Epistemic institutions

A joint epistemic action-based account

Miller, Seumas

DOI
10.1111/phis.12227

Publication date
2022

Document Version
Final published version

Published in
Philosophical Issues

Citation (APA)
Miller, S. (2022). Epistemic institutions: A joint epistemic action-based account. *Philosophical Issues*, *32*(1), 398-416. https://doi.org/10.1111/phis.12227

DOI: 10.1111/phis.12227

ORIGINAL ARTICLE

Philosophical Issues
A Supplement to NOÛS

Epistemic institutions: A joint epistemic action-based account

Seumas Miller[1,2,3]

[1]Charles Sturt University, Canberra, Australia

[2]University of Oxford

[3]TU Delft

Correspondence
Charles Sturt University, Canberra, Australia
Email: semiller@csu.edu.au

1 | INTRODUCTION

Contemporary social institutions include complex organizations, or systems of organizations such as governments, police services, business corporations, universities, welfare institutions and the like; they also include, criminal justice systems (comprised of a police organization, courts, correctional facilities etc.), legal systems (comprised of a legislature, the law, courts, legal firms etc.), financial systems (comprised of retail and investment banks, a stock exchange, regulators, auditing firms etc.) and so on. Accordingly, on the one hand, there is a need for a *general* theory of social institutions and, on the other hand, a need for *special* theories of particular institutions, e.g., a theory of universities. So far so good. However, there is a further distinction that needs to be accommodated, namely, that between epistemic institutions and non-epistemic institutions. The raison d'etre and core business (so to speak) of some institutions is epistemic, e.g., knowledge acquisition and dissemination. Epistemic institutions include universities, news media organizations and intelligence agencies (of which more below).

Issues that need to be addressed in relation to epistemic institutions include: (1) An outline of one's favored general theory of social institutions, given that epistemic institutions are a species of social institution; in the case of this article, an outline of a joint action-based teleological theory (Section 2); (2) An analysis of the core constitutive concept(s) upon which the favored more specific theory of epistemic institutions is to be based and the relationship of this concept(s) to that of an epistemic institution; in the case of the joint action-based teleological theory, the core concept in question is joint epistemic action (Section 3); (3) An elaboration of the distinction between epistemic and non-epistemic institutions, but also of distinctions between different sub-

species of epistemic institution and in particular, in this article, between universities, news media organizations and intelligence agencies (Section 4).

It should be noted that the theories in question while they are descriptive in that they need to be anchored in contemporary social reality, they are also normative and, in the case of epistemic institutions, normative in that knowledge (and its cognates) are (typically) desirable human goods.

2 | SOCIAL INSTITUTIONS

As indicated above, our concern in this article is principally with institutions understood as organisations or systems of organisations, as opposed to less complex social forms, such as conventions or social norms, on the one hand, and more complex and complete social forms, such as societies or nation-states, on the other hand. The term "institution" tends to be used to refer to those organisations, or systems of organisations, that play a central and important role in a society and that, therefore, reproduce themselves. Accordingly, in this article we understand institutions to be organisations, or systems of organisations, that reproduce themselves.

Roughly speaking, an institution (in this sense) consists of an embodied (occupied by human persons) structure of differentiated roles (Ludwig, 2017; Miller, 2010, Ch. 2)—as well as, typically, additional non-human components, e.g., buildings and other artefacts. These roles are defined in terms of tasks, and rules (including conventions and social norms, as well as explicit laws and regulations) governing the performance of those tasks. Moreover, there is a degree of interdependence among these roles, such that the performance of the constitutive tasks of one role cannot be undertaken or, at least, cannot be undertaken except with great difficulty or considerable inefficiency, unless the tasks constitutive of some other role or roles in the structure have been undertaken or are being undertaken. Further, these roles are often related to one another hierarchically, and hence involve different levels of status and degrees of authority. Finally, on teleological and functional accounts and, in particular, on the joint action-based teleological account favoured here, these roles are related to one another in part in virtue of their contribution to the ends or functions of the institution; and the realisation of these ends or functions normally involves interaction not only between the members of the institution in question, but also between these internal institutional actors and external actors who might be members of other institutions or, alternatively, might be non-institutional actors.[i] Thus police detectives interact with one another but also with judges and members of the public.

An important distinction to be kept in mind here and in what follows is that between what (or who) is constitutive of an institution, (e.g., the academic members of a university and their students performing epistemic actions qua academics and qua students (respectively)), and what might be causally required to maintain the institution in existence, (e.g., a renewal of the university's lease of land agreement with a government authority). This distinction is less clear-cut in some instances than others (which is not, of course, a sufficient reason for abandoning the distinction). Thus, acceptance by members of other institutions, such as professional groups, that the degrees issued by a university are evidence of a certain epistemic attainment, might well be a necessary causal condition for the university's continued existence (say, for financial reasons) but, nevertheless, such collective acceptance might not necessarily be regarded as constitutive of the university as an institution. On the other hand, if very few external institutions or professional groups accepted that the degrees issued by the university were evidence of a certain epistemic attainment then the status of the organisation as a university might be called into

question, notwithstanding that it continued to be well-funded and attract an adequate, if small, quantum of students. Further, if *in fact* the education on offer in the organisation was of high quality (and it was well-funded and continued to attract an adequate, if small, quantum of students) then our inclination might very well be to regard it as a university, notwithstanding that a majority of external institutions and professional groups refused to accept its degrees as evidence of a certain attainment. Accordingly, even if collective acceptance is a necessary causal condition for the continued existence of an institution it might not be a necessary constitutive feature of that institution. Constitutive features and causal conditions should not be confused, and collective acceptance is not necessarily a constitutive feature of an institution, let alone a necessary and *sufficient* constitutive feature of an institution (in our organisational sense of institution).

Acceptance, and collective acceptance in particular, is a constitutive feature of authority relationships rather than necessarily of social institutions per se (contrary[ii] to collective acceptance theories (Ludwig, 2017; Searle, 1995; Tuomela, 2002)). Thus, if most of those with respect to whom a would-be authority seeks to exercise his or her authority do not accept their authority, then the would-be authority does not in fact have the authority in question (even if, as sometimes happens in the case of governments, the would-be authority has power, and even if, as sometimes happens in the case of epistemic authorities, the would-be authority has the formal status of an (epistemic) authority). This is consistent with an individual or group maintaining their authority, notwithstanding that *some* individuals under their authority refuse to accept their authority. For instance, a criminal may refuse to accept the authority of the police. In itself, this does not extinguish the authority of the police. However, if the members of the relevant community in general refused to accept the authority of the police, then this would extinguish their authority. It is also consistent with an authority who is accepted by the relevant community (including institutional subordinates) as such, *occasionally* successfully authorising a person or action which act of authorisation the community does not accept, e.g., a police chief authorising a police spokesperson who neither the citizenry nor subordinate police officers accept (to use one of Jenny Lackey's examples (Lackey, 2018, p. 24)). Naturally, since most social institutions are constituted in part by authority relationships then collective acceptance will be a constitutive feature of most social institutions simply by virtue of these authority relationships being a constitutive feature of the institutions in question.

