

ARTICLE

Social-Computation-Supporting Kinds

David Strohmaier 

Department of Computer Science and Technology, University of Cambridge, Cambridge, UK
Email: ds858@cam.ac.uk

Abstract

Social kinds are heterogeneous. As a consequence of this diversity, some authors have sought to identify and analyse different kinds of social kinds. One distinct kind of social kinds, however, has not yet received sufficient attention. I propose that there exists a class of social-computation-supporting kinds, or SCS-kinds for short. These SCS-kinds are united by the function of enabling computations implemented by social groups. Examples of such SCS-kinds are *reimbursement form*, *US dollar bill*, *chair of the board*. I will analyse SCS-kinds, contrast my analysis with theories of institutional kinds, and discuss the benefits of investigating SCS-kinds.

Keywords: Social ontology; social kinds; institutional kinds; computation

Introduction

Social kinds are heterogeneous. Some social kinds, such as *US dollar bill*, depend on propositional attitudes toward the kind. Others, such as *recession*, do not. Some social kinds, such as *reimbursement form*, are purposive—that is, they have a function. Others, such as *inflation*, have no function. As a consequence of this diversity, some authors have sought to identify and analyse different kinds of social kinds (e.g., Khalidi 2015). One distinct kind of social kinds, however, has not yet received sufficient attention.

I propose that there exists a class of social-computation-supporting kinds, or SCS-kinds for short. These SCS-kinds are united by their role in social processes. They have the function of enabling computations that are realised by social groups. By being available to human agents, they support social processes such as organising production in a company or distributing information across a wide set of individuals. Examples of such SCS-kinds are *departmental reimbursement form*, *US dollar bill*, *chair of the board*.¹

So far, these examples of social kinds have been understood as institutional kinds (Searle 1995; Guala and Hindriks 2015), but such accounts have largely neglected the computational nature of this subset of social kinds. Consequently, their functions, their mind-dependence, and their power have been mischaracterised. Lacking the notion of social computation, other analyses of these social kinds misconstrue their unique characteristics. I propose that the theory of SCS-kinds remedies this situation. It identifies the explanatorily more powerful kind of social kinds.

The role of SCS-kinds in enabling social processes also motivates paying increased attention to them. Our interactions would be impoverished if we could not think about the fabric pieces in our pockets as dollar bills and the person at the end of the table as the chair of the board. Such kinds can contribute to social progress in similar ways as technological advances. Hence, achieving a better

¹Here and in the following I use italics to distinguish references to kinds from references to their instances.

understanding of their characteristics is not just a step forward in social ontology but of practical significance.

I start by sketching the debate and current positions on social and institutional kinds. To identify SCS-kinds, I then discuss the conditions for social computations. This theory of social computation will allow me to provide the necessary and sufficient conditions for a social kind to be an SCS-kind. At the core of the analysis is the function of SCS-kinds, which is to support social computations appropriately. I end by sketching the benefits of investigating SCS-kinds.

Social and institutional kinds

While social kinds have received increasing attention, there are no universally accepted criteria for what makes a kind social. Very few philosophers even dare to propose such criteria.² Instead the literature has produced a number of prototypical examples that are assumed to provide sufficient guidance for the debate. *US dollar bill*, *recession*, and *student* are all considered social kinds. The present paper identifies a specific kind of social kinds, and therefore rivals other accounts that have identified institutional kinds as a prominent subset of social kinds.

I will address three major points of contention in this debate on social and institutional kinds: mind-dependence, power, and functions. Mind-dependence has perhaps received the most attention.³ Some, such as John Searle (1995, 2010), have considered most or all social kinds to be mind-dependent. Searle's theory is formulated in terms of facts rather than kinds but it suggests that social kinds are characterised by status functions which human agents assign to physical objects. Status functions specify the function of instances of a kind—for example, that a screwdriver is for installing and removing screws.

In response to Searle, cases of social kinds that do not appear to depend on mental states toward them have been pointed out, e.g., recessions (Thomasson 2003). Muhammad Khalidi (2015) has distinguished three kinds of social kinds depending on whether (a) the existence of the kind depends on propositional attitudes toward the kind and whether (b) the existence of an instance depends on a propositional attitude toward the instance.⁴

First, Khalidi suggests that for some social kinds, such as *recession*, neither the kind itself nor its instances depend on propositional attitudes toward them. An ancient economy might have suffered a recession without conceiving of it as such or having developed the concept. Second, in cases, such as *US dollar bill*, the dependence only holds for the kind not the instances. Searle (1995, 31) offers the example of a dollar bill which has fallen off the printing press and through a crack on the floor before anyone had propositional attitudes toward it. It is still a dollar bill. In a third type of cases, such as *prime minister*, both the kind and the instances depend on propositional attitudes toward them. If all people stopped thinking of Boris Johnson as prime minister of the UK, he would cease to be an instance of the kind.

The diversity of social kinds leads to the question whether there are systematic reasons why some kinds belong into the categories provided by Khalidi. Second-order kinds, such as *institutional kind* and my proposal of *SCS-kind*, serve to explain the specific dependence relations, amongst other things.

Closely linked to the issue of mind-dependence is the issue of the power of social kinds. Kinds such as *departmental reimbursement form*, *US dollar bill*, and *chair of the board* seem to have a

²Margaret Gilbert's (1992, 2013) joint commitments and Searle's (1995, 2010) status functions are two of the best known features that have been supposed to distinguish social kinds. Neither has found anything close to universal approval.

³Some of the attention to mind-dependence results from the importance of this issue for realism about social kinds (Mason 2016) and the question of individualism (Epstein 2015).

⁴Khalidi only discusses three kinds and neglects a fourth potential kind of kind, where only the instances but not the social kind itself depends in its existence on propositional attitudes toward it. Whether any such kinds exist is an open question that I will not address.

special role in enabling social interaction that is different from that of a screwdriver. Searle described the screwdriver as a social object to which we have assigned a function, but it can fulfill its function in virtue of its physical makeup. From such cases Searle distinguishes institutional kinds which can fulfill their function independently of their physical makeup. They appear to have a special power.

Currency kinds such as *10 dollar bill* are perhaps the most prominent example of such institutional kinds. Though 10 dollar bills are printed on cotton-linen fabric, this physical substrate might be replaced with plastic in the future and, in principle, the bills can fulfill their function if they were made of stone (*cf.* Friedman 1994). The general presumption is that the minds of individuals lead to a dissociation between physical makeup and social power of the kind instances. Our representations of the objects, be they made of cotton-line or stone, enables them to exert power in social interactions.

