

BARBARA TRYBULEC*

VIRTUE EPISTEMOLOGY
AS A THEORY OF KNOWLEDGE-HOW:
APPLICATION OF INDIVIDUAL EPISTEMOLOGY
TO THE DEBATE OVER THE EXTENDED SUBJECT OF KNOWLEDGE**

Abstract

Knowledge is often regarded as the main concern of epistemology. It has gained the status of the most valuable epistemic aim. However, it is not easily explained why this should be the case. In this paper, I evaluate an answer to this question offered by virtue epistemology and show that the latest analyses imply that virtue epistemology cannot explain the special value of knowledge-that. The definitions of knowledge proposed by virtue epistemologists either apply only to knowledge-how, or cannot address the value problem. The aim of this paper is to show how virtue epistemology can account for the high epistemic status of knowledge-how of both individual and extended cognitive systems. I argue that, among extended cognitive systems, only those composed of two or more people can be regarded as subjects of knowledge-how. As far as systems that consist of a single person and her cognitive enhancement are concerned, virtue epistemology is a reductive theory of knowledge-how, as it attributes knowledge-how only to that person.

Keywords: virtue epistemology, knowledge-how, the value of knowledge, epistemic achievement, wide cognitive system

Although people try to attain a variety of cognitive aims, such as understanding, rationality, justification, truth, and empirical adequacy, knowledge is often regarded as the main concern of epistemology. But what is it about knowledge that makes it more respectable than other epistemic aims? In this paper, I evaluate an answer to this question offered by virtue epistemology. According to virtue epistemology, knowledge is valuable, because it is an

* Instytut Filozofii, Uniwersytet Marii Curie-Skłodowskiej, pl. M. Curie-Skłodowskiej 4, 20-031 Lublin, barbara.trybulec@poczta.umcs.lublin.pl.

** This work was realized as a part of the grant "Sonata" founded by National Science Centre, Poland, 2014/13/D/HS1/00689.

epistemic achievement. Unfortunately, this answer is inadequate when applied to knowledge-that, and succeeds only with respect to knowledge-how.

The aim of this paper is to show that virtue epistemology can account for the high epistemic status of knowledge-how, whether it is possessed by a single individual or an extended cognitive system. I argue that only those extended cognitive systems composed of two or more people can properly be regarded as subjects of knowledge-how. Virtue epistemology handles systems consisting of a single person and her cognitive enhancement by attributing knowledge-how only to that person. Should the analyses presented in this paper prove correct, they will constitute an example of how epistemological reflection can shed light on the philosophical debate over technologically and socially extended cognitive systems.

The paper has been divided as follows: In the first part, I present the general problem of the value of knowledge. In the second part, I discuss the debate between intellectualists and anti-intellectualists over the relation between knowledge-that and knowledge-how. I focus on the intellectualists' claim that knowledge-how does not require an ability to actually perform the known activity. Although I agree with this claim, I show that it does not challenge the anti-intellectualist position, because it is compatible with Ryle's construal of knowledge-how as a broadly understood complex disposition whose possession is consistent with not being able to actually perform the known activity. The third part of the paper introduces virtue epistemology as a theory of knowledge-how and explains why knowledge-how possesses intrinsic value. It also clarifies the relation between two kinds of epistemic achievement: knowledge-how and understanding-why, which are both defined by virtue epistemologists in terms of epistemic achievement. This leads to confusion, unless one accepts the false claim that knowledge-how and understanding-why are identical. My solution to this problem is to conceive of knowledge-how and understanding-why as two aspects of cognitive achievement that appear in different circumstances of an agent's cognitive activity. In the final part of the paper, I show how to apply the proposed solutions to extended cognitive systems, i.e. groups of people and systems composed of a single person together with her cognitive enhancement. I argue that virtue epistemology proves effective not only as an account of knowledge-how of individual cognitive agents, but also as a theory of extended cognitive systems.

1. THE PROBLEM OF THE VALUE OF KNOWLEDGE

The problem of the value of knowledge was first noticed by Plato in *Meno* and *Theaetetus*. It boils down to a single question: What is it about knowledge that makes it more valuable than mere true belief? In contemporary epistemology, some philosophers argue that this question is ill-conceived, as it rests on the unjustified assumption that knowledge constitutes a distinctively valuable epistemic standing (Kaplan 1985). Nonetheless, it is clear that we do indeed value knowledge more than mere true belief, and every viable theory of knowledge should explain this common intuition, even if it also aims to undermine it.

The problem of the difference in value between knowledge and mere true belief is referred to as “the Meno problem”. According to Plato, knowledge is more valuable than mere true belief, because it gives one confidence in the truth of one’s belief. In other words, in adverse circumstances, the subject can easily abandon mere true belief, which will not happen if she possesses knowledge (Pritchard, Turri 2014). This means that having knowledge, rather than mere true belief, can increase the likelihood of achieving one’s goals. To put it another way, knowledge possesses greater instrumental value than mere true belief. This intuition is still strong in epistemology, so much so that it lies at the core of virtue epistemology as the condition of safety.

Two more value problems concerning knowledge, besides the Meno problem, have been identified by Duncan Pritchard. *The secondary value problem* is the question of why knowledge is more valuable than any proper subset of its parts (Pritchard 2007: 87). It arose after Gettier’s (1963) seminal paper, which motivated many epistemologists to add another component to the standard tripartite definition of knowledge. This, however, raised the question of why knowledge with this additional component should be more valuable than justified true belief?¹

Another value problem identified by Pritchard concerns the question of why knowledge is distinctively valuable. If one wants to solve the value problem, it is not enough to show that knowledge is more valuable as a matter of degree than true belief and true justified belief. What we need is an account of why knowledge is more valuable in kind than other epistemic standings. The problem of a special and distinctive value of knowledge is called *the ter-*

¹ Mark Kaplan argues that justified true belief with an additional component does not possess special epistemic value over justified belief. The search for a solution to the Gettier problem is, by his lights, of no import to the understanding or improvement of rational inquiry (Kaplan 1985).

tiary value problem (Pritchard 2007: 104, note 4). To sum up, a viable account of knowledge should be able to resolve all three value problems mentioned above.

A tendency to undermine the value of knowledge is one of the main features of post-Gettier epistemology. The number of failed attempts at a definition of knowledge has led many epistemologists to ask why we should care about knowledge in the first place. The first theory of knowledge accused of failing to address the Meno problem was process reliabilism, which defined knowledge as true belief that arises out of a reliable belief-forming process (Goldman 1979). Linda Zagzebski argued that the reliability of a belief-forming process has no bearing on the value of the resultant true belief (Zagzebski 1996, 2003). The value of true belief is not determined by the qualities of the process from which it originates, no matter how reliable it may be. The only value of a reliable belief-forming process is that it produces true belief. If a belief is true, the reliability of the process that formed it, or lack thereof, is of no consequence. Reliabilism explains the value of knowledge by appealing to the value of a reliable belief-forming process, but fails to account for why true belief produced by a reliable process should be more valuable than one produced by a process that is not reliable.

