

Social Knowledge and Supervenience Revisited

Forthcoming in *Erkenntnis*

Mark Povich

Abstract: Bird's (2010; 2014) account of social knowledge (**SK**) denies that scientific social knowledge supervenes solely on the mental states of individuals. Lackey (2014) objects that **SK** cannot accommodate 1) a knowledge-action principle and 2) the role of group defeaters. I argue that Lackey's knowledge-action principle is ambiguous. On one disambiguation, it is false; on the other, it is true but poses no threat to **SK**. Regarding group defeaters, I argue that there are at least two options available to the defender of **SK**, both taken from literature on individual defeaters and applied to group defeaters. Finally, I argue that Lackey's description of the case of Dr. N. – as a case in which the scientific community does not know but is merely in a position to know – is mistaken. It assumes that Dr. N.'s publication is not scientific knowledge. An analogy to the individual case shows that it is plausible that the scientific community is not merely in a position to know, although its members are. This leaves intact a conception of social knowledge on which it does not supervene on the mental states of individuals.

1. Introduction

Bird (2010; 2014) has defended a radically extended account of scientific social knowledge. It is radically extended in the sense that it denies that group knowledge supervenes on the mental states of individuals. Let us, following Lackey (2014), call this account **SK**. Lackey (2014) objects that **SK** cannot accommodate 1) a knowledge-action principle (when combined with another plausible principle linking individual and group action) and 2) the role of group defeaters. I argue that Lackey's knowledge-action principle is ambiguous. On one disambiguation, it is false; on the other, it is true but poses no threat to **SK**. Regarding group defeaters, I argue that there are at least two options available to the defender of **SK**, both taken from literature on individual defeaters and applied to group defeaters.

Finally, I argue that Lackey's description of the case of Dr. N – as a case in which the scientific community does not know but is merely in a position to know – is mistaken. Lackey's analogy to the individual case can be redescribed to show that it is plausible that the scientific community is not merely in a position to know, although its members are.

In Section 2, I present more completely **SK**; this Section also spells out the central example of non-supervenience under dispute – the case of Dr. N. In Sections 3 and 4, I present and respond to Lackey's action and defeater objections, respectively. In Section 5, I show how Lackey's distinction between knowing and merely being in a position to know, while useful, is misapplied to the case of Dr. N.

2. Scientific Social Knowledge

When does a social epistemic subject, such as the scientific community, know?¹ According to **SK** (Bird 2010; 2014), a social epistemic subject *S* knows that *p* if and only if *p* is true, the information that *p* is accessible to relevant members of *S* (e.g., other scientists, but not necessarily the lay public), and the following three conditions hold²:

- (i) [*S* has] characteristic outputs that are propositional in nature (*propositionality*);
- (ii) [*S* has] characteristic mechanisms whose function is to ensure or promote the chances that the outputs in (i) are true (*truth-filtering*);
- (iii) the outputs in (i) are the inputs for (a) social actions or for (b) social cognitive structures (including the very same structure [the structure that produces the output]) (*function of outputs*). (Bird 2010, 42–4)

The three conditions are supposed to be social analogues of the functions of and conditions required for individual knowledge. I will give them independent defense here. My focus instead will be on responding to Lackey's (2014) objections.

Bird argues that **SK** implies that scientific knowledge does not supervene on the mental states

¹ My focus will not be on the ontology of social epistemic subjects. See Bird (2010; 2014) for an argument that the division of cognitive labor in science enforces an organic interdependence that constitutes it as a social subject.