That some instances of social kinds enable complex social interactions independently of the physical makeup is one of their most intriguing features. Searle seeks to explain this capacity of a subset of social kinds by attributing deontological powers to them—that is, facts about institutional kinds are supposed to give us desire-independent reasons for action.⁵ The status of someone as chair of the board gives members of the relevant organisation desire-independent reasons to follow the chair's orders.⁶

Another approach to the power of a subset of social kinds has been game theoretical. Such an alternative account of institutions has been developed by Francesco Guala and Frank Hindriks (2015; Guala 2016; Hindriks and Guala 2019). They seek to unify a game-theoretical equilibrium-based account of institutional kinds with an account like that of Searle, where they interpret the status functions as providing rules for action. Their idea is that institutional kinds such as *property* and *state border* help to change games in a way that opens up different equilibria.

Guala and Hindriks discuss how two groups can avoid a fight over a scarce resource, e.g., grazing land, by creating rules about who can use and transfer the resource. Such rules allow the agents to settle on a better equilibrium, enabling action without costly fighting and thereby increasing overall payoffs. For the sake of simplicity, the groups can then call the resources “property” insofar as they are subject to such rules.⁷ According to Guala and Hindriks, such institutional kinds can be understood as the result of introducing theoretical terms to reduce the cognitive load of the apparatus of rules (status functions) we have created to achieve better equilibria.

As a consequence, Guala and Hindriks account for the power of institutional kinds in quite a different way from Searle. There is no need for desire-independent reasons; instead, institutional kinds have their power in virtue of their role in changing the equilibria of social interactions. Put differently, institutional kinds have their power because of the features of the game they transform. If there were no possibility of the behavioural rules to affect the equilibrium, then the kinds would lack any power.⁸

⁵Although the details differ, Gilbert's (1992) account in terms of joint commitments gives a similar role to normativity in explaining the power of social kinds. The notion of deontological powers has also been used by Searle to accommodate a challenge. As has been argued by authors such as Barry Smith (2003) and Amie Thomasson (2003), not all instances of social kinds are realised by physical objects. Thomasson points to the US Constitution—not to be confused with the document on which it was first written—which exists but is not realised by any physical object. Searle (2010) has accepted the examples of his critics and argued that these nonphysically realised social kinds are still cases of we-intentionally assigned powers.

⁶Consideration of other examples makes such reasons look less plausible as candidates for explaining the power of social kinds. It seems unlikely that the status of being a US dollar bill gives you a reason to accept it as a form of payment independently of your desires.

⁷The case of money which Guala (2016) discusses in the third chapter of his book *Understanding Institutions* is also illustrative of this approach.

⁸The rules do not have to do all the work on their own. There might be coordination devices, ranging from harmless signals to announcements backed by government force. In the end, however, everything comes down to rules in equilibrium.

The two accounts also differ on the third issue concerning social kinds. Searle has argued that the functions of social kinds, be it the function of a screwdriver to screw or a dollar bill to serve as payment, are the result of what one might consider projection, i.e., the mental assignment of a status. Hindriks and Guala (2019) have accepted this as only one part of the story. They distinguish what they call the etiological function from what they call the teleological function of institutional kinds. The etiological function is determined by the reasons for the kinds existence and persistence.⁹ Teleological functions are normatively evaluative and projected upon the kinds by us. This part of their account resembles Searle's, although Hindriks and Guala draw upon Dennett's (1987) notion of a design stance from which the institutional kinds are interpreted to explain the projection and do not require any we-intentionality.

According to Guala and Hindriks, the etiological function of institutional kinds is to generate cooperative benefits and their teleological function is to secure values—that is, support or promote values of members in society. In sum, institutional kinds have their causal source in the game theoretical benefits they create via cooperation, and we project upon them the function to secure values.

These three issues should receive attention when analysing a kind of social kinds. Is the kind of social kinds mind-dependent? Why do instances of this kind have power in social interactions? Are the kinds purposive and, if so, what is the source of their functions? The theory of SCS-kinds will have its distinctive place in these debates, and I will explicate connections and differences to prior work throughout the analysis. It will compete with theories of institutional kinds as explaining what is special about a subset of social kinds. While I do not deny that there might be instances of institutional kinds as sketched by Searle or Guala and Hindriks, *SCS-kind* is the explanatorily more significant second-order kind of social kinds.

The analysis will offer a functional account of a specific kind of social kinds just as theories of institutional kinds do, but it will identify a different function than Guala and Hindriks. It will refer to the minds of agents, but it will not be a projective account of the sort proposed by Searle. The key to this kind of social kinds and their power lies in their role in social computations—more specifically, computational social processes which the kinds support. In the next section, I discuss these computations.

Social computations

Not all cases of social computations involve SCS-kinds, but the analysis of SCS-kinds requires the notion of a social computation. The type of social kinds is analysed as supporting social computations.

Social computations are computational processes realised by social groups. A computational process is a social computational process if and only if it is instantiated by a social group. The groups serve as the physical computing systems manifesting the computational processes. A team of researchers engaging in a calculation—for example, calculating the potential impact of a nuclear bomb during the Manhattan project—is one instance of social computation. Here a team, not just one member on the team, instantiates a computational process.

A more ambitious example of social computation is a family organising their next vacation by allocating various subtasks to its members. The parents might check flights and implement a sorting algorithm, while a daughter identifies available hotels according to specific search criteria. Later the results are combined in a rule-governed manner. As long as the family follows a procedure that is correctly described as a computation spread out across its members, it will serve as another example of social computation, albeit not necessarily of SCS-kinds.¹⁰

⁹Hindriks and Guala combine etiological and selectionist theories of functions into one category, while I will later specifically commit to a generalised selectionist account of functions.

¹⁰Under the assumption that there are social groups with no members, as proposed by Brian Epstein (2015), social computations can also occur in this limiting case of groups. I will not address such cases in the following.

I do not endorse a specific account of physical computation (*cf.* Piccinini 2017). On the one end of the spectrum are extremely permissive accounts, such as various forms of pancomputationalism (Putnam 1988; Searle 1992; Chalmers 1996) and Dennett's design stance theory (1987). According to these theories, either all or an overwhelmingly large number of social groups are computing systems. Pancomputationalism turns all physical systems, which I take to include all social groups, into computing systems because their behaviour can be mapped to formal specifications of computing systems. For Dennett's theory, the question is whether interpreting a physical system as a computing system from a design stance brings a predictive benefit. If it helps to interpret the family as implementing computational algorithms to predict the decision they reach, then it is a computational system.