The basic building block of our favoured teleological theory of social institutions is the concept of a joint action. Roughly speaking, a joint action is performed by multiple agents, each of whom performs a constitutive individual action, and each of whom in doing so is aiming to achieve a common goal (Bratman, 2014; Ludwig, 2016; Miller, 1992, 2001, Ch. 2, 2007a, 2010, Ch. 1.). These constitutive individual actions are interdependent by virtue of the common goal since, other things being equal, this goal will not be achieved unless each or most perform their contributory action. Examples of joint action are: two people dancing together, a number of tradesmen building a house and a team of researchers conducting an attitudinal survey. Joint action is to be distinguished from individual action on the one hand, and from the 'actions' of corporate bodies on the other. Thus, an individual walking down the road or shooting at a target are instances of individual action. A nation declaring war or a government taking legal action against a business corporation are instances of *corporate* action (French, 1984; Gilbert, 1989; Miller & Makela, 2005).

So joint actions are interdependent actions directed toward a common goal or end (collective end, in our parlance). Moreover, the notion of a collective end is a construction out of the prior notion of an individual end. Roughly speaking, a collective end is an individual end more than one agent has, and which is such that, if it is realised, it is realised by all, or most, of the actions of the agents involved; the individual action of any given agent is only part of the means by which

the end is realised[iii]. The realisation of the collective end is the bringing into existence of a state of affairs. Each agent has this state of affairs as an individual end. (It is also a state of affairs which may be somewhat underspecified, aimed at under more or less the same description by each agent or, at least, the agents have overlapping individual ends such that there is a possible state of affairs which would adequately satisfy each of these individual ends.) So a collective end is a species of individual end (Miller, 1992, 2001, pp. 56–71, 2007a, 2010, pp. 41–46). Accordingly, this analysis of joint action is individualist[iv].

Note that one important species of joint action is joint assertion. A detailed analysis of the practice of assertion is provided below. Here let us briefly consider Jennifer Lackey's examples of, firstly, what she refers to as a co-ordinated group assertion and, secondly, what she refers to as authority-based group assertion (Lackey, 2018, p. 22). A tour group stranded on a desert island jointly write in the sand the words, "We need help" (an example of Lackey's co-ordinated group assertion). The collective end account of joint action provides an individualist analysis of this example that is contrary to Lackey's group assertion analysis. The analysis based on the collective end account is that each member of the tour group asserts that we (the members of the tour group) need help and asserts this by means (in part) of a jointly produced utterance (the writing in the sand); but each asserts this interdependently with the others doing so and does so having as a collective end (let us assume) that a passing plane notices the message and comes to their rescue. Similarly, (to use another of Lackey's examples of coordinated group assertion) if a research group worked together to formulate a sentence expressing an assertion or to determine the content of an assertion, prior to the act of assertion, but each attached their name to the sentence, then each member is asserting the content of the sentence, but doing so jointly and, therefore, interdependently in the service of some collective end, e.g. to display their solidarity on the issue in question, to give the assertion greater weight than the same assertion performed by a single member might have, or simply to indicate (as Lackey seems to be suggesting) that they all had a role in the formulation of the sentence or the determination of its content.

Lackey's notion of authority-based group assertion essentially concerns spokespersons for groups. Roughly, her idea is that a group per se asserts something through its spokesperson. and the spokesperson in saying things qua spokesperson does not assert anything. Contra Lackey, the spokesperson does, I suggest, perform assertions while neither the group nor its members necessarily assert anything. Specifically, the spokesperson asserts things about the group per se or its members, including that it, or they, assert (let us assume) p or otherwise commit themselves to p. However, there is a convention in force such that under certain conditions if the spokesperson asserts that the group or its members have asserted that p or have otherwise commit themselves to p, then the group or its members can justifiably be treated *as if* they have asserted p or otherwise committed themselves to p. But from this it does not follow that the group or its members have *in fact* asserted p or have *in fact* otherwise committed themselves to p. They might have; but they might not have. Specifically, this notion of justifiably being treated *as if* one has asserted that p does not entail that one has in fact asserted that p or that one is believed to have asserted that p, although it does entail something weaker, such as that one can be held morally, and perhaps legally[v], liable for (let us assume) the falsehood that p being disseminated and for the untoward consequences of this dissemination.

Note also that if the individualist analyses of organisational action (undertaken below) in terms of multi-layered structures of joint action, joint mechanisms and chains of institutional action are successful then there are clear implications for corporate action. Specifically, it is likely that it will turn out to be unnecessary to invoke irreducibly collectivist notions of corporate action.[vi]

Note further that in the case of institutions and, in particular, epistemic institutions, the collective end is or, at least, ought to be, a human good which is also a collective good in some sense (but not necessarily in the economists' sense of a non-rival and non-excludable good). In the case of epistemic institutions, a strong candidate for such a collective good is collective knowledge of which more below (Goldman & Whitcomb, 2011; Schmitt, 1994; Smith, 1982).

2.1 | Multi-layered structures of joint action

Organizational action typically consists in a multi-layered structure of joint actions (Miller, 1992, 2001, Ch. 5, 2010, Chs. 1 & 2). One relevant illustration of the notion of a layered structure of joint actions is a cybersecurity department comprised of three (let us assume for purposes of simplification) cyber teams: a cyber threat intelligence team (TI); an incident response team (IR), and an engineering team (EN). Suppose at an organizational level a number of joint actions ('actions') are severally necessary[vii] and jointly sufficient to achieve some collective end, e.g., to prevent or mitigate malware attacks. Thus, the joint epistemic action of the TI team gives early warning to the IR team (which can act to prevent or, let us assume in this instance, mitigate a cyberattack) and, if necessary (as we assume it is in this instance), to the EN to enable it to 'patch' a defect in the system which the cyberattack is exploiting. Assume that the 'action' of TI is, in fact, a joint action, as is the 'action' of IR and the 'action' of EN. Moreover, assume also that the 'action' of TI, the 'action' of IR, and the 'action' of EN are severally necessary and jointly sufficient to achieve the collective end of preventing or mitigating the ongoing cyberattack, e.g., a virus; as such, these 'actions' taken together constitute a fourth joint action which is comprised of the three joint actions of TI, IR and ED (respectively).

At the first level there are individual actions directed to three distinct collective ends: the collective ends of (respectively) collecting and disseminating cyber threat intelligence, responding to the cyberattack, and removing the cyber system vulnerability. Thus, at this level there are three joint actions (of TI (a joint epistemic action), IR, and ED, respectively). However, taken together these three joint actions constitute a single (second level) joint action. The collective end of this second level joint action is to mitigate the effects of the ongoing cyberattack; and from the perspective of this second level joint action, and its collective end, these (first level joint) constitutive actions are (second level) individual actions.

Note that typically in organizations not just the nature, but also the extent, of the individual contributions made to the collective end will differ from one team member to another. Note also that (as mentioned above) the collective end of the organisation (and of particular joint actions) will exist in the minds of the participants under different descriptions; indeed, in some instances it might be more accurately characterised as a set of overlapping individual ends. In addition, as is often the case with long term ends or with the ends of complex actions, the content of these collective ends is initially underspecified and only receives further specification as the joint activity proceeds.