² Bird does not give necessary and sufficient conditions for social knowledge. This biconditional is Lackey's (2014) reconstruction of Bird, which I endorse here, both as a reconstruction of Bird and as an account of social knowledge.

of individuals. This is illustrated by Bird's cases of Drs. Q. and N. These cases are identical with respect to the mental states of individuals, but differ with respect to social knowledge. Dr. Q. is a reclusive scientist who dies with her results tucked away from the scientific community. Even though Dr. Q. knew that p , the scientific community did not, in virtue of her lack of publication. This is because publication is what makes a proposition accessible to the scientific community (Bird 2010, 32). Dr. N., on the other hand, is a scientist whose results, though published in a reputable journal, every individual forgets once she dies. Bird (2010, 32) claims that, in virtue of publication and, so, accessibility, the scientific community knows Dr. N.'s conclusion that d , even though no individual scientist knows that d .³

The cases of Drs. Q. and N. suggest that the set of propositions known by the scientific community is independent of what individuals know. Individual knowledge is insufficient, as shown by the case of Dr. Q., and unnecessary, as shown by the case of Dr. N., for social knowledge. I defend Bird's treatment of the case of Dr. N. from objections in the next two sections.

3. Lackey's Objections: Action

Lackey (2014) argues that the following two principles conflict with **SK**.⁴

(KAP) Knowledge/Action Principle: S knows that p only if S is epistemically rational to act as if p or, equivalently, S is epistemically rational to act as if p if S knows that p . (287)⁵

(GMAP) Group/Member Action Principle: For every group, G, and act, a , G performs a only if at least one member of G performs some act or other that causally contributes to a . (286)⁶

³ As Bird (2010) describes the case, eventually Dr. N.'s publication is rediscovered, but its rediscovery is unnecessary for social knowledge. All that is required is its rediscoverability (i.e., accessibility).

⁴ If Carter (2015) is right, then Lackey's objections extend beyond **SK** to some mainstream views on group knowledge. Thanks to an anonymous reviewer for bringing this to my attention.

⁵ One might respond to Lackey's objection by denying **KAP** (Brown 2008; Lackey 2010), though I will not do so here. Thanks to an anonymous reviewer for pointing this out. Lackey considers this response (2014, 290).

⁶ I note a potential problem, or at least an imprecision, with Lackey's formulation of **GMAP**. Since a group (and its actions) and its members (and their actions) are at different mereological levels, a better formulation of **GMAP** might be: For every group, G, and act, a , G performs a only if at least one member of G performs some act or other that (*partly*) constitutes a . This formulation seems to make better sense of common examples of group action, including Lackey's negligent homicide case (286). It is also telling that List and Pettit (2011, 65) begin their discussion of the supervenience of group agents on their members with a constitution example (i.e., the supervenience of the shapes made by dots on a grid on the positions or coordinates of those dots). I leave open the possibility that there are cases of group

According to Lackey, these two principles are inconsistent with **SK**, as shown by applying them to the above-mentioned case of Dr. N., whose results, though published in a reputable journal, everyone forgets once she dies. In this case, Bird and Lackey agree that **SK** implies that the scientific community knows that d even though no individual scientist does. Applying **GMAP**, if the scientific community were to act, it would be through members unaware of d . Lackey (2014, 288–9) argues that it is therefore epistemically irrational for the scientific community to act as if d , for example, by asserting⁷ d or approving drugs that depend on d ⁸. But, given **KAP**, the epistemic irrationality of acting as if d shows that the scientific community does not know d . Hence, **GMAP** and **KAP** show that the scientific community does not know that d ; therefore, **SK** is false.

In response, I argue that **KAP** is ambiguous between two different sense of “act as if”. I distinguish “acting on” knowledge that p from merely “acting in accordance with” knowledge that p . On the latter disambiguation, **KAP** is false. On the former disambiguation, **KAP** is true, but does not show that **SK** is false.

The disambiguation has to do with knowledge's role in action production. Consider Hawthorne and Stanley's reasoning when defending their version of **KAP**: “When someone acts on a belief that does not amount to knowledge, she violates the norm, and hence is subject to criticism” (Hawthorne and Stanley 2008, 577). The phrase “acts on” seems to imply that the belief plays some causal role in producing the subject's action⁹. This is the plausible disambiguation of “act as if”. The less plausible disambiguation treats acting as if p as merely acting in accordance with p , where it could be sheer luck

action in which Lackey's causal contribution formulation of **GMAP** is more appropriate.

⁷ Lackey uses assertion as an example, but note that knowledge-action principles are usually distinguished from knowledge norms of assertion. Thanks to Julia Staffel for bringing this to my attention.