On the other end, we find more restrictive accounts, e.g., syntactic (Fodor 1975, 1981) and mechanistic accounts (Piccinini 2015; Coelho Mollo 2018). On these accounts, computing systems are typically distinguished by a stance-independent function. Independently of the interpretation by human agents, physical computing systems have either the function of syntactic manipulation or the function of mechanistic rule-application to medium-independent vehicles.¹¹

I assume that even on the more restrictive accounts there will be sufficient cases of social computation to render SCS-kinds of interest (see Strohmaier, manuscript). For example, on a mechanistic account like that of Piccinini, the function of a team of researchers might be to manipulate strings of formulas according to rules to arrive at a result fulfilling certain constraints. These strings would be medium-independent vehicles.

Assuming that social computations are widespread might appear audacious, especially given that restrictive accounts of computation are one option, but it is not without precedent. In the following, I will discuss previous proposals of social computation in cognitive science and philosophy of mind before I point out similar theories in the social sciences.

In philosophy, the existence of social computation is usually implied by even stronger cognitive claims. Theories of the extended mind conceive of smartphones as extensions of human cognition because they are integrated in one computational process (Hutchins 1995; Clark and Chalmers 1998; Gallagher 2018). For example, the cognitive process of geographical orientation happens partially through the extension of Google maps.

Such theories of extended cognition have also been applied to social groups and institutions. Shaun Gallagher has suggested that cognitive processes are "constituted in various social practices that occur within social and cultural institutions" (Gallagher 2013, 4; see also Gallagher and Crisafi 2009). On this account, some institutions—for example, institution in legal proceedings—serve to support the extended cognition of human agents. As long as these cognitive processes are computational processes, they are instances of social computation (*cf.* Huebner and Jebari 2019).

Coming from a different direction, theories of group agency also postulate social computations under reasonable assumptions. These accounts argue that organisations, and even entire nation states, are agents with propositional attitudes. Assuming that propositional attitudes can be analysed as computational states, theories of group agency also postulate social computation (List and Pettit 2011; Huebner 2014; Tollefsen 2006, 2015). The internal cognitive processes of group agents, if there are any, would serve as prime examples of social computation.¹²

Both theories of the extended mind and group agency apply cognitive concepts to social processes. Talk about computation avoids such strong claims about cognition, because it does not claim the processes are cognitive. It commits us neither to the cognition of individuals extending

¹¹One might raise the question whether only computing systems can implement computations. Perhaps a group can implement a computation without having the function of doing so? I leave this issue open. All that is required for social computation is that a social group implements the computation.

¹²An exception has to be made for group agents who are dominated by a single dictator in the game theoretical sense of the term (see List and Pettit 2011). Then the cognitive processes would not necessarily be an instance of social computation. But these would be the limiting cases of group agency and not the prime examples of the debate.

beyond their cranium nor to the agency of groups. Consequently, the analysis of SCS-kinds avoids many controversies surrounding the extended mind and group agency.

For example, Adams and Aizawa (2008) have argued that purported instances of extended minds lack the mark of the cognitive, which they proposed was nonderived representational content. But for social computation no such content is required. While the representational content of the computations across a team in the Manhattan project might be derived, the process would still be a computation implemented by the group. The mark of cognition is not a mark of computation. Thus, even if Adams and Aizawa's argument were to succeed, it would not undermine my proposal.

Similarly, the theory of social computation does not require any groups to be full-blown agents with desires and beliefs or equivalent mental states. If mental states such as beliefs and desires are computational states, they are special ones and not realised in the course of all computational processes (cf. Huebner 2014). A pocket calculator has no desires. As far as the proposal of SCS-kinds goes, the groups might in their capacities be closer to pocket calculators than to human agents. While the assumption of social computations is far from trivial, it isolates the less controversial part of some of the proposals in the social cognition and social ontology literature.

In addition to cognitive science and philosophy, conceiving of social processes as computational processes has a history in the social sciences. One of the first approaches along these lines was the Carnegie School of organisational behaviour. Under the influence of Herbert Simon (March and Simon 1958), it described organisations as implementing procedures. Organisations were understood in terms of computational processes.

More recently, the field of computational social sciences follows "an *information processing paradigm of society*" (Cioffi-Revilla 2017, 2; emphasis in the original). It describes various social arrangements as complex adaptive systems that process information and are subject to a computational methodology. While it is not always clear whether the social processes are merely computationally modelled or in fact computational themselves, there is a close affinity to the present proposal.

As can be seen, there is sufficient support for my assumption of widespread social computation in cognitive and social science, and I am relying only on a moderate version of the claim.

In addition to the concept of social computation, the analysis of SCS-kinds will also require the notion of a realising agent. Intuitively, when a team engages in a calculation and one person engages in one part of the overall process, say calculating one variable, they are a realising agent of this computation.

A group member, *A*, is a realising agent¹³ of a social computational process, *P*, if and only if:

- (1) *P* is composed of the subprocesses p_1, p_2, \dots
- (2) At least one of these subprocesses is realised by *A*.

I assume here that computational processes allow a mereological analysis, i.e., that they can be understood as either atomic or composed of multiple subprocesses.¹⁴

Not all subprocesses have to be instantiated by group members for a process to be a social computation. Some might be implemented using digital devices such as smartphones or pocket calculators, as suggested by the extended mind literature. In some cases, the entire social computation might be automated so that there are no realising agents, for example in a quantitative trading

¹³There is no requirement that this agent has to be a human individual. Some authors have accepted that groups can have other groups as members (e.g., Ritchie 2020). Then these member groups can also serve as realising agents. For the sake of simplicity, I will stick to individuals as realising agents in the examples.

¹⁴Social computation can be atomic and still have a subprocess if we are using "parthood" such that every object is trivially a part of itself.

computation. That being said, SCS-kinds will have their effect on social computations in virtue of their specific relation to the realising agents, as the next section describes.

SCS-kinds

Among social kinds there is a class of social-computation-supporting kinds, or SCS-kinds for short. Paradigmatic examples include *departmental reimbursement form*, *US dollar bill*, and *chair of the board*. An SCS-kind has the function of enabling social computations by serving as a cognitive resource to the participants. Thus, a social-computation-supporting kind is distinguished by its function.

One aspect of the function of the kind *departmental reimbursement form* is to figure in the cognitive processes of agents instantiating social computations. For example, the kind can help a philosophy department to implement a more efficient accounting process than unsorted receipts. That the personnel interprets the sheet of paper accordingly is a condition for it fulfilling its function.

Formally, a social kind, K , is an SCS-kind if and only if instances of K have a function that for a social computational process P :

- (1) The instance is represented as K in subprocesses of P by realising agents.
- (2) In virtue of this representation, P is computationally more powerful or efficient than the social processes Q, R, \dots ,
- (3) where Q, R, \dots are the most similar possible social processes where no instance of K is represented as such by the realising agents.