Note finally that here, as elsewhere in institutional arrangements, the role structure within each of the sub-joint actions is maintained in part by the commitment of each or most of the participants to the collective end constitutive (respectively) of each of these sub-joint actions. Likewise, there is a need for a coordinating structure comprised in part (let us assume) of a committee consisting of the leadership of the cyber department and of each of the three teams (TI, IR and ED, respectively). This structure and, therefore committee, exists to try to ensure that each of the sub-joint actions do in fact contribute to the larger joint action. Needless to say, without appropriate

ongoing coordination the larger joint action could not be successfully performed. Importantly, this coordination consists in large part in ensuring that the pursuit of the collective ends of each of the sub-joint actions meshes appropriately with the pursuit of the overall collective end of the larger overall joint action. This appropriate meshing relies on three factors in our example, as it does in most cases of complex joint actions conducted in institutional settings. First, and most obviously, there is the above-mentioned coordination committee. Self-evidently, those in charge of each of the sub-joint actions (the leaders of each of the three teams performing the joint actions, TI, IR and ED (respectively)) need to coordinate with one another and with the overall leader of the cyber security department to ensure the pursuit of the collective ends of the sub-joint actions aligns with the larger, overall collective end. Accordingly, all of these leaders (although, especially the leader of the department) need to have as their ultimate collective end the mitigation of the effects of the cyberattack. Second (less obviously), all (or, at least, most) of the participants in the larger overall joint action are aware that their participation in their sub-joint actions is ultimately in the service of the ultimate collective end of (in this case) mitigating the effects of the ongoing cyberattack and, therefore, that they need to willingly (including occasionally as a result of their own discretionary decisions) adjust their contributory individual and sub-joint actions, accordingly, i.e., they need to have adopted the ultimate collective end of the larger joint action (even if this is not often in the forefront of their minds because it does not need to be—rather their focus needs to be on their own individual contributory action and their own local collective ends). It is an illusion to imagine that the actions to be performed by most institutional role occupants can be reduced to mechanically performed tasks under the complete control (at least in principle) of those in authority. Hence, the need for role occupants engaged in the core activity of an institution to understand (at least to some extent) and pursue (even if indirectly and often unconsciously) the collective end(s) of the institution to which they belong. Certainly, this is necessary if an institution is to be successful over time and, in particular, if an *epistemic* institution is to be successful over time (of which more below). Third, there are those not directly participating in the joint action (including its sub-joint actions) who, nevertheless, have a role in ensuring that those directly participating in the joint action realise its collective end and who, therefore, also have as their collective end the mitigation of the effects of the cyberattack. These would include the CEO or other senior manager of the organisation of which the cyber department is a constitutive organisational element. But it might also include the members of external authorities, e.g., cyber security regulators.

Obviously, given the crucial role of institutions and institutional actions in the prevention of cyberattacks, it is important for the purposes in this article that the activity of organizations that are institutions can be understood in purely individualist terms and by recourse to the core notion of joint action (including that of joint epistemic action); hence the significance of the technical notion of a multi-layered structure of joint action.

2.2 | Joint institutional mechanisms

These mechanisms are often embedded in organizations, although this is not necessarily the case. Consider the ubiquitous joint institutional mechanism of voting (Miller, 1992, 2001, Ch. 5, 2010, Chs. 1 & 2, 2018a). There is voting for political office, voting in the cabinet of a parliamentary democracy, shareholder voting in corporations, voting in committees, including voting in relation to cybersecurity measures, etc. Consider, for instance, shareholder voting. Voting rights belong to shareholders and, let us assume, each share gives the shareholder one vote. Shareholder

A exercises their institutional right (if not duty) by casting their vote in an election and A does so only if other shareholders, B, C, D etc. also vote, and only if there is something or someone, e.g., candidates to be directors on the board of directors, to vote for. Thus, in addition to the actions of voting there are the actions of the candidates, X, Y, Z, etc. to be directors. That they stand as candidates is (in part) constitutive of the input to the voting mechanism; after all, voters vote *for candidates*. So there are interlocking and differentiated actions (the inputs). Furthermore, there is some result of the operation of the mechanism: some candidate, say, Smith, is voted in by virtue of having secured the required number of votes (the output). What of the mechanism itself? A key constitutive feature of this voting mechanism is as follows: to receive the required number of votes *is* to be successful in the election.[viii] Importantly, that Smith, in particular, is voted in is not something necessarily aimed at by all of the participants; specifically, those who voted against Smith were not aiming at getting Smith elected. Since we are assuming Smith did in fact receive the required number of votes it follows that those who voted for him have realized the collective end of their joint action. Importantly, there is also a collective end of *all* the voters and *all* the candidates (or at least all those voting and standing for election in good faith). This is the collective end that the those who get the required number of votes—whoever they happen to be—are, thereby, members of the board of directors. This is a collective end of all bona fide participants in the joint institutional mechanism and reflects the commitment of the participants to the above-mentioned key constitutive feature of the institutional mechanism, i.e., that a candidate with the required number of votes is thereby entitled to be a board member. Accordingly, participants in this joint institutional mechanism perform the individual actions of casting a vote and/or standing as a candidate, and they have as a collective end that those who get the required number of votes—whoever they are—are thereby members of the board. So voting is a species of joint action and, more specifically, a joint institutional mechanism. Moreover, *if* each act of voting consists of (an expressed) judgement in the service of a collective epistemic end (e.g., to determine the best candidate) then voting is a species of joint epistemic action (Miller, 2018a).

As we saw in relation to multi-layered structures of joint action, it is important for the purposes in this article that the activity of organizations that are institutions can be understood in purely individualist terms and by recourse to the core notion of joint action; hence the significance of the technical (individualist) notion of a joint institutional mechanism. This notion has particular importance in the context of the influential argument based on the so-called Judgement Aggregation Paradox to the conclusion that groups make reasoned decisions, yet none of the individual members of these groups has individually made these decisions on the basis of a process of individual reasoning and that, therefore, we must acknowledge processes of irreducibly collective reasoning involving irreducibly collective intentions and judgments, and, indeed, irreducibly collective minds possessed of these collective attitudes to engage in this collective reasoning (Copp, 2006; List & Pettit, 2011; Miller, 2007b, 2018a).

2.3 | Chains of institutional action

There is an institutional phenomenon that we refer to as a *chain of institutional action* (Miller, 2014, 2016b, Ch. 5). Chains of institutional action involve, firstly, an extension of the notion of a multi-layered structure of joint action to a *diachronic* series of joint actions directed to a single, ultimate collective end and, secondly, joint institutional mechanisms applied to such diachronic series of joint actions. Moreover, as we saw in the case of many multi-layered structures of joint action and of joint institutional mechanisms, chains of institutional action bring with them

institutional responsibility and, in particular, collective institutional responsibility. Consider a criminal investigation team, including interviewers, cyber forensic officers et al., investigating a major cybercrime. Let us assume that the team is engaging in a joint institutional action, namely, that of determining who is responsible; often a difficult *epistemic* undertaking given the problem of attribution. Moreover, they do so having as a collective end to determine the *factual* guilt or innocence of this and other suspects. At some point these police investigators complete this process and provide a brief of evidence to the prosecutors according to which, and based on all the evidence, the members of an organized cybercrime group (CCG) are the offenders. So far so good, but the criminal justice processes do not terminate in the work of the investigators. For there is now the matter of the trial; that is, the determination by the members of a jury of the legal guilt or innocence of the members of CCG. Let us assume that the members of the jury perform the joint action of deliberating on *legal* guilt or innocence of CCG, and jointly reach the verdict of guilty. The question that now arises concerns the institutional relationship between the joint institutional action of the investigators and the joint institutional action of the members of the jury. It is here that the notion of a chain of institutional action (and responsibility) is illuminating.

Let us assume in what follows that the collective end of the criminal justice process comprised of both the criminal investigators *and* the members of the jury (as well as others, but here we simplify), is that the factually guilty be found legally guilty (and the factually innocent not be found legally guilty). Note that from the perspective of this larger institutional process the collective epistemic end of the investigators (that of determining the factual guilt or innocence of a suspect) is merely *proximate* whereas that of the members of the jury is *ultimate*.