⁸ Note that this implies that it is irrational to act as if d while being unaware of d . I am not sure that this is always true, especially on my “acting on” disambiguation of “acts as if,” but I will assume it here. It does not imply – what is likely false – that it is irrational to act as if d while being unaware of one's knowledge that d .

⁹ This role is more explicit in other formulations of knowledge-action principles that stress knowledge's role in practical reasoning (e.g., Hawthorne and Stanley 2008), where knowledge must interact with other mental states in rational and causal processes. Lackey's formulation of **KAP** downplays this.

that one's action so accords. In merely acting in accordance with p , one's knowledge that p plays no role in the action.¹⁰ Compare two cases: 1) agent A knows that p and merely acts in accordance with p , for example, asserting that p , though her knowledge plays no causal role in the production of her assertion, and 2) agent B knows that p , but in asserting that p , her knowledge that p does play a causal role in the production of her assertion. A's action is not epistemically rational, but agent B's action is. This suggests that the proper formulation of **KAP** is something like: If S knows that p , then S is epistemically rational to act *on* p (not merely in accordance with p). **SK** is consistent with this.

Now we can apply this distinction to the case of a social epistemic subject. The epistemically irrational scenario above (i.e., the case of agent A) is analogous to Lackey's description of the case of Dr. N. In Lackey's example, the reason the scientific community is irrational for asserting d or approving drugs that depend on d is that, in my terms, the scientific community is merely acting in accordance with d , rather than acting on d . In Lackey's example, the scientific community's knowledge that d seems to play no causal role in the production of its actions. I agree that this is epistemically irrational, but it is also irrational in the analogous case of individual knowledge, as shown above. In both the social and individual cases, the irrationality stems from the fact that acting in mere accordance with p would be a matter of epistemic luck (Lackey 289).

I claim that in the case under dispute the scientific community is merely acting in accordance with d , rather than acting on d , because the knowledge that d is not playing a causal role in producing social actions. What does it mean for knowledge that d to play a causal role in producing the scientific community's actions? Let us briefly look at the individual case first. For knowledge that d to play a role in producing an agent's action, her knowledge that d must play an appropriate (i.e., non-deviant; this

¹⁰ Cf. Turri (2011), who distinguishes between having a reason and believing for a reason.

qualification applies throughout)¹¹ causal role in producing her action¹².

Presumably, what this means is something like that her brain and motor system are wired up such that the brain state on which her knowledge supervenes is (non-deviantly) causally relevant to her motor output¹³. So, to apply this to the social case, we must have some account of the supervenience base of scientific social knowledge, so that we can see how it may be causally relevant to social action. In social cases, the supervenience base of knowledge can include the mental states of individuals, but it also includes social institutions and objects like articles, laboratories, and the internet (Bird 2010; 2014). For knowledge that *d* to play a causal role in producing the community's action, then, is for the community to make use of these knowledge-supervening social structures in its decision-making (see condition iii of **SK** above). Since one of the relevant structures – perhaps the only structure – on which the social knowledge that *d* supervenes in the case of Dr. N. is her publication, the only way for the community to act *on d* is for her publication to play some role in the community's decision-making. This is what is left out of Lackey's description of the case, which is why the scientific community is irrational; Dr. N.'s publication seems to play no role in the community's social actions. The way Lackey describes the case of Dr. N. is not as a case of the scientific community acting *on* its knowledge that *d*, but merely acting in accordance with it. Were the scientific community to act *on d*, even if no individual were aware that *d*, it would be rational.

I claim that it would be rational for the scientific community to act *on d*, even if no individual were aware that *d*, but how could it so act? That is, how could the social knowledge that *d* play a non-deviant causal role in the community's social action without any of its members being aware that *d*?

¹¹ I set aside worries about causal deviance, i.e., whether and how her knowledge must cause her action *in the right way*. Whatever the right account is in the individual case, it can likely be extended to the social case (e.g., Turri 2011). Note also that a non-deviance condition should presumably be added to Lackey's construal of **GMAP** as well (i.e., at least one member of G must perform some act or other that *non-deviantly* causally contributes to *a*).