This analysis describes SCS-kinds as difference-making with regard to social processes, i.e., to processes realised by social groups. As can be seen, this description captures the intuitive case of the *departmental reimbursement form*. Such forms make the operation of the department more efficient because the department members represent it appropriately as demanded in the second condition. While many kinds may support social computations, SCS-kinds do so in virtue of being represented by realising agents.

In the remainder of this section, I will unpack the analysis of SCS-kinds by first discussing the notion of function at play, then providing more details about the role of social computations, and finally discussing the requirement of representation by realising agents.

Function

SCS-kinds have a particular function that distinguishes them from other social kinds. I will call this function the SCS-function since it is the distinguishing mark of SCS-kinds. Functions have received much attention in the philosophy of biology and rival accounts have been proposed. We have already encountered the two types of function Hindriks and Guala distinguished.

The first type of account is projective. Searle's functions, as well as Hindriks and Guala's teleological functions, are the result of assignments by agents. The analysis of SCS-kind is better served by a nonprojective account of functions, similar to Hindriks and Guala's etiological account, for at least two reasons.

First, the analysis does not require the strongly normative force desired by Searle in general and by Hindriks and Guala for teleological functions. If a dollar bill fails to enable a transaction because it is too crumpled for the self-checkout register, it is a bad bill. According to Searle, we all imposed a function on the object and it just doesn't live up to our expectations. That justifiably makes it the object of our disapproval. But justifying such normative practices is not the goal of introducing SCS-kinds, which instead concern the persistence of certain kinds in virtue of their role in computation.

Second, the analysis does not require the transparency of projective accounts of functions. If one participates in the projection of a function on a screwdriver, then one knows what the function of a

screwdriver is. Or to take Hindriks and Guala's account, if an individual assigns functions from the design stance, these functions are transparent to them.¹⁵ By contrast, no one needs to know about the computational function of the SCS-kind instances for them to have this function. My proposal also covers social kinds which are opaque to us.

Hence a Millikan-style approach to functions is more appropriate for the proposed analysis of SCS-kinds. Approaches of this type¹⁶ have identified the history of evolutionary selection as the source of the heart's function to pump blood and the eye's function to see. These organs have what Millikan (1984, 28) calls their proper direct functions because they belong to a reproductive family with the appropriate history of evolutionary selection.¹⁷

For the present analysis, however, it is preferable to avoid the commitment to evolution. While the application of evolutionary theory to social processes is far from novel (Boyd and Richerson 2005; Skyrms 1996), it is still controversial, and especially hard to establish for kinds that lack a long history, e.g., *chair of the board*. If avoidable, the functions of social kinds should not be tied to social evolution.

Luckily, the selectionist approach to functions has been generalised beyond evolutionary processes.¹⁸ Rather than a process of trait selection through inheritance, all that is required for creating a function is the selective persistence of instances of the kind. According to the work of Justin Garson, "in order to acquire a new function, a trait need not have contributed to its own reproduction (i.e., making copies of itself). It need only have done something that allowed it to persist better (longer, more effectively) than some alternative trait within a population" (2017, 524; see also Garson 2011, 2012). Trait-differential persistence of instances is the source of their function.

To give an example, Garson points to the possibility of neuronal synapses acquiring a function by persisting in virtue of fulfilling a certain task. The synapses with the selected trait are more likely to persist than a random synapse and, hence, the trait becomes their function. Such a generalised account avoids assumption about specific evolutionary processes and the reproduction of social kinds. Therefore it is the account best suited for the analysis of SCS-kinds.¹⁹

With this notion of function in hand, the necessary condition for SCS-kinds can be spelled out as requiring them to have been subject to a general selection process that allowed the kind to persist either through differential retention—that is, because instances which have the feature are more likely to persist—or differential reproduction—because instances which have the feature are more likely to be reproduced. In these cases, the difference-making trait is their specific type of contribution to social computation.

As can be seen, Guala and Hindriks's etiological functions are much more closely related to the functions used in the analysis of SCS-kinds than Searle's normative functions. Both SCS-kinds and GH-kinds are enabling kinds, i.e., they generally are supposed to enable a social process, but the goal states of the functions differ. As discussed, Hindriks and Guala (2019) proposed that institutions have the etiological function to generate cooperative benefits by enabling better equilibria.

The SCS-function and the GH-function can overlap in the sense that enabling cooperation between agents can also enable social computation. For a simple example, consider a team engaging in a shared calculation where the steps need to be taken in a certain order. A kind similar to traffic

¹⁵One might change their account such that the projection from an idealised individual, rather than an actual one, is decisive for the teleological function. Then the teleological function would not be transparent to anyone actual, but it would still be transparent to this idealised individual and therefore the difference would remain.

¹⁶Millikan 1984, Wimsatt 1972. See also the etiological account of Wright 1973, 1976.

¹⁷In the following, I will speak broadly of such evolutionary selected functions and not draw the finer distinctions (e.g., direct proper vs. adapted function) proposed by Millikan.

¹⁸Millikan (1984, 40) had already taken steps in this direction in her early work. Garson and others have built upon her work on adapted functions. Millikan's (2014) own treatment of social kinds does not refer to evolution.

¹⁹One further alternative I do not discuss here are objective goal accounts, such as found in Piccinini (2015, chap. 6). They are less suitable because the objective goals are typically taken relative to some organism, but it is not clear which would be the organism in question—the individual human agents or the social group? These issues being resolved, Piccinini's account would allow the same analysis.

lights, an example Hindriks and Guala discuss, would enable this social computation and at the same time count as a GH-kind. Some SCS-kinds are GH-kinds, but this is not guaranteed.

Some of the SCS-kinds which are not GH-kinds pose a problem for Guala and Hindriks because they are intuitive examples of institutional kinds but lack the function of enabling cooperative benefits. Consider a scenario in which a company has a kind of form, say the *customer acquisition form*, which has to be filled out and seen by only one employee. In this case, it is hard to see how the form could be GH-type institutional kind since it is not obvious how it enables cooperation.

To push the example even further and remove all doubt, assume that the company is a one-person business. The business is run by one employee who fills out the forms and implements the group's organisational processes.²⁰ The *customer acquisition form* can still be specific to the group and fulfill its computational function, but its well-functioning can no longer be analysed using the account by Hindriks and Guala. There is no game inside the company. So Hindriks and Guala are apparently committed to the kind's malfunctioning. Any one-person company employing institutional kinds internally would be misusing them.