In chains of institutional action (and responsibility) all the participants aim (or should be aiming) at the further (ultimate) end in addition to undertaking their own roles (and, therefore, aiming at the end definitive of their own particular role). Naturally, as mentioned above, the collective ends of institutions can be aimed at under somewhat different descriptions and, in some instances, may consist in overlapping individual ends. However, if in the case of institutional actors engaged in the core activity of an institution, and of an epistemic institution in particular, this overlap is not significant then coordination, mistrust and other problems arise and, eventually, institutional dysfunction. Moreover, given the moral significance of this activity, all the participants (at least, in principle) share in the *collective moral responsibility* (understood as joint responsibility (Miller, 2006)) for achieving that further end (or for failing to do so). In our cybercrime example, presumably the end in question is for the factually guilty to be found legally guilty (and the factually innocent not to be found legally guilty[ix]) and this is an end (a collective end) that is realised by the investigators working jointly with the members of the jury (and the other relevant institutional actors). It is not an end that the investigators could achieve on their own; they can only arrive at knowledge of factual guilt[x]. But equally it is not an end that the members of the jury could realise on their own; for they rely on the knowledge provided by the investigators.[xi] Further, this interdependence in the service of the ultimate end, namely, the collective end that the factually guilty be found legally guilty and the factually innocent be found legally not guilty, is something that both investigators and jurors are typically aware of, and certainly ought to be aware of. If investigators thought that their investigative efforts would never or hardly ever terminate in the factually guilty being found legally guilty and the factually innocent being found legally not guilty it would seriously undermine their commitment to conduct good quality investigations; what, after all, would be the point other than that they continued to be paid? Likewise, if jurors typically did not believe that their determinations of guilty and not guilty were unlikely to be correct because, for instance, the evidence on which these determinations were to be made probably should not be relied upon then this would seriously undermine their commitment to

their role as jurors; what, after all, would be the point other than that they discharged their obligation to undertake their role as jurors in the minimal sense of going through the motions (so to speak).

As we saw in relation to multi-layered structures of joint action and joint institutional mechanisms, it is important for the purposes in this article that institutional actors engaged in institutional processes can be understood in purely individualist terms and by recourse to the core notion of joint action, including joint epistemic action; hence the significance of the technical (individualist) notion of a chain of institutional action. Let us now turn to the analysis of the core notion in the construction of epistemic institutions: joint epistemic action.

3 | JOINT EPISTEMIC ACTION AND EPISTEMIC INSTITUTIONS

Our starting point is the invocation of a familiar threefold distinction made in respect of knowledge[xii]. First, there is *knowledge-by-acquaintance*: knowing someone or something. Knowledge-by-acquaintance is a genuine relation; the knower is conscious of some existent. Thus, knowledge-by-acquaintance does not consist in a truly believed proposition (justified or otherwise), but rather in the direct experience of an object.

Second, there is *propositional knowledge*; knowledge of the truth of some proposition (Ichikawa & Steup, 2017). This is knowledge that, for example, some state of affairs obtains. Propositional knowledge, it is here assumed, is expressible in a public language by utterances of sentences with a subject and a predicate. Whereas propositional knowledge is expressible in a public language, it is not necessarily *asserted* and, thereby, communicated to others; it might remain in the realm of inner thought. However, assertion is fundamental to individual and collective knowledge acquisition and dissemination.

Third, there is *knowledge-how*; knowing-how to do something e.g., knowing-how to ride a bike, knowing-how to read an x-ray film. Whereas knowledge-by-acquaintance and propositional-knowledge are cognitive states, knowing-how is essentially practical in character as I have argued elsewhere (Miller, 2022)[xiii]—although this is controversial—and, as such, more closely aligned with conative rather than cognitive states.

Evidently, we need to have all three sorts of knowledge. We need to verify certain claims by direct observation (knowledge-by-acquaintance). We also must have, and be able to obtain and communicate, propositional knowledge. In addition, we need to know how to do various things, e.g., drive a car, read a map. Moreover, these three different categories of knowledge are *interdependent* (Miller, 2022). This interdependence is both conceptual and means-end. However, for our purposes here we only need to make the weaker claim of means-end interdependence. Know-how, (e.g., knowing-how to drive a car) typically depends on knowledge-by-acquaintance (e.g., seeing and grasping the steering wheel). And the methods of acquiring new propositional-knowledge often depend on knowledge-by-acquaintance (e.g., observation), and know-how, (e.g., how to use scientific equipment such as microscopes), as do the latter two types on propositional- knowledge (e.g., a written manual describing scientific equipment and how to use it). Notice that with respect to so-called knowledge-how there is a distinction to be made between knowing how to perform an epistemic action, e.g., read an x-ray film, and knowledge how to perform a non-epistemic behavioural action, e.g., knowledge how to ride a bike (Miller, 2022). Here an epistemic action is one directed to, and constituted by, an epistemic end, whereas this is not so in the case of non-epistemic behavioural actions.

What of collective knowledge (Miller, 2017, 2018b, Ch. 3; Schmitt, 1994)? The salient notions of collective knowledge in the philosophical literature tend to be species of propositional knowledge. These are often referred to as common knowledge, mutual knowledge, mutual true belief and the like.[xiv] These notions are typically constructed out of the notion of mutual true belief. Thus, two agents, A and B, mutually believe truly that p if A believes truly that p, B believes truly p, A believes truly that B believes truly p, B believes truly that A believes truly that p, and so on.

Mutual knowledge—in the sense of mutual true belief—is closely related to another concept, namely, that which we will refer to as openness (Miller, 2015, 2016a). Openness is the social or interpersonal analogue of knowledge-by-acquaintance and, as such, is not necessarily propositional in character. For openness is mutual sensory awareness (hereafter mutual awareness) of an object and of oneself and the other person(s) as having awareness of that object. In the case of linguistic 'objects', speakers and hearers have mutual sensory awareness of utterances of sentences, i.e., of certain sorts of structured sounds and marks.

If openness is the social or interpersonal analogue of individual acquaintance-knowledge, *joint* knowing-how is the social or interpersonal analogue of individual knowing-how. Joint knowing-how finds expression in joint action; joint actions are the exercise of joint knowledge-how. However, as was the case with individual knowing-how, we need to distinguish an epistemic from a non-epistemic joint knowing-how.

Consider, for instance, the building of a skyscraper. This involves architects, engineers, bricklayers, carpenters, electricians etc., all of whom have specific forms of individual know-how (both epistemic and non-epistemic), but none of whom are individually possessed of all the different forms of knowing-how. Accordingly, their *collective* i.e., *joint*, know-how is required in order to realise the collective end of constructing the skyscraper.

We have distinguished three forms of collective knowledge, namely, propositional (mutual knowledge), acquaintance (mutual awareness), and practical (joint knowledge-how). However, there are two additional species (or, perhaps, sub-species) of collective knowledge that should be mentioned since they are highly relevant to epistemic institutions. The first of these we will refer to as public knowledge, the second as expert knowledge. These two species of collective knowledge have a propositional and a practical form.[xv]

In its propositional form public knowledge consists of true propositions that are matters of individual knowledge in the ordinary sense for some persons[xvi], i.e., it is 'in their heads', but for many or most these propositions are only knowledge in the sense that they are available for acquisition. Thus, much of the information stored in hardcopy format in books in libraries, in softcopy format in electronic databases, in public records, (e.g., court records) is public-*propositional*-knowledge. Again, the propositional knowledge in the heads of relevant public officials, such as those serving in information counters at railway stations, is public knowledge in our sense.