¹² She does not have to be aware of her knowledge that *d*.

¹³ Of course here I am assuming something about the metaphysics of (individual) knowledge, namely that it supervenes (at least partially) on states of the brain. If this false, I find it hard to make sense of knowledge's causing action at all, but that is beyond the scope of this paper.

One of Bird's (2010) examples of scientific knowledge without awareness could be contrived to allow for this, especially if reliable automation is involved (35). For example, consider Lackey's description of the case of Dr. N. again, where no one is aware of Dr. N.'s finding that d , though the community makes d -relevant decisions. Now, let us add that the scientific community has created a reliable, automated system that checks past results for certain truth-conducive properties like statistical significance, statistical power, effect size, etc. If the automated system detects a result that passes the community's standards, it sends the result to another automated system, reliable and trusted by the community, that uses the result to help the community make social decisions. The first system detects Dr. N.'s result and sends it to the next system for processing in social action. This system might, on the basis of d , suggest that a certain drug be developed, which is then approved by the community (via some causal contribution of at least one member, such as some member's clicking an "approve" button). Here it is plausible that the knowledge that d plays a non-deviant causal role in the production of social action, such as approval of a drug, without any member being aware that d .¹⁴ As Bird (2010, 35) notes, cases like this are becoming increasingly likely as the scientific division of labor becomes ever more divided and automated. The Large Hadron Collider may provide a real example (Bird 2010, 45).

I conclude that **SK** is consistent with **GMAP** and **KAP** properly construed to include a causal role for knowledge in action production; it would be rational for the scientific community to act *on d*, even if no individual is aware that d .

4. Lackey's Objections: Defeaters

Lackey's (2014) second objection to Bird's treatment of the case of Dr. N. has to do with

¹⁴ Lackey (2014, 289) considers a case of social action without awareness, "Distributed Information," that is somewhat like mine, but it seems to run afoul of **GMAP**, since the group assertion itself is automatic; so, this cannot be an example that shows that if **GMAP** and **KAP** are true, then **SK** is false. Whether it does run afoul depends on exactly what "causally contributes" means in **GMAP**. Considering my constitution construal of **GMAP** (see fn. 6), there does not seem to be any action of any member of the group that partly constitutes the group's assertion.

defeaters. She distinguishes between psychological defeaters and normative defeaters¹⁵ (2014, 292). A psychological defeater is a mental state of S's that rebuts her belief that p or undercuts her justification for believing that p . A normative defeater is a psychological defeater that S should have, given available evidence. The objection begins with an addition to the case of Dr. N. The addition is that the vast majority of individual members of the scientific community, via some undescribed process, come to believe that not- d . Scientists collectively agree that not- d in conferences and workshops and they act on not- d by asserting it in lectures and publications and by approving drugs that depend on not- d (291). Lackey then argues that this implies that the scientific community itself believes that not- d and its belief that not- d defeats its putative knowledge that d .¹⁶

Lackey (293) draws two conclusion from this. First, social knowledge may be a lot less common than we would like. For many scientific propositions, there could be a number of scientists who hold conflicting beliefs that function as psychological defeaters. Second, if the belief that not- d defeats the knowledge that d , this implies an arbitrary asymmetry: group beliefs cannot justify other group beliefs, but they can defeat other group beliefs, which, she claims, is arbitrary without a story about why.

However, it is not clear why Lackey believes that in the case described, it is the belief that not- d that must defeat the knowledge that d . She does not, for example, claim that the belief that not- d is more justified than the belief (putative knowledge) that d . (I will not dispute that the scientific community believes that not- d in this example.) Nor does she explain why everyone believes that not- d . It is open for the proponent of **SK** to claim that the knowledge that d (or the structures on which it supervenes) defeats the belief that not- d . After all, the belief (putative knowledge) that d is, we are supposing, true, reliably formed, and easily accessible to all the relevant scientists, while the belief that

¹⁵ She does not mention externalistic defeaters. I return to this below. Here and throughout, when I speak of beliefs being defeated, this is shorthand for beliefs' justification is defeated.