By contrast, my analysis grants that one-person groups can implement social computations. As long as the kind *customer acquisition form* adequately contributes to a social computation of the one-person group it can still be a well-functioning SCS-kind. While their game theoretical approach forces Guala and Hindriks to conceive of all institutional kinds as concerning the interaction between strategic agents, the SCS-function comes with no such commitment.

Enabling social computation does not have to be the sole function of an SCS-kind's instances. Even on the selection account, instances of one kind can face selection pressures for multiple traits. As selected by evolution, the nose serves both the function of respiration and the function of smelling.²¹ Similarly, a *bureaucratic form* might persist in an organisation because it serves a computational process and because it creates a paper trail if the management seeks to fire someone.

An SCS-kind might also have additional projected functions, as suggested by Hindriks and Guala. For example, the SCS-kind *limited liability corporation* might also have the projected function of shielding owners from liability. While a limited liability corporation might fulfill its SCS-function by fulfilling its projected function, they are distinct.

Supporting social computations

In a nutshell, the function of SCS-kinds is to support social computations. I have already addressed one question this characterisation raises, namely what the nature of social computations is. There are, however, two more questions that deserve to be addressed. First, what does it mean to *support* social computation? Second, do the examples I have used really fit this description?

The analysis suggests that to support a social computation is to make it more powerful or efficient. For one social computation to be more powerful than another means that it achieves a result, a computational end state, that the other cannot achieve. For example, consider a group of administrators who fail to supervise an organisational process until a filing system allows them to implement a memory-intensive process that would otherwise be beyond their capacities. The system of organisation enables them to achieve a result that would otherwise be unreachable.

For a computational process to be more efficient means that it achieves a result while requiring less computational resources than another.²² For an example, consider the transaction cost theory

²⁰The one employee might change over time while the business and the *customer acquisition form* remains the same, ensuring that the company is not identical with the individual.

²¹One might seek to draw various distinctions between these functions, for example Millikan's (1984) distinction between direct proper and adapted functions might apply. For the present purposes, these distinctions are not decisive because all these types of functions are sufficient to fulfill the conditions for being an SCS-kind.

²²If one considers the resources used by the computational process as input, then one can also describe the process as more powerful. It can do more with less input.

of the firm (Coase 1937). This theory is a response to the fact that market economies are populated by firms which are not organised by market principles, but by institutional organisation and command. If contract-based market economics generally outperform command economies, as both introductory economics courses and history seem to show, then why are companies not outperformed and driven off the market by individuals linked by contracts that do not establish command relationships?

According to the transaction-cost approach, firms are established because the efficiency costs of external transactions consisting in a series of contracts exceed those of internal command transactions. If we understand arranging exchange as a computational process, the transaction-cost theory of the firm describes the circumstances in which various SCS-kinds associated with firms fulfill their SCS-function.

The support criterion in the analysis of SCS-kinds is relative to other social processes. A well-functioning SCS-kind instance improves the computation relative to processes lacking it. The contribution can be either constitutive or causal, i.e., the instances of the social kind can be constitutive to the process or cause it to be more efficient/powerful. The condition is that one way or another the representation of instances of the SCS-kind creates a difference from the comparison class of social processes lacking appropriate representations.

One might be surprised that the analysis makes reference to multiple alternative social processes rather than just the most similar one. Why explore multiple possible worlds rather than the closest one fitting the condition? The reason is that the most similar process, *Q*, might include another social kind that would make it more powerful or efficient. In this case, one might say that the kind still lives largely up to its SCS-function as long as it helps to outperform all the other comparison processes *R*, *S*, *T* The comparison with just one other social process would hide this partial enabling. In other words, this formulation allows for an easy interpretation of how the instances of an SCS-kind might partially succeed.

Note also that the analysis selects the most similar *social* processes, not *computational* processes, for comparison. The processes relevant for comparison might not be computational. What matters for the functioning of an SCS-kind is that it actually improves the processes of the social group, not that it does so over other computational processes. If computational process are in general an ineffective way of solving a problem, then an SCS-kind cannot function well.

Having clarified what it means to support a social computation, I turn to whether my examples fit this description. Establishing that *departmental reimbursement form*, *US dollar bill*, and *chair of the board* support social computations might be the greatest challenge for showing that they are SCS-kinds. Most will grant that they have the function to support social processes, but are these processes computational? While I will not settle on a specific account of physical computation, a few general considerations about these key cases are in place.

The simplest of the cases is the *departmental reimbursement form*. Clearly the instances of this kind support an organisational process if they function well. There is also little doubt that they are selected to play a role in a group's information processing. The reimbursement forms serve to summarise the required information for the functions of the department. They provide an input and transformation mechanism of relevant information into the organisational system.

Only an extremely restrictive analysis of physical computation, more demanding than some of the most restrictive proposals (e.g., Piccinini 2015; see also Strohmaier, manuscript), would rule out these processes. They manipulate information vehicles according to rules to arrive at an output. Hence, the departmental reimbursement form is the most straightforward example of an SCS-kind.

US dollar bills are more challenging because there is not a single organisation, a distinct and formalised group, they are supposed to serve. Their function is not just to serve the processes inside the US Federal Reserve, but to circulate and have effects on computations beyond these organisations. The departmental reimbursement forms are specific to a department, but the presence of the

bills benefits the whole economy.²³ Having currency bills available makes market processes run more smoothly and that is why the government provides them.

However, the requirement for fulfilling the SCS-function is not to support a social computation of a specific organisation. Instead of an organisation, the group implementing the social computation will consist of market participants. As long as we conceive of the allocation processes implemented by this group in computational terms, *US dollar bill* will be an SCS-kind.²⁴

A chair of a board plays a specific role in organisational processes. The kind is, thus, closer to *departmental reimbursement form* than *US dollar bill*. A chair of a board is supposed to contribute to organisational decisions, in the limiting case imposing them. These decision processes are excellent candidates for social computations since they involve the manipulation of information, often encoded in strings of natural language, and are supposed to produce an output, the decision of the organisation.

Before moving to the next condition, a few words are in place on the power of SCS-kinds—that is, on their effectiveness in social processes. According to Searle, institutional kinds have deontological power—they are supposed to provide desire-independent reasons, but no such normativity and no such reasons are required for SCS-kinds. Instead, SCS-kinds are characterised by their power of enabling computational social process by being represented.²⁵

This perspective also differs from that of Guala and Hindriks. They understand the power of institutional kinds as deriving from their effects on the equilibria of games. More specifically, institutional kinds are based on behavioural rules that are supposed to lead to increased payoffs. This view does not require special desire-independent reasons and instead relies on the availability of rules that can change the game. Significantly, however, the social processes are not understood in computational terms. The power is not described as stemming from kinds' role in social computations.