There is a kind of reliabilism — John Greco’s agent reliabilism (Greco 1999, 2000) — that seems resistant to Zagzebski’s argument. According to Greco, it is not the reliability of a process that confers value on true belief, but its ability to constitute the agent’s “cognitive character”. A reliable belief-forming process is valuable not because it leads to true belief, but because it stems out of the agent’s cognitive processes, abilities, and dispositions. The argument is that true belief formed in virtue of the agent’s cognitive character is more valuable than true belief resulting from a lucky chance or the cognitive character of someone else. This line of thought is developed by proponents of virtue epistemology. Since I present virtue epistemology as a theory of knowledge-how, let me turn to the distinction between knowledge-that and knowledge-how before discussing the theory in more detail.

2. KNOWLEDGE-THAT AND KNOWLEDGE-HOW

The distinction between knowledge-that and knowledge-how overlaps with other distinctions, for example the ancient division between *episteme* and *techné*, or the contemporary distinctions between theoretical and practical or declarative and procedural knowledge. Roughly speaking, the first ele-

ments in the above distinctions are explained in terms of a knowledge of facts, expressible as a proposition about how things are. The second elements are associated with knowledge necessary to perform certain activities, which sometimes cannot even be expressed as a proposition, but is manifested in action. Discussed below are the main theses of two rival standpoints concerning the relation between knowledge-that and knowledge-how.

2.1. ANTI-INTELLECTUALISM

The proponents of anti-intellectualism argue that knowledge-that and knowledge-how are distinct in kind. The defining feature of knowledge-how is not that it corresponds with the subject's knowledge of a certain proposition, but rather with her ability to do something. In other words, if someone has an ability to perform an activity, she knows how to do it. Riding a bicycle is often cited as a case in point. What does it mean for someone to know how to ride a bicycle? Intuitively, knowing how to ride a bicycle implies having an ability to ride a bicycle. The statement that one knows how to ride a bicycle, but does not have an ability to do it, or that one has an ability to ride a bicycle without knowing how to do it, is recognized by anti-intellectualists as false, for knowledge-how always coexists with an ability, and vice versa.

The idea that knowledge-how is different in kind from knowledge-that is due to Gilbert Ryle (1949), who opposed the position called intellectualism, according to which knowledge-how is constituted, at least in part, by knowledge-that. Ryle showed that knowledge how to do something cannot be reduced to propositional knowledge of the proper way to perform the activity involved. And, contrary to intellectualists, he argued that to perform a certain activity skillfully, one need not first consider facts about how it is done.

Ryle points out, however, that knowledge-how is not always an ability the subject can utilize under any circumstances. Sometimes, it involves a sort of disposition to engage in different activities related to the action that the subject knows how to do. This means that knowledge how to, for example, tie a clove-hitch manifests itself not only in the action of actually tying clove-hitches, but also in "imagining tying them correctly, in instructing pupils, in criticizing the incorrect or clumsy movements and applauding the correct movements that they make, in inferring from a faulty result to the error which produced it, in predicting the outcomes of observed lapses, and so on, indefinitely" (Ryle 1949: 55).

Ryle incorporates the notion of responsibility into his concept of knowledge-how. He needs the concept of responsibility to be able to distinguish intelligent, skillful action from automatic behavior. Intellectualists explain

the difference between intelligent and automatic behavior by asserting that the former, unlike the latter, involves having propositional knowledge. Anti-intellectualists are forced to find another way of distinguishing these two kinds of behavior. Ryle argues that the term “intelligent” should be reserved for people who are responsible for their performance:

to be intelligent is not merely to satisfy criteria, but to apply them; to regulate one’s actions and not merely to be well-regulated. A person’s performance is described as careful or skillful, if in his operations he is ready to detect and correct lapses, to repeat and improve upon successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right (Ryle 1949: 29).

Knowledge-how, then, can be attributed only to responsible agents, capable of controlling and correcting their behavior. Ryle presented his arguments in the middle of the 20th century. Since then, most epistemologists have accepted the thesis that knowledge-how is different in kind from knowledge-that.

2.2. INTELLECTUALISM

The intellectualist turn in the debate over knowledge-how originated, at the beginning of the twenty-first century, with the work of Jason Stanley and Timothy Williamson entitled *Knowing How* (2001). By appeal to linguistic arguments, the authors demonstrate that to know how to perform a certain activity is just to know a set of specific propositions. For example, to know how to tie a clove-hitch is to know the method of tying a clove-hitch, or a set instructions dictating appropriate behavior. In other words, propositional knowledge-that governs the subject’s performance. The claim that knowledge-how is reducible to propositional knowledge is known as reductive intellectualism.² The somewhat more moderate intellectualist position, represented by John Bengson and Marc Moffett, does not identify knowledge-how with knowledge-that, but rather grounds the former in the latter (Bengson, Moffett 2007, 2011). The authors argue that to know how to perform an activity is to understand the method of doing it. Failures in this understanding lead to misguided propositional attitudes that deprive the subject of knowl-

² Stanley and Williamson argue that a subject has knowledge-how only if she knows a proposition concerning how a particular activity should be performed, and entertains this proposition under the practical mode of presentation. This means that the appropriate method of performing the action is presented to the subject not descriptively — by showing her instructions — but rather demonstratively — by showing the appropriate performance, for example: this thing that you are doing right now is how you ride a bicycle. If the subject knows that a certain method will allow her to ride a bicycle, under a practical mode, she knows how to ride a bicycle (Stanley, Williamson 2001).

edge-how. Bengson and Moffett define understanding in terms of a reasonable mastery of certain concepts (Bengson, Moffett 2007). For example, if a subject knows how to add, then she understands the concept of adding, which means that she knows its meaning and can guide her action accordingly. Understanding many concepts requires the ability to correctly employ them in the performance of an action. However, Bengson and Moffett argue that this is not always the case. There are activities, such as performing a salchow or complex ski stunts, such that knowledge how to perform them does not imply having the actual ability to perform them. Nevertheless, such actions require understanding, that is knowing the meaning of the relevant concepts, and the ability to use this knowledge in different corresponding activities. As an example, the authors cite world-class figure skaters who certainly know how to perform a quintuple salchow, yet since the jump is extremely athletically difficult, not all of them are actually able to execute it. They understand the concept of a quintuple salchow, which allows them to teach the jump, mark it, and do many different things that reveal their knowledge-how, but they are not able to do the salchow themselves (Bengson, Moffett 2007).

Other philosophers have also made and illustrated, if somewhat gruesomely, the observation that ability is not necessary for knowledge-how. Stanley and Williamson discuss the case of a pianist who has lost her arms (Stanley, Williamson 2001), Carl Ginet describes a violinist with damaged fingers (Ginet 1975), and Paul Snowdon presents a skilled chef who lost his arms (Snowdon 2004). All these authors argue that, intuitively, people who have lost their abilities still know how to do the activities they performed skillfully in the past. Knowledge-how does not manifest itself only in the performance of a known action, but also in many other associated activities. Hence, anti-intellectualists are wrong, when they claim that possession of the relevant ability is necessary for knowing how to perform a certain action.

It is doubtful, however, whether anti-intellectualists would argue in favor of this thesis. As I pointed out earlier, Ryle himself understands ability rather broadly, so it is not at all obvious that he would reject the above examples as illustrations of knowledge-how.