¹⁶ Like Lackey, I will ignore the distinction between *prima facie* and *ultima facie* justification. This should not affect the points that follow.

not-*d* is a widespread falsehood. There are accounts of defeat where easily accessible but unpossessed evidence can be a defeater (e.g., Harman 1973) or where the availability of an alternative reliable process that, had it been used in addition to the actual process, would have led to a different belief can be a defeater (e.g., Goldman 1979).

SK seems to preclude the application of Harman's account of defeat. For, according to **SK**, Dr. N.'s result is known by the scientific community; but if Dr. N.'s result is knowledge that the scientific community already has, it is obviously not *unpossessed* evidence. However, it would still be unpossessed evidence for the community's members, in which case their *individual* beliefs would be defeated. Furthermore, this is evidence that, as relevant members of the scientific community, they *should* possess. If every member's belief that not-*d* were defeated, could the corresponding belief of the group they compose remain undefeated? **SK** seems to imply an affirmative answer, so I will not pursue Harman's account further.

What made the application of Harman's account of defeat seem to falter (i.e., that Dr. N.'s result is not unpossessed evidence) makes Goldman's account of defeat applicable and attractive. For, Goldman (1979) explicitly excludes gathering new evidence from alternative reliable processes and this might seem to preclude the application of Goldman's account to this case. However, while accessing Dr. N.'s result would be gathering new evidence for individual scientists, it would not be gathering new evidence for the scientific community. According to **SK**, accessing Dr. N.'s result that *d* is not the scientific community gathering new evidence; it is accessing knowledge it already has, though its members might not have. This social situation is somewhat akin to introspection, such as an individual accessing a memory of previously acquired evidence, which Goldman endorses as an alternative reliable process. Of course, this process of “social introspection,” accessing previous scientific results, must itself be a reliable process for the reliabilist account of defeat to work in this case. However, I think it is plausible that it is, given the kinds of truth-filtering mechanism that ensure condition ii of **SK**

(Bird 2010, 43–4). There is therefore a reliabilist story that is not ad hoc about why the belief that not-*d* is defeated in this example.¹⁷ This would also avoid Lackey's objection that, on **SK**, social knowledge is less common than we would like, since the knowledge that *d* is not defeated.

Suppose we adjust the case of Dr. N. so that all the members of the scientific community possess misleading evidence supporting their beliefs that not-*d*. Would not these individuals' beliefs or their group belief also constitute a defeater of the community's knowledge that *d*?¹⁸ The proponent of **SK** could still insist that the belief that not-*d* is defeated for the reasons given above (i.e., the presence of the alternative reliable process of social introspection; see also fn. 17), but I also do not think that claiming that there is mutual defeat *per se* in this case is a problem for **SK**. Recall that Lackey presents her example to suggest that there is an arbitrary asymmetry for the defender of **SK**: group beliefs cannot justify other group beliefs, but they can defeat other group beliefs. Since that is the conclusion I am trying to avoid, I think I can consistently accept that there is mutual defeat in this example, as long as such defeat is not conceived internalistically as effected by group (or individual) beliefs. Rather, the availability of Dr. N.'s publication may serve as a defeater for the group belief that not-*d* and the availability of other publications misleadingly supporting not-*d* may serve as a defeater for the group belief (putative knowledge) that *d*. Another way to put it is that there are two reliable processes that epistemically cancel out.¹⁹

However, if I admit the possibility of mutual defeat, I am still open to Lackey's objection that

¹⁷ I do not wish to claim that the existence of alternative reliable processes accounts for every case of group defeat. See, for example, Baker–Hyttch and Benton (2015), Beddor (2015), and Lasonen–Aarnio (2010) for arguments that this will not work in every case of defeat. Relatedly, there are also well-known problems of method individuation for reliabilism (i.e., the generality problem; see Conee and Feldman 1998; see Comesaña 2006 and Heller 1995 for some possible answers). I only claim that it is plausible that the defender of **SK** can use the reliabilist story to respond to Lackey's example (see also Goldman 2014). See Grundmann (2009) for a defense of a proper function version of the alternative reliable process account of defeat. Note that on Beddor's (2015, 157) non-reductive reasons account of defeat, according to which S's belief that *p* is defeated at *t* iff S has sufficiently good reason to abandon her belief that *p* at *t*, Dr. N.'s result that *d* (or the scientific community's knowledge that *d* that supervenes on it) is plausibly a sufficiently good reason for the scientific community to abandon its belief that not-*d*.