The types of power proposed in the literature are not mutually exclusive, SCS-kinds could also have deontological powers and shift equilibria, but any overlap is contingent. The differences show that identifying this kind of kinds captures a hitherto neglected position in the debate on social kinds.²⁶

Being represented by realising agents

The analysis of SCS-kinds ascribes to them as part of their function that their instances be represented as belonging to the kind by realising agents. The departmental reimbursement forms only fulfill their SCS-function if they improve social computations by being represented as belonging to the kind *departmental reimbursement form*. The social computations are more powerful or efficient in virtue of instances of an SCS-kind being represented as belonging to the respective first-order kind.²⁷

Why is it not sufficient for the instances to support social computation without any representation as belonging to some kind? Consider a more extensive type of kind, let's call them SCS*-kinds, where

²³One might argue that this is no longer the case for bills because they are increasingly displaced by electronic means of exchange. Even if this were the case, *US dollar bill* might still be an SCS-kind, just one that no longer fulfills its functions. If the kind were to lose the function altogether, the example should be considered historical.

²⁴Economists of the Austrian persuasion (e.g., Hayek 1948) have a tendency to take such a computational perspective when they discuss the market process in terms of knowledge or information being used in decentralized planning. Other authors have argued that “[m]oney is technologically equivalent to a primitive version of memory” (Kocherlakota 1998, 250) or serves to implement a record protocol (Smit, Buekens, and Plessis 2016).

²⁵A related tradition links the existence of institutions to cognitive limits of human agents (cf. Mantzavinos 2011). While my proposal does not imply that all SCS-kinds serve to overcome such limitations, some might help to overcome motivational rather than cognitive limitations, it is likely the case for many.

²⁶As mentioned in the discussion of social computation, Gallagher and Crisafi 2009 making a closely related proposal that does not address the issues specific to the debate on social kinds.

²⁷The “in virtue” is to be construed broadly to include all explanatory relations, causation as well as grounding.

a social kind, *K*, is an SCS*-kind if and only if instances of *K* have a function that for a social computational process *P*:

- (2*) *P* is computationally more powerful or efficient than the social processes *Q*, *R*, ... ,
- (3*) where *Q*, *R*, ... are the most similar possible social processes which do not involve an instance of *K*.

These conditions do not require any realising agents. However useful the concept of such SCS*-kinds, one should realise how vastly more expansive it is.

For a start, this analysis includes all kinds of artifacts which have the function of supporting social computations, including *pen* and *paper*.²⁸ Clearly having a pen allows us to engage in social computations that would otherwise require larger human effort or would be infeasible altogether. The research teams of the Manhattan Project would agree that their calculations relied on pen and paper which thereby fulfilled their functions, but this does not make them SCS-kinds because the representation of *pen* instances as belonging to this kind does not make the difference to the computation. It is only their ability to be used for writing.

Compare this with the case of *departmental reimbursement form*, which is also an artifact kind. In contrast to *pen* and *paper*, it fulfills its function by agents thinking about its instances as belonging to this kind. The agents in the department keep track of resources by thinking about the sheet of paper with fields as being a departmental reimbursement form, and by filling it out and reading it accordingly. The organisation's information processing relies on the members treating these forms as belonging to this specific kind shared between them, not just as notepaper. As in the case of an object-class in a digital software programme, being represented as belonging to the kind is what makes the difference.

The role of representation also explains why SCS-kinds should not be extended to the kind *human brain*.²⁹ The human brain clearly also plays a crucial role in enabling social computations, but *human brain* is not an SCS-kind. It is not a function of brains to enable computational processes by being represented as belonging to the kind *brain*. The distinguishing contribution of SCS-kinds is dependent on being represented by agents realising a social computation, which is not the case for brains.

As the examples show, the representation by the realising agents is the crucial link to cognition, which also distinguishes the contribution of these kinds to social computations. Such a link is present in the case of the *departmental reimbursement form* but not the *pen* and not of the right type in the case of *brain*. The representation by agents makes the difference in the comparison with other contributions to computational processes. Thus, SCS-kinds have their place in social cognition, although the theory remains neutral in debates on the extended mind and group agency.

Given this connection to cognition, the question of mind-dependence arises. Following Khalidi, one might ask:

- (1) whether it is necessary for an SCS-kind to exist that individuals have a propositional attitude toward it, and
- (2) whether it is necessary for the instance of an SCS-kind to exist that individuals have a propositional attitude toward this instance.

Since it is easier to answer, I will begin with the second question regarding instances. The selectionist analysis of the function of SCS-kinds does not establish the mind-dependence of their

²⁸I assume that such artifact kinds are social kinds. This treatment is not supposed to imply that they would be institutional kinds in the sense of Searle or Guala and Hindriks. As discussed, institutional kinds form a subset of social kinds.

²⁹On most conceptions, *human brain* is not a social kind and therefore is not covered by the present analysis, but what is at stake is whether *human brain* would otherwise fulfill the conditions for being an SCS-kind.

instances, because not all instances of an SCS-kind need to fulfill their function. Assume that someone prints a departmental reimbursement form by accident and never even looks at what they have produced. The piece of paper might be thrown out without anyone representing it as the SCS-kind to which it nevertheless belongs.³⁰

It is not necessary for an SCS-kind that for all its instances someone represents it as belonging to this first-order kind. It is necessary, however, that at least one instance has been represented this way since the instances have been selected for this trait.³¹ This brings us to the first, and more challenging, question of whether the existence of an SCS-kind necessarily depends on propositional attitudes toward it.

Following the analysis, the SCS-function requires the representation of some instance as belonging to the kind. To be selected for, the function has to be fulfilled at least on one occasion, otherwise the kind could not persist in virtue of this trait. Although selectionist functions do not generally require any propositional attitudes, the specific trait that is selected for in this case does. Thus, for a social kind to be an SCS-kind, it is necessary that at least one realising agent has a propositional attitude toward the first-order kind. They have to represent an instance as belonging to this social kind.

While for a social kind *K* to be an SCS-kind it is necessary that there existed a propositional attitude toward *K*, this does not directly imply that the existence of *K* itself depends on such an attitude. This further dependence, which is the target of Khalidi's first question, relies on the assumption that *K* is necessarily an SCS-kind. For example, if *US dollar bill* is necessarily an SCS-kind, then it depends on propositional attitudes toward it as discussed by Khalidi.³² Whether any social kind is necessarily an SCS-kind is a question the present analysis leaves open to future investigation.