To strengthen the thesis that ability and knowledge-how come apart, intellectualists evoke situations where the subject has the ability to perform a given action, yet she seems not to possess the relevant knowledge-how. Bengson and Moffett present the case of Irina, who has a completely wrong idea about how a salchow is performed. However, she is also suffering from a severe neurological disorder due to which what she actually does is not what she believes herself to be doing. Whenever she attempts to perform a salchow her way, i.e. the wrong way, her disorder actually causes her to

perform it correctly. Hence, Irina is able to do a salchow, but, due to her mistaken belief about how to execute it, she does not in fact know how to do it. In terms of Bengson and Moffett's theory of knowledge-how, Irina, despite her ability to do a salchow, does not understand how to do it, she does not know the meaning of the concept "salchow", and this error in understanding deprives her of knowledge-how (Bengson, Moffett 2007). This observation fits in nicely with Ryle's broad understanding of an ability. Irina is only able to execute a salchow, yet she is not able to do any other corresponding activity. She is unable to teach someone else how to do a salchow, or identify someone who does it correctly. Her knowledge how to do a salchow is therefore largely diminished, even on Ryle's anti-intellectualist position.

The cases of knowledge-how without the ability to perform the known activity, and vice versa, pose a problem for anti-intellectualism only if it asserts that the subject's knowledge-how is grounded in nothing more than the fact that she can perform the known action. Indeed, in most circumstances, this is how things are. Counterexamples such as Irina's case are used by Bengson and Moffett to argue that knowledge-how requires intentional action, which in turn requires understanding, whereas understanding is explained in terms of knowing a relevant concept (Bengson, Moffett 2007). Arguing in favor of the connection between knowledge-how and concept possession places these authors in the company of moderate intellectualists when it comes to knowledge-how. In sum, they stress that it is not the relation between knowing-how and ability that epistemologists should focus on, but rather the relation between knowing-how and concept possession.³

The proponents of anti-intellectualism reply to these analyses by focusing on the condition that knowledge-how always coexists with propositional knowledge of how a given activity is performed. They challenge this thesis by citing evidence from cognitive psychology, suggesting that knowledge how a certain activity is done is not always linguistically represented (Devitt 2011, Adams 2009, Wallis 2008). A person riding a bicycle, for example, is rarely able to describe how she does it. On the other hand, one can be perfectly aware of how a bicycle ought to be ridden — that one needs to maintain proper balance on the seat while pedaling — although the fact that one can provide this theoretical description is hardly sufficient for one to actually know how to ride a bicycle. Hence, even if sometimes knowing how to do something entails propositional knowledge, one knows how to perform a given ac-

³ An interesting moderate intellectualist theory of knowledge-how is presented by Yuri Cath (2012). He argues that knowledge-how involves the relation of "seeming" to the relevant propositions. For example, a subject knows how to ride a bicycle only if there is a way of riding a bicycle such that it seems to be the way for the subject to ride a bicycle.

tion not in virtue of having that propositional knowledge, but rather in virtue of possessing certain abilities one possesses.

Proponents of virtue epistemology also argue against grounding knowledge-how in knowledge-that. To this end, they point to the discrepancy between knowledge-that and knowledge-how as far as Gettier-style cases are concerned. If knowledge-how is grounded in or reducible to knowledge-that, it should have the same epistemic properties as knowledge-that, which means that it should be equally vulnerable to the value problem. Virtue epistemologists claim, however, that this is not the case. Knowledge-how proves more resilient to Gettier-style counterexamples than knowledge-that, and this is where these two types of knowledge come apart (Carter, Pritchard 2013, 2015a, b).

3. VIRTUE EPISTEMOLOGY AS A THEORY OF KNOWLEDGE-HOW

Virtue epistemology emerged partly as a conciliatory position in response to the conflict between internalism and externalism toward epistemic justification, which dominated epistemological discussion in the second half of the past century. Epistemologists who appeal to cognitive virtues, which include broadly understood dispositions and traits of character, are able to combine the thesis of the virtues' internal location with their inaccessibility to reflection, or, inversely, the thesis of their partially extended location with the condition of their conscious endorsement (Carter, Pritchard forthcoming).⁴ The strategy has been introduced, in order to find new solutions to the Gettier problem and block the regress in the justification of beliefs. In an effort to explain our practice of epistemic evaluation, virtue epistemologists analyze whether knowledge is distinctively valuable, and if so, why. They show that successful inquiry depends upon the subject's epistemic character, which includes wisdom embodied in the judgments, cognitive habits, skills, and dispositions acquired through education and experience.

Interest in virtue epistemology was partly motivated by the growing dissatisfaction with post-Gettier epistemology. Instead of incorporating an additional condition into the definition of knowledge, virtue epistemologists proposed that, in order to know, the subject should form a true belief in vir-

⁴ Adam Carter, a proponent of virtue epistemology, argues that when cognitive enhancement is used, the subject's cognitive character — her epistemic virtues — extends partly beyond her body. However, the necessary condition of such an extension is that the subject, at some point, consciously endorse the enhancement as part of her cognitive character (Carter, Pritchard forthcoming). More about this in (Trybulec 2017b).

tue of the stable and reliable cognitive abilities making up her cognitive character (Greco 1999, 2000).

Unfortunately, not every true belief formed in virtue of the subject's cognitive ability constitutes actual knowledge. Imagine, for example, the following situation, originally proposed by Roderick Chisholm (1977). A shepherd is looking into the field and, using his cognitive abilities, forms the true belief that there is a sheep in the field. Suppose, however, that the shepherd is not looking at the sheep as such, but at a big, hairy, sheep-looking dog which is obscuring the view of the actual sheep in the field. Intuitively, the shepherd lacks knowledge, for the truth of his belief is just a matter of luck. Nevertheless, he is using his reliable cognitive faculties to form a true belief.

To address this problem, virtue epistemologists introduce an anti-luck condition — the safety principle that denies the subject's knowledge in situations where her true belief could have just as easily been false; in other words, in cases where her cognitive success derives from a lucky coincidence (Greco 2000). The shepherd in our example lacks knowledge, for his true belief could very easily have been false, which means that it is not a safe belief. Proponents of a more robust virtue epistemology suggest, however, that the anti-luck condition is redundant. It is sufficient, they claim, to introduce the requirement that the subject's cognitive success be achieved not only through her cognitive abilities, but because of them (Greco 2009, Sosa 1991, 2007, Zagzebski 1996).

This seemingly small change enables virtue epistemologists to deal with Gettier-style cases without the need to evoke the safety condition. A true belief that has been acquired because of an ability cannot be unsafe. The shepherd does not know that there is a sheep in the field, because his true belief has resulted not from exercising his cognitive abilities, but rather from a lucky coincidence.

According to John Greco, this is also the only way to answer the value problem. His proposal is based on two assumptions. Firstly, knowledge is a kind of achievement, where achievement is understood as success resulting from cognitive ability. Secondly, the value of achievements is final (non-instrumental). Hence, if knowledge is defined as a kind of achievement, it is intrinsically valuable. Greco recognizes this strategy as an answer to the three value problems presented in part 1 of this paper. Unfortunately, deeper analyses conducted more recently by virtue epistemologists reveal that knowledge cannot be defined in terms of cognitive achievement.