¹⁸ I thank two anonymous reviewers for pressing this objection.

¹⁹ I thank an anonymous reviewer for this way of putting it. I note that there is a problem of individuating reliable processes, but I cannot address this here (see the references in fn. 17).

social knowledge may be less common than we would like. It seems to me, though, that, given the extent of inconsistent evidence and controversy in scientific practice, this may be a result of any plausible account of group knowledge in science. It is also not clear that Lackey's (2016) positive account of justified group belief fares better than Bird's when it comes to the amount of scientific social knowledge it allows. In particular, the prevalence of disagreement in science seems to suggest that it will be rare that “a significant percentage of the operative members” of the scientific community justifiedly believe that p and that full disclosure of and rational deliberation about all the relevant evidence for p “would not result in further evidence that when added to the bases [i.e, reasons, reasons for reasons, and so on] of [the scientific community]’s members’ beliefs that p , yields a total belief set that fails to make sufficiently probable that p ” (381). However, an in-depth discussion of Lackey's positive account is beyond the scope of this paper.

The second response to Lackey is to deny the existence of group defeat, not because there is something wrong with *group* defeat, but because there is something wrong with defeat itself. Lackey says that if there is no defeater in this case, then this would remove groups “from the realm of the rational altogether” because “rebutting defeaters are precisely what rule out this combination of states from being epistemically permissible” (294). However, the defender of **SK** can account for the irrationality of this case by appeal to a kind of knowledge norm of theoretical reasoning similar to **KAP** above (Baker–Hyttch and Benton 2015; see also Lasonen–Aarnio 2010).²⁰ The norm would be something like Baker–Hyttch and Benton's knowledge norm of belief (Williamson 2000):

(**KNB**) One must: not believe that p if one does not know that p . (Baker–Hyttch and Benton 2015, 56)

The upshot of this strategy is that the defender of **SK** need not say that the scientific community's

²⁰ Bird would likely approve of this response to Lackey since he explicitly compares his functionalist approach to knowledge-first epistemology, writing that, “It is preferable to try to understand the analogy [between individual and social knowing] without such a commitment [to any particular analysis of individual knowledge] – particularly as there may be no true analysis of knowledge at all” (2010, 40).

knowledge that d defeats its belief that not- d (or vice versa). The intuition that something has gone wrong rationally in this case is better explained by the community's failure to adhere to **KNB**. The community possesses evidence that its belief that not- d is false (namely, Dr. N.'s publication), which means that it possesses evidence that its belief that not- d is not knowledge. In such a situation, **KNB** rules that the community must not believe that not- d . It is **KNB**, not rebutting defeaters, that rules out this combination of states from being epistemically permissible. Note that this strategy also avoids Lackey's objection that, on **SK**, social knowledge is less common than we would like, since the knowledge that d is not defeated.

Indeed, although the following claim is not necessary for this response to be an adequate one, the scientific community appears actually to adhere to something like **KNB**. (The norm could apply even if the scientific community does not adhere to it.) However, the scientific community seems consistently to engage in behavior that follows the norm, for example, by issuing retractions of published results that come under dispute and by not letting disputed beliefs guide important social actions.

5. Is the Scientific Community Merely in a Position to Know?

Lackey, having argued against the claim that the scientific community knows that d in the case of Dr. N., argues that the case is more accurately described as one in which the community is merely in a position to know that d . According to Lackey, the distinction between knowing and merely being in a position to know “is grounded, at least in part, in the difference between information that has been accessed and information that is merely accessible” (2014, 294). She gives the following example of an individual merely being in a position to know and claims that it is analogous to the community in the case of Dr. N. On my desk there is an unopened letter my friend's confession to a crime. I know nothing of the crime beforehand. In this case, I do not know that my friend committed the crime, though I am in a position to know by opening it and reading. The scientific community, Lackey claims, is in a similar

position with respect to Dr. N.'s publication.