The significance of SCS-kinds

I have put forward SCS-kinds as an alternative kind of social kinds, competing with theories of institutional kinds, but isolating SCS-kinds does not only serve to satisfy our curiosity about what makes the social realm special. SCS-kinds have an important place in explaining the achievements and shortcomings of social arrangements. I have already mentioned the theory of the firm as one example in economics. The theory of SCS-kinds is of wide-reaching significance for human society.

Previous research into social kinds either focusses on what makes them metaphysically peculiar (e.g., Searle 1995, 2010) or on their role for achieving social justice (e.g., Haslanger 1995, 2000, 2005). Independently of the value these projects have, they only cover part of the impact of social kinds and SCS-kinds in particular. The historical growth of societal capacities has its roots at least partially in SCS-kinds innovation. SCS-kinds, such as various types of forms, are a key component of bureaucratic organisation and thereby explain increased state and firm capacities.

Future research should provide the basis for modelling the societal impact of SCS-kinds. Their computational nature should make such a project easier. Despite the differences between digital computations and social computations, building an approximate computational model of the latter should be feasible.³³ Importantly, the theory of SCS-kinds justifies an interpretation of these

³⁰Searle's (1995, 31) story of a US dollar bill which falls into a crack after being printed is another example of an SCS-kind instance not being mind-dependent.

³¹For a subset of SCS-kinds, however, instances depend in their existence on propositional attitudes toward them. For a judgment to be a US Supreme Court judgment requires that the judges have an attitude toward it.

³²This mind-dependence introduces the threat of a circularity, but this threat is no stronger than for any other case that Khalidi discusses and can be averted by denying that any kind is necessarily an SCS-kind.

³³Assuming the Church-Turing thesis (cf. Copeland 2017), the exact same computations should be possible. The models might of course be computationally prohibitive.

computational models as directly reflecting the computational aspects of social reality. Social ontology and computational social science meet in the analysis of SCS-kinds.

Research does not have to stop at analysing the current state of SCS-kinds, it can seek to improve upon them. Even well-functioning SCS-kinds rarely optimise the power of social computations because they only outperform the most similar social processes. The lack of optimisation opens a novel approach to research of great significance for human progress. Which innovation in social kinds could increase the computational power of our social arrangements? Which SCS-kinds could make processes in organisations computationally more efficient, and would their adoption free resources? Such questions are not solely the domain of economics and organisational theory. By investigating SCS-kinds and how they are to be modelled, philosophers can contribute to the well-functioning of social processes.

For example, philosophers could investigate whether design patterns from software design are applicable to SCS-kinds. Can we improve social processes by drawing on best practices in object-oriented programming for the creation of social kinds? Can we group social kinds in similar ways as software engineers have grouped styles of programming? Philosophers with awareness of the fundamental differences between computations as they occur in social groups and in digital devices are especially well placed to evaluate such proposals. If philosophers redirect their investigations into social ontology as suggested, they are well placed to facilitate such cross-disciplinary innovation.

Of course, the functions of social computations can also be investigated with a critical eye (see also Gallagher 2013). A powerful bureaucratic organisation might implement a social computation very well, but at the same time achieve a morally dubious goal or might simply, in virtue of its effectiveness, be a cage to human individuality along the lines of Max Weber's ([1921–1922] 2013) iron cage of modernity. The actions of individuals become minuscule contributions to an overall computational process and are thereby reduced to subprocesses. Especially by drawing on the science and technology studies, research into SCS-kinds can shed light on these issues.

One investigation that appears especially pressing in the current moment is the effect of introducing digital devices into social groups. Understanding the computational nature of SCS-kinds can help to see how these changes reinforce Weber's worries. While SCS-kinds used to be employed by human agents, their usage is increasingly tracked and enforced by digital devices. No longer can the accounting department gloss over an empty field in a departmental reimbursement form because now the accounting software requires an entry. As a consequence, the social computations SCS-kinds support become more restrictive. With help of the theory of SCS-kinds, philosophers can investigate how to redesign these kinds to promote human autonomy in social processes.

That SCS-kinds lie at the touching point of multiple inquiries promises renewed relevance to social ontology. The significance of these kinds is interdisciplinary. Research into them will require methodological innovation, at least in applying methods from other fields to social ontology. Which formalisms for describing digital computations are applicable to describe roles in organisations? Where do the abstractions fail? New research directions and methodological challenges are waiting for those willing to view social kinds from a computational perspective.

Conclusion

Because it has neglected the role of social computations, the debate on social and institutional kinds has failed to acknowledge an important kind of social kinds. Departmental reimbursement forms, US dollar bills, and chairs of the board have the function to support social computations by being represented as belonging to their kind. They are instances of social-computation-supporting kinds.

SCS-kinds make a distinct contribution to processes realised by social groups. Recognising this contribution opens new roads of inquiry to social ontology. If they use this opportunity, philosophers will be at the forefront of innovating social kinds.

Acknowledgments. I would like to thank the anonymous referees, Brian Epstein, Glady Tyen, Simon Wimmer, and the audience members of my talk at the 2019 Social Ontology conference for helpful comments and discussions. This paper reports on research supported by Cambridge Assessment, University of Cambridge.

David Strohmaier is a research associate in the department of computer science and technology at the University of Cambridge. His work crosses the disciplinary boundaries between philosophy, computer science, and the social sciences.

