions of individual cognition. When Hutchins argues that the interaction processes of a navigation team on a bridge are computational, he does so to establish the interaction as an extended cognitive process. Not all computational processes, however, are instances of extended human cognition. Thus, there is at least a conceptual difference between the proposal of organisations as computing systems and the claim that organisational processes are instances of extended cognition.

The difference becomes apparent by considering that various arguments against extended cognition do not apply to a computational ontology of organisations. For example, one argument against extended cognition claims that there is no full parity between the extended processes and the internal cognitive processes (Spaulding 2012). As the objection goes, the functional profiles of the processes are not sufficiently similar. This is not a problem for the computational theory of organisations. While such a parity to human minds might be required for extending human cognition, it is not needed for organisational processes to be computational processes. A computing system can have a very different functional profile from a human brain.

Another challenge to extended cognition is that an extended taxonomy of cognitive kinds does not bring sufficient explanatory benefits to justify abandoning the individualistic approach to human cognition (Rupert 2004, 2005, 2009). Again, this problem does not afflict the computational theory of organisations. In contrast to theories of extended cognition, I do not rely on such a revision of cognitive science. I can be agnostic about human cognitive theory.

20 One interesting application has been in the field of human-computer interaction, see Hollan, Hutchins, and Kirsh 2000. Again, my own proposal has a broader scope than this suggestion, insofar it is not limited to interaction with silicon-based computing artefacts.

Despite close connections, the computational ontology of organisations thus remains distinct from the endorsement of extended cognition. Cases of extended cognition in organisations form a subset of computational organisational processes.

6.2 Group Agency

In recent years, multiple authors have argued that some groups are agents with desires and beliefs of their own. Most examples in this literature have been of organisations, such as the Ford Motor Company (Tollefsen 2015: 106) or tenure committees (List and Pettit 2011: 92).[21] Given a computational theory of mind, one would expect connections between the theory of group agency and the computational ontology of organisations.

Does the claim that all organisations are computing systems entail that all organisations are group agents? Are all group agents also computing systems? The answers to these questions depend on the account of group agency one adopts. The literature is roughly dominated by two theories of how to attribute minds to groups, the first being functionalism, the second interpretivism (cf. Strohmaier 2020a; Tollefsen 2015).

There are varying formulations of both approaches, but the main difference lies in whether the mental states of a group are identified with internal states of the group or with overall behavioural dispositions. Functionalist proponents of group agency consider beliefs and desires to be internal states of a group with the appropriate functional profile, that is the right kind of causal connections to other states and behaviour.[22] Interpretivists draw instead on Dennett's (1987) theory of the intentional stance, according to which a system has beliefs and desires if interpreting its behavioural patterns in this way is useful. On both approaches, there are interesting connections between group agency and computation.

Following Putnam's (1960, 1967) ground-breaking work, many functionalists think of mental states as computational states. Piccinini developed his analysis of computing system partially to offer a better grasp of such a computational approach to the mind. While alternative formulations of functionalism have been developed, at least for such types of functionalism all group agents would be computing systems.[23]

21 One of the three parts of List and Pettit's influential book *Group Agency* is in fact entitled "The Organizational Design of Group Agents".

22 One could also introduce non-causal relations, e.g. grounding, as part of an unusual functional profile, but for the present purposes I leave this possibility aside.

23 Again, this does not imply that all group agents are organisations, because of the modal aspect of the proposed theory. The group might be an agent contingently, but organisations are necessarily structured to be computing systems.

Functionalism does not entail the converse, that all groups which are computing systems are group agents. No functionalist has yet claimed that all organisations are group agents, and this is unlikely to change even if the proposed theory was accepted. Many organisations are not set-up to have the kinds of complex internal states that functionalists expect to realise beliefs and desires (cf. Huebner 2014 on this issue). Not all computing systems are agents. If the sports centre were exactly the sketched finite state machine, then it would not be a full-blown agent with beliefs and desires, despite its computational nature, because the computational states do not have the functional profiles of these mental states – for one, beliefs and desires respond to a much larger range of input since they are connected to more states.