In its practical form, public knowledge consists of individual know-how (e.g., how to bake a cake, how to drive a car, how to read and write) that is either actually possessed, or is available for acquisition, by all or most members of some 'public'. Thus, the widespread availability of 'how to' manuals, driving lessons, primary school education and, in the end of the day, the widespread access to human persons possessed of the relevant 'know-how' and capable of inducting others into it, ensures that there is public-*practical*-knowledge.

Expert-propositional-knowledge is knowledge that is typically[xvii] 'in the heads' of the members of some group (the experts) in the form of mutual knowledge, but this knowledge is not 'in the heads' of another group (the non-experts). Expert-knowledge, like public-knowledge, is frequently stored in libraries, databases and so on that are, at least in theory, accessible to the public, i.e., the non-experts. However, expert-knowledge is *not readily understandable* by

ordinary members of the public, and so it is not in a substantive sense available to them. Thus, much scientific knowledge in academic journals is expert-propositional-knowledge, but not public-propositional-knowledge.

Expert-*practical*-knowledge is actually possessed by experts or is readily available to them, e.g., by way of professional top-up training courses. Expert-*practical*-knowledge is akin to expert-propositional-knowledge in that it is not in a substantive sense available to the public. For example, the surgeon's knowledge-how to perform open-heart surgery is limited to those who gain access to medical schools, pass examinations, and so on.

A further category of knowledge is what might be referred to as *secret knowledge*. Secret knowledge may well be understandable by experts or even by members of the public, but it not readily available to them. For instance, it might be knowledge classified as secret national security intelligence or knowledge possessed by members of an organisation concerning its criminal activities. Intelligence officers and investigative journalists are often focussed on unearthing secret knowledge.

Let us now address the question of the action-based character of knowledge. Naturally, some epistemic states, such as knowledge that there was a loud bang a moment ago, do not result from actions (in the sense of intentional actions). However, many do, such as the knowledge acquired by the police as a result of their investigation that Sutcliffe is the Yorkshire Ripper.

A fundamental kind of epistemic action is that of judgement. Beliefs and, therefore, knowledge are often the terminal point of an act of judgement, and evidence-based acts of judgement are typically freely performed (Frankish, 2007; Miller, 2015; Montmarquet, 1993; Walker, 1996). For example, an examinee comes to believe on the basis of a series of calculations that the answer to a complex mathematical problem is zero; the examinee is not absolutely certain, after all she could have made a mistake, but after checking she is fully confident of her own judgement. As it turns out the examinee gave the right answer based on valid mathematical reasoning. Surely, the inference-based judgement that terminated in her belief that the answer was zero was freely performed. By this I do not simply mean that she freely chose to try to answer the mathematical problem, although this is also true; rather I mean that in being 'compelled' by logic her act of judgement was, nevertheless, freely performed. In this respect judgements are akin to actions in general; an action that is 'compelled' by reason does not thereby cease to be a freely chosen action. In short, judgements are epistemic actions (Miller, 2015). What of *joint* epistemic actions?

As noted above, epistemic action is action directed to an epistemic end and, if the epistemic end is a collective end, then the action in question may well be joint epistemic action. Examples of such joint epistemic action would be a team of detectives seeking to determine the identity of the Yorkshire Ripper (Miller, 2014) or a team of scientists seeking a cure for cancer.

Naturally, epistemic action can involve behavioural action; consider the evidence gathering activities of the detectives mentioned above. However, this is not necessarily the case. For example, mental acts of judgement are epistemic actions because directed at truth, knowledge, understanding or some other epistemic end; but they are not necessarily instances of behavioural action.

Accordingly, let us refer to actions with a defining epistemic end (whether they involve bodily action or not) as epistemic actions; what makes these actions epistemic is that their description necessarily involves a reference to an epistemic end. By contrast, non-epistemic actions do not have a defining epistemic end; the description of a non-epistemic action does not necessarily involve a reference to an epistemic end. Some behavioural actions are also epistemic. For instance, if someone shouts out loud or writes the words, 'The king is dead' he or she might be performing

an assertion but also making a bodily movement, e.g., a vocalisation or a hand-writing movement. Which brings us to that most fundamental of joint epistemic actions: assertion.

A fundamental kind of joint epistemic action is that of assertion (Greco, 2020; Miller, 2016a). On Miller's account (Miller, 2016a), an assertion *qua joint action* is an instance of a social practice which has as its point or collective end (normatively speaking) the communication of (relevant) truths. Thus, a particular assertion by a speaker qua instance of the social practice of assertion is unsuccessful it is not true (or is irrelevant), and/or it is a lie, and/or it is not believed by the hearer. According to Miller's analysis (Miller, 2016a), a successful and felicitous[xviii] act of assertion by speaker S to hearer H that p (by means of utterance U) will consist of the following.

1. S and H perform a joint action of communication of p (with intended assertoric force);
2. S aims at truth (by making a judgment) and H trusts S (by making a judgment);
3. S has as an end (in saying that p with intended assertoric illocutionary force) that H judge that p: i.e., S asserts that p (given conditions (1) and (2));
4. p is true;
5. H judges that p (on basis of an inference from S's assertion that p);
6. S and H mutually truly believe that p[xix];
7. Collective end, E, of practice of assertion is realized, namely, that S and H mutually truly believe that p*.

Here it is important to understand the meaning of the asterisk attached to the proposition that p. When the collective end is realized then S and H mutually truly believe that p. However, prior to its realization the content of the collective end is not the proposition that p; since H, in particular, does not know, believe or judge that p. Rather the content of the collective end is something to the effect that—in the case where it is simply a matter of determining whether or not p is true or false—either p or not p (whichever is true, assuming both are relevant). In the case where the question at issue is open-ended, e.g. Who is the Yorkshire Ripper? then it is not simply a matter of determining whether p is true or false; rather the answer is some proposition the content of which is unknown by H (as opposed to the truth or falsity of known content being unknown by H). For simplicity, I shall use the asterisk to refer to both kinds of case.

Note that this phenomenon of the content of the collective end necessarily being unknown by at least some of the participants is a feature of joint epistemic action but not of joint non-epistemic behavioral action. In joint non-epistemic behavioral action, the aim is to bring into existence the state of affairs specified in the content of the collective end: the point is to bring the world (so to speak) in alignment with the prior and, therefore, known content of the collective end. By contrast, in the case of joint epistemic action the aim is in fact to provide the required specification of the content of the collective end: the point is to provide the (as yet unknown) content of the collective end by investigating the world and doing so in a manner that enables us to align this collective end content with the way the world is in the respect in question. In conclusion, then, successful and felicitous assertion is a joint epistemic action (having a constitutive joint action of communication as part of the means to its collective end).

Assertions (and related truth-seeking speech acts) and, therefore, joint epistemic actions are the basic building blocks of structured discourse and written material the principal purpose of which is epistemic; indeed, the purpose of which is a collective epistemic end such as mutual true belief or, more likely a structure of mutual true beliefs, with respect to some matter. Naturally, the practice of assertion does not exist in isolation from other social phenomena. The practice of assertion is itself governed by various social norms, notably social norms of truth and trust

(Grice, 1975), but also others such as the social norms governing the evidential bases for assertions (Goldberg, 2020). Moreover, obviously assertions interact with other speech act types, such as questions and commands. Further, these linguistic practices are themselves embedded in non-linguistic strategic and institutional (in our organizational sense) contexts.