In most cases, knowledge and cognitive achievement are inseparable. One possesses knowledge if one's cognitive success (true belief) is the result of exercising one's cognitive abilities. This harmonizes with our intuitions, for

we are unwilling to attribute knowledge to someone who has acquired true belief by accident (as in Gettier's cases). Achievements seem to be valuable in their own right, simply because they require a kind of effort from the subject. Knowledge as true belief that has resulted from a subject's cognitive abilities is something for which the subject deserves credit, because she must exercise her cognitive virtues to achieve it.

It turns out, however, that knowledge cannot be defined in terms of cognitive achievement, for achievement constitutes neither a necessary nor a sufficient condition of knowledge. There are rare cases where the subject cannot be credited with knowledge in spite of cognitive achievement, and there are also cases where, intuitively, the subject possesses knowledge without having exercised the relevant cognitive abilities. Let us analyze these two situations more closely.

We can introduce the first by exploiting the shepherd's case mentioned above as an example of Gettier-style epistemic luck, also known as intervening epistemic luck. The lucky fact that there was a sheep behind the dog "intervened" between the subject's cognitive performance and his cognitive success. If not for this convenient coincidence, the subject would not have been able to form a true belief. The subject's cognitive success was not accomplished in virtue of his cognitive abilities, which means that it did not constitute knowledge.

Imagine, however, a slightly different situation. This time the shepherd saw a genuine sheep, although he might very easily have been wrong, because there are many sheep-shaped objects in the vicinity, any one of which would have led the shepherd to form a false belief that he saw a sheep. The shepherd was very lucky to look at the only genuine sheep in the field. Again, our subject is cognitively successful, for he has formed a true belief and, in order to do that, he has used his reliable cognitive abilities. Nevertheless, under these circumstances, his true belief could very easily have been false, so, like in the first case, we should conclude that the shepherd does not know that there is a sheep in the field.

However, this conclusion is somewhat less intuitive than in the first case, as it is hard to deny the shepherd's cognitive achievement. Although he is still very lucky, his cognitive success is due to his stable and reliable cognitive abilities rather than sheer chance. The lucky chance is not needed in this case. This situation is one of environmental epistemic luck, where cognitive achievement does not guarantee the safety of one's true belief. Knowledge, however, has to be immune to both kinds of epistemic luck, for otherwise it will be vulnerable to Gettier's counterexamples. Robust virtue epistemologists believe that this immunity can be secured by the condition that cogni-

tive success stem from ability. Unfortunately, this is not enough to guarantee knowledge. The shepherd in the second scenario is credited with an achievement, but not with knowledge (his true belief is not safe).

The second type of case where knowledge and cognitive achievement come apart relates to situations where the subject is credited with knowledge, although she deserves very little credit for acquiring true belief. An example of such a case is offered by Jennifer Lackey (2008). Imagine a person who, upon arriving at the station in Chicago, asks the first passer-by for directions to Sears Tower. The passer-by gives her the correct directions, so she forms a true belief about the location of Sears Tower. The source of information is reliable, the true belief is safe and, hence, intuitively we have no reason to deny the visitor's belief the status of knowledge. But can we regard it as a cognitive achievement? Does the visitor deserve credit for forming a true belief? This is not so obvious. The visitor uses her epistemic abilities to a very limited extent, it is in virtue of the informant's cognitive abilities that the visitor's belief is true.⁵ The same is true of basic perceptual beliefs. Since the subject of such beliefs uses only her perceptual processes, it is doubtful whether this cognitive success might count as an achievement. However, if perceptual beliefs are true and safe, there seems to be no reason not to recognize them as knowledge.

The problematic cases mentioned above show that achievement is neither necessary nor sufficient for knowledge. Hence, it seems that it is not achievement that confers the distinctive value on knowledge, but rather immunity to epistemic luck (Pritchard 2007).⁶ Unlike knowledge, achievement is compatible with environmental epistemic luck, and this is why virtue epistemologists are unable to formulate an adequate response to the value problem in terms of achievement. It seems, therefore, that virtue epistemologists will have to give up on the intimate connection between knowledge and cognitive achievement. They can venture at best that the close connection allows knowledge to be treated as valuable.

The definition of knowledge as an achievement accomplished through ability was introduced by robust virtue epistemologists as a remedy to the Gettier problem (Sosa 1988, Greco 2009, Zagzebski 1996). But, as Pritchard

⁵ The subject in this case is employing certain kinds of epistemic virtue. She does not ask any person on the street, but rather she seeks an informant who seems to be reliable. Nevertheless, the true belief is still of very minimal credit to her. For this reason, it is hard to attribute cognitive achievement to her.

⁶ As I have already mentioned, this observation was first made by Plato in *Meno*, where he points to stability and being "tethered to the fact" as the features distinguishing knowledge from mere true belief.

suggests, we can solve the Gettier problem without identifying knowledge with cognitive achievement. The necessary condition can be expressed in the form of a safety principle that blocks the attribution of knowledge to subjects in problematic Gettier's cases (Pritchard 2009). Hence, Pritchard proposes a return to Greco's early idea where knowledge was depicted as true and safe belief formed via the subject's stable and reliable cognitive abilities that make up her cognitive character (Greco 1999, 2000). This definition releases knowledge from the problematic requirement of cognitive achievement.

Such a solution, however, calls into question the intrinsic value of knowledge, for what could now provide knowledge with value? There are two possible ways to answer this question. Firstly, one could point to the safety principle as the essence of the epistemic value of knowledge. Namely, it ensures that the subject of knowledge, unlike the subject of mere true belief, could not have easily been wrong, which is a very valuable feature not only in everyday life but also as far as Gettier-style cases are concerned. Secondly, virtue epistemologists could choose a negative response to the value problem, and claim that it is not knowledge that is distinctively valuable but cognitive achievement — a distinct epistemic standing that could, in some rare situations, come apart from knowledge (Pritchard 2009). Recently, however, a new movement has emerged within virtue epistemology with the aim of justifying the intrinsic value of knowledge by renewing the intimate relation between knowledge and cognitive achievement. In the remaining part of this section, I will discuss the most recent work addressing the intrinsic value of knowledge within virtue epistemology.

The above reflections, which lead to a divorce of knowledge from cognitive achievement, appeal to a standard notion of knowledge as a kind of true belief. Their results concern, therefore, only knowledge understood in this sense. From the anti-intellectualist point of view, however, a virtuously formed belief is not the only kind of knowledge there is. And the other kind — knowledge-how — possesses very different epistemic features that may very well secure its intrinsic value. According to recent work by virtue epistemologists, knowledge-how differs from knowledge-that with respect to resilience to epistemic luck (Carter, Pritchard 2013, 2015a, b). In fact, they argue by appeal to intuitions, knowledge-how is indistinguishable from cognitive achievement. The recognition of this distinctive feature of knowledge-how motivates the claim that it is distinct in kind from propositional knowledge. The difference becomes evident in situations of environmental epistemic luck.

Imagine Charlie, who wants to learn how to change a light bulb, so he consults *The Idiot's Guide to Everyday Jobs* and finds there an accurate set of instructions. On this basis, he forms a true belief about how light bulbs are

changed. However, Charlie is extremely lucky to have formed a true belief, for the disgruntled author of the guide had filled the book with misleading instructions. Only a random computer error at the printing house caused the text under the entry for “Changing a Light Bulb”, in just this one copy of the book, to be replaced by new text. By incredible coincidence, this new text provided the clear and accurate set of instructions that Charlie would later consult (Cath 2012). Does Charlie know how to change a light bulb? Surely, he does not know that the instruction in the guide is correct, for his true belief to that effect is not safe. The intuition whether he knows how to change a light bulb is not clear, though Cath — the author of the scenario — concludes that he does. This is a situation of intervening epistemic luck. If not for the lucky chance, Charlie’s belief about the correctness of the instructions in the guide would be false.