Since interpretivists do not require appropriate internal group states for the realisation of beliefs and desires, they tend to be more liberal in their attribution of group agency. Consequently, interpretivists might postulate group agents which fall short of the requirements for being computational systems according to the analysis I have used. The notion of function employed by Piccinini's theory of computation is thicker than the one typically employed by interpretivists.[24]

In virtue of this permissive attribution practice, interpretivists might also assert the converse that all organisations are group agents. Interpretivism opens the door to claiming that when organisations work towards a goal state, they *want* to reach it. On this interpretation, the sports centre wants to register the team and therefore it provides the relevant forms to fill out. After all, Dennett (1987) happily attributes minimal representational states to thermostats.

Interpretivists could avoid this conclusion by claiming that in specific cases attributing agency does not bring sufficient predictive value. For example, an interpretivist could argue that attributing desires and beliefs to the sport centre does not bring any additional benefits over interpreting it as a finite state machine. Thus, the fate of the converse depends on the relative predictive value.

But while the details of the connection between the computational nature of organisations and group agency depend on the account of the latter, neither functionalism nor interpretivism makes the modal claim which distinguishes organisations from other groups. No theory of group agency distinguishes organisations from other kinds of groups by describing being structured for computing as a necessary feature of organisations.

24 That being said, interpretivists might also tend towards a thinner realism regarding computing systems. Drawing once more on Dennett's work, they might propose that all that is needed for something to be a computing system is that it would be useful to interpret it as such a system from the design stance. As I discussed, my proposal is committed to a relatively restrictive analysis of computation, at the danger of being trivialised.

6.3 Group Computation

A number of contributions to the literature endorse neither specifically the extended mind, nor group agency, but instead offer a picture of groups as engaging in computation. One example of this is Goldstone and Theiner's (2017) "multiple, interacting levels of cognitive systems perspective on group cognition". They judiciously avoid the term "mind" and instead talk about cognitive systems in a very broad sense (Goldstone and Theiner 2017: 339, 348–351).

The aim of this theory differs from mine. Goldstone and Theiner do not focus specifically on a kind of group but instead seeks to establish a common framework applicable to different levels of analysis. The goal is to show that considering systems like groups to implement cognitive systems, such as lateral inhibition networks, can be a useful perspective. Hence, the paper is best taken as supporting my claim that groups can implement computing processes. It does not bear, however, on whether organisations have a specific relation to computation.

While I have proposed a necessary computational nature as a distinguishing mark of organisations from other kinds of groups such as families, Goldstone and Theiner seek to apply their perspective as broadly as possible and accordingly avoid venturing into the group ontology debate. Accordingly, they also do not postulate a necessary connection between being an organisation and being a computing system that distinguishes organisations from other groups such as families.

Perhaps closest to the present proposal is Huebner and Jeabri's (2018; see also Huebner 2014), "Computational Theories of Group Behaviour". Like the proposed analysis of organisations, they explicitly draw on Piccinini's theory of computation.[25] However, while they reference Piccinini prominently in the introduction of their paper, Huebner and Jebari spend significantly less time on establishing that human groups fit this particular analysis of computation. Instead they resort to a more generic notion of information processing that makes it easier to apply to various human and non-human phenomena.

Furthermore, the contribution of Huebner and Jebari is not a proposal regarding group ontology, but group cognition. They treat group level information processing as a step towards establishing group cognition and even group agency.

Similar to Goldstone and Theiner, the work of Huebner and Jebari can be taken to support the claim that groups implement computations, without addressing the

---

25 Huebner and Jebari do not draw on the 2015 monograph by Piccinini, but earlier work. This leads to some subtle differences with my proposal. This earlier work (Piccinini and Scarantino 2011), however, already refers to medium-independent vehicles. That this term does not appear in Huebner and Jebari illustrates that they do not cover the conditions of Piccinini to the same extent as the present paper.

question whether kinds of groups can be distinguished on the basis of their relation to computation. As a consequence, all these contributions do not cover the specific innovation of organisations over other groups, such a kinship groups. A kinship group is not by necessity a computing system; an organisation is. The group ontology debate can learn from these contributions to cognitive science, but they are not engaged in the same project.

7 Conclusion

Examples of organisations have figured prominently in debates on group ontology, but the contemporary debate lacks a theory of organisations and their functions. The present paper put forward an ambitious ontology of organisations as groups necessarily structured for computing.