In relation to epistemic institutions, we noted above the importance of collective knowledge in its various forms, including mutual true belief and joint know-how. We also need to stress the importance of (above-described) multi-layered structures of joint *epistemic* action, joint *epistemic* institutional mechanisms and chains of institutional *epistemic* action since, as is the case with their non-epistemic counterparts, they are in part constitutive of social institutions in general, but especially of epistemic institutions. Consider, for instance, the Human Genome Project : an essentially epistemic project.

The Human Genome Project (HGP) was a large scientific project conducted by a number of cooperating organisations (principally 20 universities and research centres) and hundreds of scientists over many years (roughly from 1988–2001) HGP was the international, collaborative research program whose collective epistemic end was the complete mapping and understanding of all the genes of human beings, i.e., the human genome. According to the National Human Genome Research Institute, 2004, "The HGP has revealed that there are probably about 20,500 human genes. The completed human sequence can now identify their locations. This ultimate product of the HGP has given the world a resource of detailed information about the structure, organization and function of the complete set of human genes. This information can be thought of as the basic set of inheritable 'instructions' for the development and function of a human being."[xx] Accordingly, the realised collective end of the project was collective expert knowledge of the human genome, i.e., a complex, integrated structure or web of knowledge. This web of knowledge consists of fragments of knowledge and these fragments were the epistemic contributions of multiple researchers working in multiple different organisations world-wide and doing so over in incremental stages over a lengthy period of time. So HGP involved realizing multiple, nested, collective epistemic ends (fragments of knowledge) in the service of the larger collective epistemic end of mapping and understanding the human genome (web of knowledge), and multiple layered structures of joint epistemic action undertaken in coordinated stages in various institutional settings in order to realize this larger collective epistemic end.

HGP and, for that matter, most modern scientific enterprises, are a species of organisational action involving multi-layered structures of joint epistemic action and chains of institutional action. Moreover, the organisations in question are, for the most part, hierarchical institutions comprised of task-defined roles standing in authority relations to one another, designed in accordance with principles of division of labour and governed by a complex network of conventions, social norms, regulations and laws. Further, scientific enterprises make extensive use of joint institutional mechanisms. Consider, for instance, the independent replication of experiments. These are instances of joint epistemic action in which the participants must exercise joint abilities by virtue of institutional requirements that could have been otherwise. The reliability of experimental results is assured by the independent replication of experiments. While the ability to replicate an experiment is an individual ability, the *independent* replication of experiments presupposes multiple scientific experimenters. Since one scientist acting on his own cannot meet the institutional requirement of independence, this procedure manifests a joint epistemic ability. Moreover, the result might not be confirmation of the original experiment. However, participants agree to abide by this result in a manner akin to abiding by the results of the joint institutional mechanism of voting discussed above.

An institutional mechanism made much of in relation to what we are referring to as joint institutional mechanisms (and, indeed, in the case of Searle in relation to the ontology of institutions more generally (Searle, 1995)) is that of constitutive rules at the core of which, at least on Searle's account, are linguistic performatives. Roughly speaking, the idea is that these rules have the form 'X counts as Y in some context' and by virtue of collective acceptance of these rules they *create* institutional entities, rights and duties. There are, of course, some limited instances, such as J. L. Austin's famous example of the naming of a ship (Austin, 1962), in which performative utterances by virtue of conventions bring about new institutional states of affairs (Miller, 1984).

However, the extension of this performative model to organisations that are institutions, and to their constitutive institutional roles in particular, goes way too far; *in general*, fundamental (as opposed to ornamental) institutional roles cannot simply be created by convention-governed performative utterances or otherwise by conventional fiat. For instance, someone who has no medical knowledge or practical skills in surgery is not a surgeon, even if declared to be such by the Royal College of Surgeons (or other relevant authority) and, therefore, accepted to be such by the wider community; conversely someone with a high level of the relevant medical knowledge and practical skills who successfully engages in surgical operations is a surgeon, even if he is not recognised as such by the Royal College of Surgeons (or other relevant authority) or, for that matter, by the community in general. In short, the ability to perform actions, including epistemic actions, in institutional settings lies at the core of institutional actors and action rather than the existence of convention governed performatives (or like conventional devices) (Miller, 2001, Ch.5).

4 | TYPES OF EPISTEMIC INSTITUTIONS

Given that epistemic activity is conducted in *all* institutions and that epistemic action is implicated in non-epistemic behavioral action, there is a need for an elaboration of the distinction between epistemic and non-epistemic institutions.

We suggested above that the notions of collective knowledge (propositional and practical) understood as the principal and ultimate collective end(s) of an institution—might suffice to distinguish essentially epistemic institutions from essentially non-epistemic ones. Thus, Toyota Corporation is not an essentially epistemic institution, notwithstanding that it undertakes research into, say, electric cars, because its principal and ultimate collective end is the production of cars, not knowledge about cars. On the other hand, a university which conducts research into electric cars remains an essentially epistemic institution if it stops short of producing cars (other than perhaps to demonstrate how this can be done, i.e., its research yields knowledge-how).

It might be argued against this that the *ultimate* collective end of the research undertaken in universities[xxi] while it is paradigmatically expert-knowledge, as opposed to public-knowledge, is, nevertheless, some benefit to the wider community such as, for instance, an improvement in the community's transport arrangements (perhaps in part by rendering some expert-knowledge in the form of public-knowledge). However, it would not follow from this that implementing these transport arrangements, as opposed to providing knowledge about how to do so, was a *principal* collective end of the university. Moreover, the ultimate collective end in such a case would be one to which the university was only one type of contributor; an epistemic contributor, so to speak.

What of the different sub-species of epistemic institutions? Let us consider the differences between universities, news media organizations and national security intelligence agencies. We

have suggested that universities have collective knowledge, notably *expert* knowledge, as their principal and ultimate collective ends. What of news media organizations? Presumably, these organizations have, or ought to have, *public* knowledge in relation to, for instance, important political issues, and certain sorts of *secret knowledge*,[xxii] as their principal and ultimate collective end[xxiii]. However, unlike universities, perhaps this knowledge is not typically new knowledge[xxiv]; rather it is knowledge possessed by others that is in need of public *dissemination*, i.e., in need of being rendered into public-knowledge in the public interest. In short, unlike universities, the principal and ultimate collective end of news media organizations is the *public dissemination* of existing knowledge (including some secret knowledge). Naturally, this typically requires that this knowledge be acquired by journalists from knowledgeable sources following on a process of not simply collection but verification or other form of analysis; and perhaps only acquired with great difficulty, as in the case of investigative journalists.

What of national security intelligence agencies? National security intelligence is an epistemic notion and, ideally it consists of knowledge, i.e., it is true or correct or accurate or probably true, or some such. Thus, intelligence officers aim at, or ought to aim at, knowledge as is the case with both academics and journalists. However, unlike academics, but like journalists, generally speaking, this knowledge is not new knowledge but is *secret knowledge* acquired from knowledgeable sources. Moreover, in the case of intelligence officers, unlike academics and (to a lesser extent) journalists, the sources of this knowledge are typically highly resistant to it being acquired, as in the case of secret intelligence acquired about foreign military organizations. Indeed, intelligence officers confront foreign *counter-intelligence* operations. Further, unlike knowledge acquired by journalists for public dissemination, national security intelligence is typically secret intelligence acquired (again, following on a process of collection and analysis) in the service of the collective end of national security and, as such, is not for public consumption. Indeed, even within a national security intelligence organization intelligence (or knowledge resulting from the collection and analysis of intelligence) might only be disseminated on a 'need to know' basis, although intelligence officers are often said to be prone not to provide such intelligence or knowledge to their fellow officers even when they need to possess it. Finally, national security intelligence needs to be actionable by political leaders and security agencies. For instance, national security intelligence in relation to a planned foreign military attack might require, at the very least, putting in place defensive measures and, indeed, might require a pre-emptive attack. In this respect national security intelligence is somewhat different from both the knowledge sought by academics and that sought by journalists.