Let us try, however, to follow Pritchard and Carter’s cue and modify this case a little bit, so as to change it into an example of environmental epistemic luck. Imagine that Charlie consults a reliable guide, though surrounded by fakes. Had he consulted one of the fakes, he would have learned incorrect instructions. In this case, the intuitions are clear — Charlie knows how to change a light bulb, although he does not know that the instructions in the guide are correct (his true belief about it is not safe). Intuitively, we are unwilling to deny him the ability to change a light bulb. He can skillfully change a light bulb without having the relevant propositional knowledge. Here, the two types of knowledge come apart, in the same manner as knowledge-that and cognitive achievement come apart in cases of environmental epistemic luck. In the modified version of the lucky bulb scenario, Charlie can be credited with cognitive achievement regardless of the lucky circumstances, his success has resulted not from a lucky chance, but rather from his own cognitive abilities. Cognitive achievement, although conceptually unrelated to knowledge-that, is thus sufficient for knowledge-how. But does it constitute a necessary condition too? As remarked above, it is not necessary for knowledge-that, for it is absent in cases of testimonial knowledge. To demonstrate how things are, as far as knowledge-how is concerned, virtue epistemologists modify the lucky bulb case one more time.

Imagine that Charlie finds out how to change a light bulb from a person he recognizes as reliable, and he passes this information onto his young son. Since the son trusts his father, he believes that the information acquired from him is reliable. Does the son know that this is the correct way to change a light bulb without knowing how to do it? Pritchard and Carter argue that he does not. In their opinion, the ability to produce a certain outcome is not sufficient for knowledge-how. What needs to be added? Intuitively, to know how

to change a light bulb, the son's successful performance of the action should be the result of his exercising his cognitive abilities. As I already pointed out, testimonial knowledge based on mere trust is not such a result. Charlie knows how to change a light bulb, because he had used his cognitive abilities to verify the reliability of his informant. His son, however, simply trusts his father. His cognitive success is not the result of his cognitive ability, but rather a mere consequence of his father's passing down the information onto him. Pritchard and Carter suggest that intuitions regarding knowledge-how are more demanding in the case of testimonial knowledge than intuitions regarding knowledge-that (Carter, Pritchard 2015a). The reason is that it seems impossible to imagine knowledge-how without cognitive achievement and vice versa, as if they were conceptually indistinguishable.

But is knowledge-how really the only kind of possible cognitive achievement? It must be, if cognitive achievement is sufficient for knowledge-how, and the lucky bulb example shows that it is. Specifically, by crediting Charlie with cognitive achievement, we automatically credit him with knowledge-how. But is there no other epistemic standing that can be conceived as a cognitive achievement? If there is, then it must be concurrent with knowledge-how and with cognitive achievement, and, at the same time, it should possess distinctive epistemic features. Is such a state at all possible?

It seems that the only way to conceive of it is to interpret knowledge-how and another cognitive standing as two different aspects of cognitive achievement — two sides of the same coin. I mention this, because virtue epistemologists often refer to understanding-why as a valuable epistemic standing, different from knowledge-that, yet sharing the essential features of knowledge-how (Pritchard 2009, Carter, Pritchard 2015a). Pritchard and Carter point out three such commonalities.

Firstly, understanding-why, just as knowledge-how, is distinctively valuable as "a type of cognitive achievement". Moreover, understanding-why, again much like knowledge-how, is not, contrary to what Bengson and Moffett as well as Kvanvig (2003) would claim, a propositional attitude, but rather a process of capturing certain relations between the phenomena, or seeing that such-and-such is the case (Pritchard 2007). Next, understanding-why, just like knowledge-how, is immune to environmental epistemic luck (Pritchard 2007, Carter, Pritchard 2015a). Hence, cognitive achievement is a necessary and sufficient condition for two distinct epistemic standings. Should they not be identical?

They are clearly not. Knowledge-how does not always coexist with understanding-why. A subject can know-how to ride a bicycle, or how to whistle, and yet be unable to say why she is able to do it, for she does not understand the

relevant relations. On the other hand, one can understand why one's house has burned down, although it has nothing to do with performing an action (Carter, Pritchard 2015a). I propose to construe understanding-why and knowledge-how as two aspects of cognitive achievement that appear in different circumstances. Where an activity is performed by a subject, cognitive achievement is identical to knowledge-how, but where one seeks to know why something has happened, cognitive achievement takes the form of understanding.

Nevertheless, if cognitive achievement is not always identical with knowledge-how, could it constitute its sufficient condition? Generally, the answer must be in the negative. Cognitive achievement is sufficient for knowledge-how only conditionally, during an activity. Hence, if a subject is performing an activity, and if she achieves cognitive success, because of her cognitive abilities, she knows-how to do it. However, if the subject's cognitive success consists in grasping the appropriate relations between relevant phenomena, she is not credited with knowledge-how, but rather with understanding-why. Hence, cognitive achievement is also, equally conditionally, sufficient for understanding-why. I allow the possibility that there are other aspects of cognitive achievement, for example wisdom, which may reveal themselves to be significant under specific circumstances, but an analysis of their epistemic features and relationship with other epistemic standings is beyond the scope of this paper. The crucial point I am trying to make is that virtue epistemology provides a solution to the value problem only insofar as knowledge-how is concerned. Thus, although cognitive achievement and knowledge-that rarely come apart, for the sake of precision, virtue epistemology should be defined as a theory of knowledge-how.

4. EXTENDED COGNITIVE SYSTEMS AS SUBJECTS OF KNOWLEDGE-HOW

Some virtue epistemologists apply solutions developed within individualistic epistemology to analyses concerning extended cognitive systems composed of one person and her cognitive enhancement, or a number of people. In the final part of this paper, I would like to bring up the issue of such systems treated as extended subjects of knowledge. The theses of extended cognition and extended mind proposed by Clark and Chalmers in *The Extended Mind* (Clark, Chalmers 1998) opened up a new field of inquiry, which virtue epistemologists have quickly recognized. Many works have been published recently on the topic of distributed knowledge, wherein the subject is no longer con-

ceived as a single person, but rather as an extended system composed of many human and non-human elements (Pritchard 2010, Palermos, Pritchard 2013, Palermos 2016, Palermos 2016a, Tuomela 2011, Goldman 2014). An overwhelming majority of these works are dedicated to propositional knowledge, understood as a kind of mental state. Very little attention is being paid to knowledge-how, which, being intrinsically valuable, should become a relevant point of interest.

The idea that a person and her cognitive enhancement can constitute a single cognitive system with its own cognitive processes or even with its own mind is called active externalism. It comes in two versions: the extended cognition thesis and the extended mind thesis. According to the former, some cognitive processes can be realized partly by elements of the world outside a person's body. The extended mind thesis is more radical and suggests that some mental states, such as beliefs, supervene partly on the cognitive artifacts people use to solve cognitive problems.⁷ Roughly, Clark and Chalmers argue in favor of the idea that cognitive processes and states can belong not only to individual people, but sometimes they should be attributed to a wider cognitive system, composed of human and non-human elements. If this is true, propositional knowledge, as a kind of belief, could, under the appropriate circumstances, be attributed to extended cognitive systems.