The proponent of **SK** is likely to find Lackey's analogy question-begging, for it assumes that the information in the letter – the analog of Dr. N.'s publication – is not known; but whether Dr. N.'s publication is known by the community is what is under dispute. Let me end by offering a reconceptualization of this analogy in light of my arguments above. The case of the unopened letter is not analogous to the case of Dr. N.; an individual analog to the case of Dr. N. would be a case where the information in the unopened letter were part of the individual's knowledge. The individual is not merely in a position to know that *p*, but knows that *p*, although she might not be aware of that fact. The accessibility of Dr. N.'s publication makes it knowledge for the scientific community. It is analogous to an individual's accessible but currently unaccessed memory that *p*.²¹ The individual members of the scientific community are merely in a position to know that *d*, but the scientific community itself already does know that *d*.

6. Conclusion

The cases of Drs. Q. and N. are intended by Bird (2010; 2014) to demonstrate the non-supervenience of group knowledge on individual knowledge. Lackey (2014) argues that this non-supervenience cannot accommodate **KAP** (plus **GMAP**) and the role of group defeaters. I argued that “act as if” in **KAP** is ambiguous between “acting on” the knowledge that *p* and merely “acting in accordance with” the knowledge that *p*. On the latter disambiguation, **KAP** is false for individuals and groups. On the former disambiguation, **KAP** is true, but does not show that **SK** is false. Regarding group defeaters, I argued that the defender of **SK** could either claim that there is a reliabilist defeater (or two such defeaters) in Lackey's example or that, appealing to **KNB**, there are no defeaters. If this is right, then non-supervenience remains a viable position in the metaphysics of social knowledge.

²¹ A reviewer notes another analogy to individual memory: just as an individual will scan her memory for relevant evidence before embarking on important action, so will the scientific community scan what it knows (e.g., its published articles). See also my discussion of “social introspection” above.

Acknowledgments: Thanks to Carl Craver, Julia Staffel, and two anonymous reviewers for invaluable comments on earlier versions of this paper.

References

- Baker–Hyth, M. and Benton, M. A. (2015). Defeatism defeated. *Philosophical Perspectives*, 29(1), 40–66.
- Beddor, B. (2015). Process reliabilism's troubles with defeat. *Philosophical Quarterly* 65 (259), 145–159.
- Bird, A. (2010). Social knowing: The social sense of “scientific knowledge”. *Philosophical Perspectives*, 24(1), 23–56.
- Bird, A. (2014). When is there a group that knows? Distributed cognition, scientific Knowledge, and the social epistemic subject. In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 42–63). Oxford: Oxford University Press.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Carter, J.A. (2015). Group knowledge and epistemic defeat. *Ergo*, 2(28), 711–735.
- Comesaña, J. (2006). A well-founded solution to the generality problem. *Philosophical Studies*, 129(1), 27–47.
- Conee, E. and Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89, 1–29.
- Goldman, A.I. (1979). What is justified belief? In G.S. Pappas (Ed.), *Justification and Knowledge* (pp. 1–25). Dordrecht: Reidel.
- Goldman, A.I. (2014). Social process reliabilism: Solving justification problems in collective epistemology. In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 11–41). Oxford: Oxford University Press.

- Grundmann, T. (2009). Reliabilism and the problem of defeaters. *Grazer Philosophische Studien*, 79(1), 65–76.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Hawthorne, J. and Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Heller, M. (1995). The simple solution to the problem of generality. *Noûs*, 29(4), 501–515.
- Lackey, J. (2010). Acting on knowledge. *Philosophical Perspectives*, 24(1), 361–382.
- Lackey, J. (2014). Socially extended knowledge. *Philosophical Issues*, 24(1), 282–298.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review*, 125(3), 341–396.
- List, C. and Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Lasonen–Aarnio, M. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24(1), 1–21.
- Turri, J. (2011). Believing for a reason. *Erkenntnis*, 74(3), 383–397.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.