Knowledge sought by intelligence officers, journalists and academics have something in common: the pursuit of knowledge as an end-in-itself. This might seem counter-intuitive in the case of journalists and intelligence officers. Naturally, any such knowledge, if it is to be collective knowledge, will be disseminated to academic peers, members of the public and national security decision-makers, respectively. Moreover, as we have just seen, national security intelligence *also* needs to be actionable, and knowledge acquired by journalists *also* should enable the citizenry to make informed political decisions. On the other hand, arguably, much expert-knowledge acquired by academics merely needs to be knowledge acquired as an end-in-itself and disseminated to academic peers (e.g., in academic publications).

Against this, it might be argued that the knowledge acquired by journalists and intelligence officers is merely a means to a collective end; the acquisition of this knowledge is not an end-in-itself. Here there are two points to be made. Firstly, it is necessary that journalists and intelligence officers have an overriding professional commitment to the epistemic purposes (collection, analysis and dissemination of knowledge) of their profession and intelligence agency (respectively)

rather than seeking to realize the public interest *outcomes* that might or might not flow from the journalists' public dissemination of this knowledge or, in the case of the intelligence officers, the national security outcomes that might or might not result from the decisions of the politicians, military leaders and other decision-makers who receive their intelligence. It is important that journalists and intelligence officers qua journalists and intelligence officers (respectively) do not engage in institutional overreach. Incidentally, an analogous point might be made in respect of academics. It is important that academics qua academics do not engage in institutional overreach by seeking to realize outcomes that, if they are to come about, ought to come about as the result of the decisions of non-academics who are in receipt of academics' research. In all three cases, such institutional overreach may well undermine the perception of epistemic independence necessary for the retention of their epistemic authority.

Secondly, the acquisition of knowledge is an end-in-itself for journalists and intelligence officers, as it is for academics, notwithstanding the further requirement that the truths acquired by journalists and intelligence officers, in particular, are also the means to further ends. For truth and, therefore, knowledge is not an external, contingently connected end of the epistemic activities of journalists and intelligence officers which those activities might be directed towards if the journalists or intelligence officers happen to have an interest in truth, rather than, say, an interest in falsity. Rather truth is internally connected to epistemic activities, including those of journalists and intelligence officers, as well as those of academics and others. For to aim at truth by, for instance, making a judgement is to aim at truth as an end-in-itself. Relatedly, to make an assertion (Dummett, 1973, Ch. 10) is to represent one's-self as aiming at the truth as an end-in-itself rather than merely as a means to an end; a means that one might abandon if, say, falsity would be a more effective means to the end in question. Aiming at truth as an end-in-itself and representing one's-self as aiming at truth as an end-in-itself is consistent with *also* aiming at truth as a means to some other further end, such as winning a war. In other words, journalistic or intelligence activity which *only* aimed at truth as a means to some other end, and presented itself as such, would not be genuine journalism or intelligence activity since it would not be genuine epistemic activity.

Of course, all institutions consist of institutional role occupants who make judgements and perform assertions, i.e., consist of institutional role occupants who aim at truth as an end-in-itself. However, it does not follow from this that all institutions are epistemic institutions. For unlike an epistemic institution, in the case of a non-epistemic institution the kinetic actions that rely on the knowledge provided by institutional role occupants who aim at truth as an end-in-itself are performed by institutional role occupants who are members of this very same non-epistemic institution itself. For instance, Toyota employees produce cars as well as conduct research. Moreover, in the case of non-epistemic institutions, although some knowledge is a necessary means to the kinetic action that is their core business (so to speak), the acquisition of new expert knowledge (as opposed to public knowledge acquisition or expert knowledge acquisition by means of the testimony of others) or of secret knowledge is not core business. Accordingly, the acquisition of new expert knowledge or of secret knowledge could, at least in principle, be outsourced by a car manufacturer to another institution without the car manufacturer ceasing to be a car manufacturer. But a university would cease to be a university if it outsourced the acquisition of new expert knowledge to another institution. Likewise, a national security intelligence agency would cease to be a national security intelligence agency if it outsourced the acquisition of secret intelligence to another institution.[xxv]

Notes

[i] Of course, institutions have other general properties such as institutional cultures.

[ii] Unless, of course, the notion of collective acceptance is simply understood as action in accordance with the constitutive rules of an institution, in which case it is trivially true that collective acceptance is a constitutive feature of institutions. See Miller (2001, Ch. 5).

[iii] This is true even in cases such as that of the firing squad example in which only one member of the squad fires a live bullet killing the victim.

[iv] This individualist account is eliminative, i.e., there are on this account no such things as institutional agents (understood as agents possessed of group intentions, beliefs etc.) per se. Moreover, a possible analogy with reductionist accounts of human agency and its constitutive mental states etc. is rejected. In the case of human agency, there are, for instance, mental states and actions, e.g., judgements, of which the agent is conscious and, therefore, directly aware. Accordingly, there is a need for a reductive analysis, if a materialist account is to be plausible. However, in the case of the alleged institutional agents (understood as possessed of mental states) there is no need for a reduction in this sense since there is no prior direct awareness of group intentions, beliefs, judgments etc.; rather there is simply a manner of speaking, e.g., "BP lied when it asserted that BP believes in protecting the environment", that might be thought to entail group intentions, beliefs, judgements etc. (see, for instance, Lackey "Group Assertion" *Erkenntnis*). However, since the lie (or other assertion) is written in a document (or otherwise produced) by a human being (taken to be representing BP corporation), such as its CEO or a spokesperson, there is no direct awareness even of the group assertion (let alone of the intentions, beliefs entailed by a group assertion).

[v] Note that in my view collective entities, such as corporations, are not moral agents and, therefore, cannot be held morally liable but can, nevertheless, be held legally liable. See Miller (2010, Ch. 10).

[vi] This would include the 'weaker' kind favoured by Tuomela (2013).

[vii] This is not strictly correct; rather, typically, some threshhold set of actions is necessary to achieve the end.

[viii] There are, of course, any number of alternative voting systems in democracies.

[ix] Assuming there are only two possible verdicts, guilty and innocent, which is not the case in some jurisdictions, e.g. Scotland.

[x] The concept of knowledge is philosophically contested and the subject of a vast literature. See, for instance, Moser (1989). Let us assume here that it is justified true belief.

[xi] Chains of institutional action, and of associated institutional and moral responsibility consist of a process in which the completion of one stage institutionally triggers the commencement of the next stage, e.g., arrest is followed either by the suspect being charged or released within a specified time-frame.

[xii] See Miller (2022) for a defence of a non-reductive account of these three species of knowledge.

[xiii] Many, if not all, of the claims made below could be made by those who view knowing-how as a form of propositional knowledge.

[xiv] For convenience, we use the term "mutual" rather than "common" when referring to the kind of phenomena in question. For definitions of some of these notions see, for example, Smith (1982).

[xv] They may well also have a knowledge-by-acquaintance form.

[xvi] Although this is generally the case, it is possible that some instances of public knowledge are not in fact in anyone's head but only available for acquisition.

[xvii] It is possible that some instances of expert knowledge are not in fact in anyone's head but only available for acquisition by experts.