This thesis, however, can be called into question. As I have argued elsewhere (Trybulec 2017b), there are two reasons why extended cognitive systems composed of a person and her cognitive enhancement are not well suited to be subjects of both knowledge-that and knowledge-how. Firstly, there are a number of objections that the extended mind thesis must face. The most common of these is known as cognitive bloat. According to this objection, Clark and Chalmers' arguments imply that every piece of information stored in an external artifact is part of the subject's beliefs. Admittedly, authors of *The Extended Mind* try to avoid these absurd consequences by incorporating four conditions known as "trust and glue conditions".⁸

⁷ To illustrate the thesis that the mind extends into the world, Clark and Chalmers present the example of Otto who possesses an external belief: Inga hears that there is an interesting exhibition at the Museum of Modern Art and decides to go there, so she retrieves the location of the Museum from her biological long-term memory. Otto has Alzheimer's and depends on his notebook for information storage. When he hears about the exhibition and decides to go there, he retrieves the address of the Museum from his notebook. For Otto, the notebook plays the same role as biological memory does for Inga. Information in the notebook functions just like the information that constitutes ordinary non-current beliefs, only its location is different (Clark, Chalmers 1998: 12).

⁸ The conditions are as follows. Firstly, the cognitive artifact must be a constant in the subject's life. Secondly, the information in the extended storage must be directly and read-

Yet critics have recognized this strategy as ineffective, showing that the conditions are neither necessary nor sufficient for extended cognition (Adams, Aizawa 2001, Rupert 2009). One of the consequences of an expansion of a subject's beliefs is the expansion of her actions entailed by such beliefs. Critics of the extended mind thesis create scenarios that reveal that this is absurd. For example, imagine a robot that serves as external storage for my beliefs. It acts according to them, even when I am sleeping. This produces new states of affairs; for example, a cake may appear on the table as a result of the robot's activity. Since the robot acted on my beliefs, I would be entitled to take full credit for making the cake (Gertler 2007).

The cognitive bloat argument is very persuasive and stimulating for the imagination, but it is not the only reason against attributing knowledge to a system composed of a single human and her cognitive enhancement. Another reason involves cognitive agency.

The standard conception of agency in contemporary analytic epistemology is most commonly associated with the work of Gertrude Anscombe (1957) and Donald Davidson (2002). Although their views differ in many respects, they both assert that agency ought to be explained in terms of intentional action. In other words, a cognitive system can be regarded as an agent just in case it has the capacity to act intentionally. According to Davidson, intentional action requires mental states such as desires, beliefs, and intentions. Critics point out, however, that this condition of agency is too stringent, for there are beings capable of genuine agency that do not possess representational mental states (Schlosser 2015). For example, we use a very effective strategy of explaining and predicting animal behavior in terms of intentional mental states. But is this appropriate? According to Daniel Dennett's instrumentalist stance, it is, because successful prediction of behavior is everything that is needed in this respect (Dennett 1987). Still, most proponents of the standard theory of agency are realists towards mental states. This means that agency can be attributed only to a system that possesses the right internal states with the right representational contents (Schlosser 2015). Clarification of this vague claim is one of the most important challenges for philosophers and cognitive scientists, and it is clearly beyond the scope of this work. Here, it is enough to conclude that, according to the standard conception of agency, to which I appeal in this paper, cognitive agency can be attributed only to humans, for only humans satisfy the required conditions of intentionality.

ily available. Thirdly, upon retrieving the information from the external storage, the subject automatically endorses it. Fourthly, the information in the artifact has been consciously endorsed at some point in the past and, indeed, is there as the consequence of such endorsement (Clark, Chalmers 1998).

I am aware, however, that cognitive scientists often explain cognitive agency in a less restrictive way, by identifying it with an ability to act and interact with the environment, which is based on information processing, and serves cognitive development and learning.

The thesis that a system composed of a person and a cognitive artifact can constitute the subject of cognitive processes is widely accepted, and I do not want to question it. Such a system can support cognitive processes that none of its elements would be able to support independently. As such, they cannot be attributed to the human element alone. As argued by the proponents of the dynamical systems theory, the subject of cognitive processes is constituted by an extended coupled system that emerges from a reciprocal causal relationship between the human and the artifact. Without prejudice to the above, however, we do not need to accept the thesis that cognitive states, the products of the system's cognitive processes, also belong to this system as a whole. If one recognizes an agent making telescopic observations or performing complex calculations as a subject of mental states, one naturally attributes an intention to this agent, along with a belief that is the product of her cognitive process. Such an agent cannot be constituted by a person and a cognitive artifact, since this kind of an extended cognitive system cannot as a whole be recognized as the subject of intentions, beliefs, desires, doubts, and other propositional states that motivate cognition. Ascribing such mental states to this kind of a cognitive system contradicts not only common intuitions, but also, as I have shown, common usage of these categories is philosophy.

Application of the standard conception of agency to the debate over extended one-person systems leads to the conclusion that cognitive agency of such a system must be reduced to the cognitive agency of one of its elements, namely the human being. The reason is that only this element in an extended one-person system exhibits intentionality and initiative, plans cognitive processes, and possesses appropriate beliefs and desires. As far as this kind of an extended system is concerned, cognitive agency is not extended. Yet there is one possible situation in which cognitive agency can be attributed to this kind of an extended system as a whole. This is the case where a person interacts with an artificial intentional system, that is a system exhibiting propositional attitudes, such as beliefs, desires, and intentions. I do not wish to determine whether an artificial system can satisfy standard conditions of agency; what I want to stress, however, is that an extended cognitive system composed of a person and an intentional artifact should be recognized as a group cognitive system.

So far, I have considered the possibility of crediting extended one-person cognitive systems with knowledge-that. My conclusion, in this respect, is

negative. Should it also be negative as far as knowledge-how is concerned? Can such a system be credited with a cognitive achievement? Elsewhere, I have presented a number of arguments supporting the negative answer to the latter of these questions (Trybulec 2017b). In short, the cognitive character responsible for gaining knowledge can only be attributed to the human component of a given system. An artifact cannot be the subject of epistemic credit, or responsibility, for it is not an intentional entity (Goldberg 2012). Hence, an extended one-person cognitive system cannot be the subject of knowledge understood as cognitive success resulting from exercising her cognitive abilities. Does that mean that virtue epistemology cannot explain knowledge-how in extended cognitive systems? As far as an extended one-person cognitive system is concerned, virtue epistemology is a reductive theory of knowledge-that and knowledge-how. It is the person who is the subject of beliefs and cognitive abilities, even if these abilities are extended, for they supervene not only on biological processes taking place inside the human body, but also on non-biological processes occurring in an artifact and on the mutual relations between them all. Such a reductive line of thinking about knowledge-how is not necessary with respect to a group subject of cognition. I will focus on the latter issue below.