The example of a sports centre registering sports teams served as an illustration of an organisation as computing system. To substantiate the proposal, I discussed Piccinini's restrictive account of physical computation. Accordingly, organisations have the function to manipulate medium-independent vehicles according to rules. After fending off challenges to this proposal and showing that the modal character allows to distinguish organisations from other groups, I discussed the relation of my proposal to existing theories. Despite connections, the computational theory of organisations is a hitherto neglected position in the group ontology debate.

Organisations are an important innovation in human society. As computing systems, they are a special kind of group that has allowed us to accomplish complex tasks. The proposed theory validates using computational explanations in organisational studies. Furthermore, like paradigmatic natural kinds, organisations have this distinguishing feature necessarily. It is time for the social ontology literature to appreciate the necessarily computational nature of organisations.

Acknowledgements: I thank Stephanie Collins, Simon Wimmer, Fabio Rojas, Brian Epstein, and others I cannot remember for discussion and comments. This paper reports on research supported by Cambridge Assessment, University of Cambridge.

References

Ashforth, B. E., and V. Anand. 2003. "The Normalization of Corruption in Organizations." *Research in Organizational Behavior* 25: 1–52.

Chalmers, D. J. 1996. "Does a Rock Implement Every Finite-State Automaton?" *Synthese* 108 (3): 309–33.

Cioffi-Revilla, C. 2017. *Introduction to Computational Social Science*. Cham: Springer International Publishing.

Clark, A., and D. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.

Coelho Mollo, D. 2018. "Functional Individuation, Mechanistic Implementation: The Proper Way of Seeing the Mechanistic View of Concrete Computation." *Synthese* 195 (8): 3477–97.

Coelho Mollo, D. 2019. "Are There Teleological Functions to Compute?" *Philosophy of Science* 86 (3): 431–52.

Dennett, D. C. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.

Effingham, N. 2010. "The Metaphysics of Groups." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 149 (2): 251–67.

Epstein, B. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. New York: Oxford University Press.

Epstein, B. 2019. "What Are Social Groups? Their Metaphysics and How to Classify Them." *Synthese* 196 (12): 4899–932.

Fodor, J. A. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.

Fodor, J. A. 1981. "The Mind-Body Problem." *Scientific American* 244 (1): 114–23.

Gallagher, S. 2018. "The Extended Mind: State of the Question." *The Southern Journal of Philosophy* 56 (4): 421–47.

Gallagher, S., and A. Crisafi. 2009. "Mental Institutions." *Topoi* 28 (1): 45–51.

Gallotti, M., and B. Huebner. 2017. "Collective Intentionality and Socially Extended Minds." *Philosophical Psychology*: 1–18. https://doi.org/10.1080/09515089.2017.1295629.

Gilbert, M. 1992. *On Social Facts*. Princeton: Princeton University Press.

Goldstone, R. L., and G. Theiner. 2017. "The Multiple, Interacting Levels of Cognitive Systems Perspective on Group Cognition." *Philosophical Psychology* 30 (3): 334–68.

Gordon, B., and G. Theiner. 2015. "Scaffolded Joint Action as a Micro–Foundation of Organizational Learning." In *Contextualizing Human Memory: An Interdisciplinary Approach to Understanding How Individuals and Groups Remember the Past*, edited by C. B. Stone, and L. Bietti, 154–86. New York: Psychology Press.

Hawley, K. 2017. "Social Mereology." *Journal of the American Philosophical Association* 3 (4): 395–411.

Hollan, J., E. Hutchins, and D. Kirsh. 2000. "Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research." *ACM Transactions on Computer-Human Interaction* 7 (2): 174–96.

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

Huebner, B., and J. Jebari. 2018. "Computational Theories of Group Behavior." In *The Routledge Handbook of the Computational Mind*, edited by M. Sprevak, and M. Colombo, 483–96. London: Routledge.

Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.

Khalidi, M. A. 2013. *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*. Cambridge: Cambridge University Press.

Khalidi, M. A. 2015. "Three Kinds of Social Kinds." *Philosophy and Phenomenological Research* 90 (1): 96–112.

Kripke, S. A. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

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

March, J. G., and H. A. Simon. 1958. *Organizations*. New York: Wiley.