References

- Adams, Frederick, and Kenneth Aizawa. 2008. *The Bounds of Cognition*. Malden, MA: Blackwell.
- Boyd, Robert, and Peter J. Richerson. 2005. *The Origin and Evolution of Cultures*. Oxford: Oxford University Press.
- Chalmers, David J. 1996. "Does a Rock Implement Every Finite-State Automaton?" *Synthese* 108 (3): 309–33. <https://doi.org/10.1007/BF00413692>.
- Cioffi-Revilla, Claudio. 2017. *Introduction to Computational Social Science*. London: Springer. <https://doi.org/10.1007/978-3-319-50131-4>.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19. <https://doi.org/10.1093/analys/58.1.7>.
- Coase, Ronald Harry. 1937. "The Nature of the Firm." *Economica* 4 (16): 386–405. <https://doi.org/10.1111/j.1468-0335.1937.tb00002.x>.
- Coelho Mollo, Dimitri. 2018. "Functional Individuation, Mechanistic Implementation: The Proper Way of Seeing the Mechanistic View of Concrete Computation." *Synthese* 195 (8): 3477–97. <https://doi.org/10.1007/s11229-017-1380-5>.
- Copeland, B. Jack. 2017. "The Church-Turing Thesis." In *The Stanford Encyclopedia of Philosophy* (Spring 2019), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2019/entries/church-turing>.
- Dennett, Daniel Clement. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Epstein, Brian. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. New York: Oxford University Press.
- Fodor, Jerry Alan. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fodor, Jerry Alan. 1981. "The Mind-Body Problem." *Scientific American* 244 (1): 114–23.
- Friedman, Milton. 1994. *Money Mischief: Episodes in Monetary History*. San Diego: Harcourt Brace.
- Gallagher, Shaun. 2013. "The Socially Extended Mind." *Cognitive Systems Research* 25–26 (December): 4–12. <https://doi.org/10.1016/j.cogsys.2013.03.008>.
- Gallagher, Shaun. 2018. "The Extended Mind: State of the Question." *The Southern Journal of Philosophy* 56 (4): 421–47. <https://doi.org/10.1111/sjp.12308>.
- Gallagher, Shaun, and Anthony Crisafi. 2009. "Mental Institutions." *Topoi* 28 (1): 45–51. <https://doi.org/10.1007/s11245-008-9045-0>.
- Garson, Justin. 2011. "Selected Effects and Causal Role Functions in the Brain: The Case for an Etiological Approach to Neuroscience." *Biology and Philosophy* 26 (4): 547–65. <https://doi.org/10.1007/s10539-011-9262-6>.
- Garson, Justin. 2012. "Function, Selection, and Construction in the Brain." *Synthese* 189 (3): 451–81. <https://doi.org/10.1007/s11229-012-0122-y>.
- Garson, Justin. 2017. "A Generalized Selected Effects Theory of Function." *Philosophy of Science* 84 (3): 523–43. <https://doi.org/10.1086/692146>.
- Gilbert, Margaret. 1992. *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gilbert, Margaret. 2013. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.
- Guala, Francesco, and Frank Hindriks. 2015. "A Unified Social Ontology." *Philosophical Quarterly* 65 (259): 177–201. <https://doi.org/10.1093/pq/pqu072>.
- Guala, Francesco. 2016. *Understanding Institutions: The Science and Philosophy of Living Together*. Princeton, NJ: Princeton University Press.
- Haslanger, Sally. 1995. "Ontology and Social Construction." *Philosophical Topics* 23 (2): 95–125.
- Haslanger, Sally. 2000. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55.
- Haslanger, Sally. 2005. "What Are We Talking About? The Semantics and Politics of Social Kinds." *Hypatia* 20 (4): 10–26.
- Hayek, Friedrich August. 1948. *Individualism and Economic Order*. Chicago: University of Chicago Press.
- Hindriks, Frank, and Francesco Guala. 2019. "The Functions of Institutions: Etiology and Teleology." *Synthese*. <https://doi.org/10.1007/s11229-019-02188-8>.
- Huebner, Bryce. 2014. *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. New York: Oxford University Press.
- Huebner, Bryce, and Joseph Jebari. 2019. "Computational Theories of Group Behavior." In *The Routledge Handbook of the Computational Mind*, edited by Mark Sprevak and Matteo Colombo, 483–96. New York: Routledge.
- Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Khalidi, Muhammad Ali. 2015. "Three Kinds of Social Kinds." *Philosophy and Phenomenological Research* 90 (1): 96–112. <https://doi.org/10.1111/phpr.12020>.

- Kocherlakota, Narayana R. 1998. "Money Is Memory." *Journal of Economic Theory* 81 (2): 232–51. <https://doi.org/10.1006/jeth.1997.2357>.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- March, James G., and Herbert A. Simon. 1958. *Organizations*. New York: Wiley.
- Mason, Rebecca. 2016. "The Metaphysics of Social Kinds." *Philosophy Compass* 11 (12): 841–50. <https://doi.org/10.1111/phc3.12381>
- Mantzavinos, Chrysostomos. 2011. Institutions. In *The SAGE Handbook of the Philosophy of Science*, edited by Jesús Zamora-Bonilla and Ian C. Jarvie, 399–412. London: Sage Publications.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Millikan, Ruth Garrett. 2014. "Deflating Socially Constructed Objects: What Thoughts Do to the World." In *Perspectives on Social Ontology and Social Cognition*, edited by Mattia Gallotti and John Michael, 27–39. New York: Springer.
- Piccinini, Gualtiero. 2015. *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- Piccinini, Gualtiero. 2017. Computation in Physical Systems. In *The Stanford Encyclopedia of Philosophy* (Summer), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2017/entries/computation-physicalsystems>.
- Putnam, Hilary. 1988. *Representation and Reality*. Cambridge, MA.: MIT Press.
- Ritchie, Katherine. 2020. "Social Structures and the Ontology of Social Groups." *Philosophy and Phenomenological Research* 100 (2): 402–24.
- Searle, John. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.
- Searle, John. 1995. *The Construction of Social Reality*. New York: Free Press.
- Searle, John. 2010. *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.
- Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- Smit, J. P., Filip Buekens, and Stan Du Plessis. 2016. Cigarettes, Dollars and Bitcoins: An Essay on the Ontology of Money. *Journal of Institutional Economics* 12 (2): 327–47. <https://doi.org/10.1017/S1744137415000405>.
- Smith, Barry. 2003. "John Searle: From Speech Acts to Social Reality." In *John Searle*, edited by B. Smith, 1–33. Cambridge: Cambridge University Press.
- Strohmaier, David. Manuscript. "Organisations as Computing Systems."
- Thomasson, Amie L. 2003. "Foundation for a Social Ontology." *ProtoSociology* 18: 269–90. <https://doi.org/10.5840/protosociology200318/199>.
- Tollefsen, Deborah Perron. 2006. "From Extended Mind to Collective Mind." *Cognitive Systems Research* 7 (2–3): 140–50. <https://doi.org/10.1016/j.cogsys.2006.01.001>.
- Tollefsen, Deborah Perron. 2015. *Groups as Agents*. Malden, MA: Polity.
- Weber, Max. 2013. *Die protestantische Ethik und der Geist des Kapitalismus: Vollständige Ausgabe*. München: C. H. Beck.
- Wimsatt, William C. 1972. "Teleology and the Logical Structure of Function Statements." *Studies in History and Philosophy of Science Part A* 3 (1): 1–80. [https://doi.org/10.1016/0039-3681\(72\)90014-3](https://doi.org/10.1016/0039-3681(72)90014-3).
- Wright, Larry. 1973. "Functions." *Philosophical Review* 82 (2): 139–68. <https://doi.org/10.2307/2183766>.
- Wright, Larry. 1976. *Teleological Explanations. An Etiological Analysis of Goals and Functions*. Berkeley: University of California Press.