As I have already mentioned at the beginning of this section, the issue of group knowledge-how, unlike the issue of group knowledge-that, is rarely taken up in contemporary epistemology. However, in everyday life, we often attribute knowledge-how to groups, saying for example that Volkswagen knows how to make good cars, Real Madrid knows how to win football games, or a musical band knows how to improvise music. Is such knowledge reducible to the knowledge of individual members of the group? It seems that it is not, for none of the individual members of these groups knows how to perform the relevant activities. In other words, none of them could perform it alone by exercising exclusively her own cognitive abilities, and hence none of them can be individually credited with the group's success. Consequently, groups such as orchestras, sports teams, and corporations are subjects of group knowledge-how. Group's cognitive ability, which determines its knowledge-how, emerges from the cognitive abilities of particular group members interacting with each other in accordance with the specific organizational structure of the group. The same can be said about group belief.⁹ Therefore, both kinds of group knowledge can be explained similarly, the only difference

⁹ Researchers addressing the issue of group intentionality often explain group beliefs by appealing to the theory of judgment aggregation, which points to group members' beliefs and on the group's organizational structure as a realization base of group beliefs (List, Pettit 2011, Pettit 2007, Tollefsen 2004, Theiner, O'Connor 2010, Tuomela 2011).

being that, in the case of group knowledge-that, the realization base is composed partly by the group members' beliefs and, in the case of group knowledge-how, by the group members' cognitive abilities.

Intellectualists would explain group knowledge-how by arguing that a group as a whole possesses propositional knowledge of how a given activity should be performed, which none of its individual member can claim to possess. Anti-intellectualists, on the other hand, are likely to depict group knowledge-how in terms of the group's abilities and epistemic responsibility, which is irreducible to the group members' aggregated epistemic character. All of this, however, is a challenge to our intuitions, since according to this holistic position, no member of a music group knows how to improvise jazz while playing in a band, and no individual working for Volkswagen knows how to make a car. Each individual in these groups knows how to do her part, but no individual knows the way the group activity should be performed, and no one has the abilities necessary to perform it individually.¹⁰ Consequently, no one should be credited with the knowledge how to make a car, or how an improvised piece is played by a band. Knowledge possessed by the group is in constant mutual interaction with actions undertaken by group members, it emerges from them and determines them, even though no member of the group personally knows the relevant propositions, or is able to perform the group activity alone.¹¹

The suggestion that group knowledge-how exists, even though none of the group members know the exact method of performing given activities (e.g., jazz band improvisation, making a car), poses a problem for intellectualists who argue that the subject of knowledge-how must know how an activity is to be performed. As I have pointed out in part 2, anti-intellectualists do not impose such a condition, which gives them an advantage in the debate over group knowledge-how. Instead of explaining how group propositional attitudes emerge from individual group member beliefs, anti-intellectualists appeal to group abilities that are distributed between the group members interacting with each other. Proponents of this concept of group knowledge-how often appeal to the dynamical system theory — a mathematical framework for studying the behavior of dynamical systems. The theory asserts that two or

¹⁰ I have adapted the example of a car company from Tollefsen and Palermos. They explain group knowledge-how by discussing how a Corvette is made (Palermos, Tollefsen forthcoming).

¹¹ As Tuomela convincingly argues, group members should jointly accept that the way they perform their activities constitutes in part a way in which group activity should be performed. Group members should jointly accept the way group activity is performed, even though, taken separately, they may not know or believe what that way is (Tuomela 1992).

more systems give rise to a distributed process or ability, when they are reciprocally related to each other (Chemero 2009, Palermos 2014, 2016). In order to attribute knowledge-how to a group, the group has to possess its own epistemic character, composed of its cognitive abilities. Since a group is understood as a dynamical, coupled system, its cognitive character is distributed over all of its parts in such a way that it seems impossible to isolate separate parts responsible for specific abilities. Hence, group knowledge-how, as a result of distributed cognitive abilities, is a cognitive achievement that cannot be attributed to any of its parts individually, but only to a group as a self-standing cognitive agent. Thus, group knowledge-how, as well as group epistemic responsibility, and other epistemic properties characteristic of the cognitive subject, should be understood as emergent phenomena arising from the members' knowledge-how embedded in the group's organizational structure.

It seems, therefore, that some groups, unlike extended one-person cognitive systems, can be recognized as genuine subjects of knowledge-how. Why is that? The difference between these two kinds of extended cognitive systems is as follows. A group is composed of two or more intentional agents, so it is possible to aggregate their intentions according to a certain strategy, so as to obtain a group belief or ability. In the case of extended one-person cognitive systems, this is not possible, as there is only one intentional agent who cannot constitute the realization base for the system's beliefs and abilities with anyone else. She is the only subject of mental states, cognitive abilities, and other epistemic virtues responsible for knowledge. Hence, the crucial difference between the two kinds of extended cognitive systems is that an extended one-person system can be split into the cognitive agent and the rest of the system, whereas a group cannot. In other words, group cognitive agency can be identified with an abstract organizational structure realized by group members interacting with one another according to the rules determined by that structure.

To recapitulate, applying the solutions of virtue epistemology to the debate over extended cognitive systems faces fewer difficulties in the case of groups. The reason is that, as far as extended one-person systems are concerned, there is no possibility to distribute mental states or epistemic credit over the system's parts. The flaws of the extended mind thesis cause problems only to those who tend to attribute knowledge to an extended one-person system as a whole. Epistemologists working on group knowledge are safe in this respect. A combination of virtue epistemology, the theory of dynamical systems, and distributed cognition results in a theory of group subject of knowledge-how that is capable of explaining and predicting groups'

cognitive activities and determining the criteria of their epistemic evaluation. Virtue epistemology is also a viable reductive account of knowledge-how of extended one-person systems.

CONCLUSION

In the present paper, I have considered the issue of the value of knowledge from the perspective of virtue epistemology. I have concluded that the proper way of thinking about this position is in terms of the theory of knowledge-how. This does not mean that virtue epistemologists have nothing to say in the debate on the conditions of knowledge-that. Their work in this field has contributed to the development of contemporary epistemological thought to a great extent. Indeed, their conditions of knowledge constitute the starting point for most epistemological analyses. Virtue epistemology is unsuccessful only in one respect — it cannot account for the intrinsic value of knowledge-that. However, it has the resources to explain the valuable epistemic status of knowledge-how. In the paper, I have demonstrated the difference between the two kinds of knowledge relative to the ongoing debate between intellectualists and anti-intellectualists. I have argued in favor of anti-intellectualism, as most virtue epistemologists do. However, the main contribution of the present work to contemporary epistemology is the thesis that virtue epistemology is a viable theory of knowledge-how for extended cognitive systems, although it takes a reductionist approach toward the knowledge-how of extended one-person cognitive systems. Owing to the analyses presented in this paper, virtue epistemology appears to be a viable theory of knowledge-how in the context of both individual epistemic subjects and wide, technologically or socially embedded, cognitive systems.

BIBLIOGRAPHY

- Adams F., Aizawa K. (2001), "Bounds of Cognition", *Philosophical Psychology* 14(1), 43-64.
Adams M. (2009), "Empirical Evidence and the Knowledge-That/Knowledge-How Distinction", *Synthese* 170(1), 97-114.
Anscombe G. E. M. (1957), *Intention*, Oxford: Basil Blackwell.
Bengson J., Moffett M. (2007), "Know-How and Concept Possession", *Philosophical Studies* 136(1), 31-57.