Mason, R. 2016. “The Metaphysics of Social Kinds.” *Philosophy Compass* 11 (12): 841–50.
Miller, J. H., and S. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton: Princeton University Press.
Morgan, G. 2006. *Images of Organization*. London: Sage Publications.
Piccinini, G. 2015. *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
Piccinini, G., and A. Scarantino. 2011. “Information Processing, Computation, and Cognition.” *Journal of Biological Physics* 37 (1): 1–38.
Preisendörfer, P. 2011. *Organisationssoziologie: Grundlagen, Theorien und Problemstellungen*. Wiesbaden: VS, Verlag für Sozialwissenschaften.
Putnam, H. 1960. “Minds and Machines.” In *Dimensions of Minds*, edited by S. Hook, 138–64. New York: New York University Press.
Putnam, H. 1967. “The Nature of Mental States.” In *Art, Mind, and Religion*, edited by W. H. Capitan, and D. D. Merrill, 1–223. Pittsburgh: Pittsburgh University Press.
Putnam, H. 1988. *Representation and Reality*. Cambridge, MA: MIT Press.
Ritchie, K. 2013. “What Are Groups?” *Philosophical Studies* 166 (2): 257–72.
Ritchie, K. 2015. “The Metaphysics of Social Groups.” *Philosophy Compass* 10 (5): 310–21.
Ritchie, K. 2020. “Social Structures and the Ontology of Social Groups.” *Philosophy and Phenomenological Research* 100 (2): 402–24.
Rupert, R. D. 2004. “Challenges to the Hypothesis of Extended Cognition.” *Journal of Philosophy* 101 (8): 389–428.
Rupert, R. D. 2005. “Minding One’s Cognitive Systems: When Does a Group of Minds Constitute a Single Cognitive Unit?” *Episteme* 1 (3): 177–88.
Rupert, R. D. 2009. *Cognitive Systems and the Extended Mind*. Oxford; New York: Oxford University Press.
Scott, W. R., and G. F. Davis. 2006. *Organizations and Organizing: Rational, Natural and Open System Perspectives*. London: Routledge.
Searle, J. R. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.
Sheehy, P. 2003. “Social Groups, Explanation and Ontological Holism.” *Philosophical Papers* 32 (2): 193–224.
Sheehy, P. 2006a. *The Reality of Social Groups*. Aldershot: Ashgate.
Sheehy, P. 2006b. “Sharing Space: The Synchronic Identity of Social Groups.” *Philosophy of the Social Sciences* 36 (2): 131–48.
Spaulding, S. 2012. “Overextended Cognition.” *Philosophical Psychology* 25 (4): 469–90.
Strauss, A. L. 1988. *Negotiations: Varieties, Contexts, Processes, and Social Order*. San Francisco, London: Jossey-Bass.
Strauss, A. L., L. Schatzman, D. Ehrlich, R. Butcher, and M. Sabshin. 1963. “The Hospital and its Negotiated Order.” In *The Hospital in Modern Society*, edited by E. Freidson, 147–69. Glencoe: The Free Press of Glencoe.
Strohmaier, D. 2018. “Group Membership and Parthood.” *Journal of Social Ontology* 4 (2): 121–35.
Strohmaier, D. 2020a. “Two Theories of Group Agency.” *Philosophical Studies* 177 (7): 1901–18.
Strohmaier, D. 2020b. “Social-Computation-Supporting Kinds.” *Canadian Journal of Philosophy* 50 (7): 862–77.
Tollefsen, D. 2006. “From Extended Mind to Collective Mind.” *Cognitive Systems Research* 7 (2–3): 140–50.
Tollefsen, D. 2015. *Groups as Agents*, Cambridge, UK; Malden, USA. Polity.
Updike, J. 2003. *The Early Stories, 1953–1975*. New York: Alfred A. Knopf.

Uzquiano, G. 2004. "The Supreme Court and the Supreme Court Justices: A Metaphysical Puzzle." *Noûs* 38 (1): 135–53.

Uzquiano, G. 2018. "Groups: Toward a Theory of Plural Embodiment." *The Journal of Philosophy* 115 (8): 423–52.

Vaughan, D. 1999. "The Dark Side of Organizations: Mistake, Misconduct, and Disaster." *Annual Review of Sociology* 25 (1): 271–305.

Weber, M. 2013 [1921–1922]. "Economy and Society." In *Two Volume Set, With a New Foreword by Guenther Roth Edition*, edited by G. Roth, and C. Wittich. Berkeley: University of California Press.