[xviii] A speech act can be an act of assertion even if it is false or a lie. However, the use of the terms, "successful" and "felicitous" is intended to rule out such cases. Note that from the Hearer's perspective, an assertion might be unsuccessful if it is a lie. Moreover, as mentioned above, an assertion *qua joint action* is unsuccessful it is not true, and/or it is a lie, and/or it is not believed by the Hearer.

[xix] Hence the paradoxical nature of G. E. Moore's assertoric form 'I believe that p but not p'. See, for example, Pruss (2011).

[xx] See Human Genome Project (2004) at the homepage of the National Human Genome Research Institute at www.genome.gov.

[xxi] For reasons of space the question of imparting knowledge, including epistemic knowledge-how to students, is not addressed here.

[xxii] Presumably, secret knowledge that it is in the public interest to disclose.

[xxiii] It might be argued that news media organisations that are corporations ought to have as their principal and ultimate collective end, profit. See Miller (2010, Ch. 8) for counter-arguments.

[xxiv] Roughly speaking, new knowledge is knowledge that no-one currently has (other than the current acquirer of this knowledge). So on this (stipulative) definition knowledge acquired of the activities of a currently existing but remote community would not be new knowledge; however, currently acquired knowledge of the past activities of a community that no longer exists would be new knowledge (if this knowledge is not currently possessed by anyone else).

[xxv] Some news/media agencies outsource their investigative journalist activity (or have never conducted such activity) and function simply as a conduit for the public communication of knowledge acquired by others, as well as being the providers of a public platform for comment (often including by their own journalists). These agencies are not so much epistemic institutions as disseminating institutions (akin to social media providers) albeit perhaps with a curatorial, e.g., editorial role, and a role as commentators.

References

Austin, J. L. (1962). *How to do things with words*. Oxford University Press. https://doi.org/10.1111/j.1468-0149

Bratman, M. (2014). *Shared Agency*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199897933.001.0001

Copp, D. (2006). On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'. *Midwest Studies in Philosophy*, *30*(1), 194–221.

Dummett, M. (1973). *Frege: Philosophy of Language*. Duckworth. https://doi.org/10.1017/S0031819100063385

Frankish, K. (2007). Deciding to Believe Again. *Mind*, *116*(463), 523–547.

French, P. (1984). *Collective and Corporate Social Responsibility*. Columbia University Press. https://doi.org/10.7312/fren90672

Gilbert, M. (1989). *On Social Facts*. London: Routledge. https://doi.org/10.1080/02691729108578620

Goldberg, S. (2020). *Conversational Pressure*. Oxford University Press. https://doi.org/10.1093/pq/pqaa075

National Human Genome Research Institute (2004). Human Genome Project. www.genome.gov

Goldman, A. and Whitcomb, D. (eds.) (2011). *Social Epistemology: Essential Readings*. Oxford University Press.

Greco, J. (2020). *The Transmission of Knowledge*. Cambridge University Press. https://doi.org/10.1017/9781108560818

Grice, H. P. (1975). Logic and Conversation. *Syntax and Semantics*, *3*, 41–58.

Ichikawa, J. J., & Steup, M. (2017). The analysis of Knowledge. *Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/knowledge-analysis/

Lackey, J. (2018). Group Assertion. *Erkenntnis*, *83*(1), 21–42.

List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford University Press. https://doi.org/10.1111/j.1468-5930.2012.00554.x

Ludwig, K. (2016). *From Individual to Plural Agency*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198755623.002.0003

Ludwig, K. (2017). *From Plural to Institutional Agency*. Oxford University Press. https://doi.org/10.1093/oso/9780198789994.002.0005

Miller, S. (1984). Performatives. *Philosophical Studies*, *45*(2), 247–60.

Miller, S. (1992). Joint Action. *Philosophical Papers*, *21*(3), 275–297.

Miller, S. (2001). *Social Action: A Teleological Account*. Cambridge University Press. https://doi.org/10.1017/CBO9780511612954

Miller, S. (2006). Collective Moral Responsibility: An Individualist Account. In (Ed.) P.A. French *Midwest Studies in Philosophy*, 30 (pp. 176–193). Blackwell.

Miller, S. (2007a). Joint Action: The Individual Strikes Back. In (Ed.) S. L. Tsohatzidis *Intentional Acts and Institutional Facts: Essays on John Searle's Social Ontology*. (pp. 73–92). Springer Press.

Miller, S. (2007b). Against the Moral Autonomy Thesis. *Journal of Social Philosophy*, *38*(3), 389–409.

Miller, S. (2010). *The Moral Foundations of Social Institutions*. Cambridge University Press. https://doi.org/10.1017/CBO9780511818622

Miller, S. (2014). Police Detectives, Criminal Investigations and Collective Moral Responsibility. *Criminal Justice Ethics*, *33*(1), 21–39.

Miller, S. (2015). Joint Epistemic Action and Collective Moral Responsibility. *Social Epistemology*, *29*(3), 280–302.
Miller, S. (2016a). Assertions, Joint Epistemic Actions and Social Practices. *Synthese*, *193*(1), 71–94.
Miller, S. (2016b). *Shooting to Kill: The Ethics of Police and Military Use of Lethal Force*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190626136.003.0012
Miller, S. (2017). Ignorance, Technology and Collective Responsibility. In (Ed.) Rik Peels *Perspectives on Ignorance from Moral and Social Philosophy* (pp. 217–236). Routledge.
Miller, S. (2018a). Joint Epistemic Action: Some Applications. *Journal of Applied Philosophy*, *35*(2), 300–318.
Miller, S. (2018b). *Dual Use Science and Technology, Ethics and Weapons of Mass Destruction*. Springer. https://doi.org/10.1007/978-3-319
Miller, S. (2022). Epistemic Action, Abilities and Knowing-How: A Non-Reductive Account. *Social Epistemology*, *36*(4), 466–485.
Miller, S., & Makela, P. (2005). The Collectivist Approach to Collective Moral Responsibility. *Metaphilosophy*, *36*(5), 634–651.
Montmarquet, J. A. (1993). *Epistemic Virtue and Doxastic Responsibility*. Rowman and Littlefield. https://doi.org/10.1093/0195128923.003.0007
Moser, P. (1989). *Knowledge and Evidence*. Cambridge University Press. https://doi.org/10.1177/0014524613494634
Pruss, A. R. (2011). Sincerely asserting what you do not believe. *Australasian Journal of Philosophy*, *90*(3), 541–546.
Schmitt, F. (Ed.) (1994). *Socialising Epistemology*. Rowman and Littlefield. https://doi.org/10.1080/02691720701673934
Searle, J. (1995). *Construction of Social Reality*. Penguin. https://doi.org/10.1086/230850
Smith, N. V. (ed.) (1982). *Mutual Knowledge*. Academic Press. https://doi.org/10.1016/S0378-2166(99)00128-9
Tuomela, R. (2002). *Philosophy of Social Practices: A Collective Acceptance View*. Cambridge University Press. https://doi.org/10.1017/CBO9780511487446
Tuomela, R. (2013). *Social Ontology: Collective Intentionality and Group Agents*. Oxford University Press. https://doi.org/10.1515/9783110618204
Walker, M. T. (1996). The Voluntariness of Judgment. *Inquiry*, *39*, 97–119.

How to cite this article: Miller, S. (2022). Epistemic institutions: A joint epistemic action-based account. *Philosophical Issues*, 1–19. https://doi.org/10.1111/phis.12227