- Bengson J., Moffett M. (2011), *Two Conceptions of Mind and Action: Knowing How and the Philosophical Theory of Intelligence* [in:] *Knowing How: Essays on Knowledge, Mind, and Action*, J. Bengson, M. Moffett (eds.), Oxford: Oxford University Press, 3-58.
- Carter A., Pritchard D. (2013), "Knowledge-How and Epistemic Luck", *Noûs* 49(3), 440-453.
- Carter A., Pritchard D. (2015a), "Knowledge-How and Cognitive Achievement", *Philosophy and Phenomenological Research* 91(1), 181-199.
- Carter A., Pritchard D. (2015b), "Knowledge-How and Epistemic Value", *Australasian Journal of Philosophy* 93(4), 799-816.
- Carter A., Pritchard D. (forthcoming), "The Epistemology of Cognitive Enhancement", *Journal of Medicine and Philosophy*.
- Cath Y. (2012), *Knowing How Without Knowing That* [in:] *Knowing How: Essays on Knowledge, Mind and Action*, J. Bengson, M. Moffett (eds.), Oxford: Oxford University Press, 113-135.
- Chemero A. (2009), *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.
- Chisholm R. (1977), *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice-Hall.
- Clark A., Chalmers D. (1998), "The Extended Mind", *Analysis* 58(1), 7-19.
- Davidson D. (2002), *Agency* [in:] *Essays on Actions and Events*, Oxford: Clarendon Press, 43-61.
- Dennett D. C. (1987), *The Intentional Stance*, Cambridge MA: MIT Press.
- Devitt M. (2011), "Experimental Semantics", *Philosophy and Phenomenological Research*, 82(2), 418-435.
- Gertler B. (2007), *Overextending the Mind* [in:] *Arguing about the Mind*, B. Gertler, L. Shapiro (eds.), New York: Routledge, 192-206.
- Gettier E. L. (1963), "Is Justified True Belief Knowledge?", *Analysis* 23(6), 121-123.
- Ginet C. (1975), *Knowledge, Perception, and Memory*, Boston: Reidel.
- Goldberg, S. C. (2012), "Epistemic Extendedness, Testimony and the Epistemology of Instrument-Based Belief", *Philosophical Explorations* 15(2), 181-197.
- Goldman A. (1979), *What Is Justified Belief?* [in:] *Justification and Knowledge*, G. S. Pappas, (ed.), Dordrecht: Reidel, 1-25.
- Goldman A. (2014), *Social Process Reliabilism: Solving Justification Problems in Collective Epistemology* [in:] *Essays in Collective Epistemology*, J. Lackey (ed.), Oxford: Oxford University Press, 11-43.
- Greco J. (1999), "Agent Reliabilism", *Philosophical Perspectives* 13, 273-296.
- Greco J. (2000), *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge: Cambridge University Press.
- Greco J. (2009), "Knowledge and Success from Ability", *Philosophical Studies* 142, 17-26.
- Kaplan M. (1985), "It's not What You Know that Counts", *Journal of Philosophy* 85(7), 350-363.
- Lackey J. (2008), *Learning from Words: Testimony as a Source of Knowledge*, New York: Oxford University Press.
- List Ch., Pettit P. (2011), *Group Agency: The Possibility, Design and Status of Corporate Agents*, Oxford: Oxford University Press.
- Kvanvig J. (2003), *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Palermos S. O., Pritchard D. (2013), "Extended Knowledge and Social Epistemology", *Social Epistemology Review and Reply Collective* 2(8), 105-120.

- Palermos S. O. (2014), "Loops, Constitution and Cognitive Extension", *Cognitive Systems Research* 27, 25-41.
- Palermos S. O. (2016), *The Distribution of Epistemic Agency* [in:] *Social Epistemology and Epistemic Agency: De-Centralizing Epistemic Agency*, P. J. Reider (ed.), London: Rowman & Littlefield International.
- Palermos S. O. (2016), "The Dynamics of Group Cognition", *Minds and Machines* 26(4), 409-440.
- Palermos S. O., Tollefsen D. (forthcoming), *Group Know-How* [in:] *Socially Extended Epistemology*, J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos, D. Pritchard (eds.), Oxford: Oxford University Press.
- Pettit P. (2007), "Rationality, Reasoning and Group Agency", *Dialectica* 61, 495-519.
- Plato (2002), *Meno* [in:] *Five Dialogues: Euthyphro, Apology, Crito, Meno, Phaedo*, tr. G. M. A. Grube, Indianapolis, IN: Hackett Publishing Company, 58-92.
- Plato (2004), *Theaetetus*, tr. J. Sachs, Newburyport, MA: Focus Publishing.
- Pritchard D. (2007), "Recent Work on Epistemic Value", *American Philosophical Quarterly* 44(2), 85-110
- Pritchard D. (2009), "Knowledge, Understanding and Epistemic Value", *Royal Institute of Philosophy Supplement* 64(64), 19-43.
- Pritchard D. (2010), "Cognitive Ability and the Extended Cognition Thesis", *Synthese* 175(1), 133-151.
- Pritchard D., Turri J. (2014), "The Value of Knowledge" [in:] *The Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.), <https://goo.gl/3SA2uk>.
- Rupert R. (2009), *Cognitive Systems and The Extended Mind*, New York: Oxford University Press.
- Ryle G. (1949), *The Concept of Mind*, Chicago: Chicago University Press.
- Schlosser M. (1915), "Agency" [in:] *The Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.), <https://goo.gl/pD6edu>.
- Snowdon P. (2004), "Knowing How and Knowing That: A Distinction Reconsidered", *Proceedings of The Aristotelian Society* 104, 1-29.
- Sosa E. (1988), "Beyond Skepticism, to the Best of our Knowledge", *Mind* 97(386), 153-189.
- Sosa E. (1991), *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- Sosa E. (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- Stanley J., Williamson T. (2001), "Knowing How", *Journal of Philosophy* 98(8), 411-444.
- Theiner G., O'Connor T. (2010), *The Emergence of Group Cognition* [in:] *Emergence in Science and Philosophy*, A. Corradini, T. O'Connor (eds.), London: Routledge, 6-78.
- Tollefsen D. (2004), "Collective Epistemic Agency", *Southwest Philosophy Review* 20, 55-66.
- Trybulec B. (2017a), "Group as a Distributed Subject of Knowledge – between Radicalism and Triviality", *Avant*, forthcoming.
- Trybulec B. (2017b), *Wiedza i jej podmiot w szerokich systemach poznawczych*, Warszawa: Wydawnictwo IFiS PAN.
- Tuomela R. (1992), "Group Beliefs", *Synthese* 91(3), 285-318.
- Tuomela R. (2011), *An Account of Group Knowledge* [in:] *Collective Epistemology*, H. Bernhard, D. Sirtes, M. Weber (eds.), Frankfurt am Main: Ontos, 20-75.
- Wallis C. (2008), "Consciousness, Context, and Know-How", *Synthese* 160(1), 123-153.

- Zagzebski L. (1996), *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.
- Zagzebski L. (2003), "The Search for the Source of the Epistemic Good", *Metaphilosophy* 34(1-2), 12-28.