



Moral Disagreement and Moral Education: What's the Problem?

Dominik Balg¹ 

Accepted: 24 May 2023
© The Author(s) 2023

Abstract

Although initially plausible, the view that moral education should aim at the transmission of moral knowledge has been subject to severe criticism. In this context, one particularly prominent line of argumentation rests on the empirical observation that moral questions are subject to widespread and robust disagreement. In this paper, I would like to discuss the implications of moral disagreement for the goals of moral education in more detail. I will start by laying out the empirical and philosophical assumptions behind the idea that widespread and robust moral disagreement undermines the prospects of transmitting moral knowledge in educational settings. Having thus provided a specific interpretation of the epistemic dynamics behind this so-called ‘challenge of disagreement’, I will proceed by discussing its didactical implications. More specifically, I will defend two claims: first, I will argue that the challenge of disagreement is not an effective challenge, because it undermines the possibility of knowledge transfer only with respect to a limited set of moral propositions. Second, I will argue that the challenge of disagreement is not a specific challenge, because the epistemically destructive effects of moral disagreement also pose a challenge for other prominent accounts of moral education that were originally proposed as promising alternatives to knowledge transmission accounts. If convincing, my arguments show that knowledge transmission accounts of moral education are in a much better position than is usually expected to incorporate the fact that moral questions are notoriously controversial.

Keywords Moral education · Moral disagreement · Moral knowledge · Moral uncertainty

1 Introduction

What are the goals of moral education? One possible answer to this question is the following: moral education is simply education with respect to a specific topic. Accordingly, the goals of moral education are simply *general* educational goals *that are applied to this specific topic*. Viewed from this perspective, there is nothing special about moral education. For example, let's assume that the transmission of knowledge is one important general

✉ Dominik Balg
dominik.balg@posteo.de

¹ Philosophy Department, University of Mainz, 55099 Mainz, Germany

goal of education. Given this, we could simply apply this goal to a variety of specific educational domains. One important goal of science education would be the transmission of scientific knowledge. One important goal of history education would be the transmission of historical knowledge. And one important goal of moral education would be the transmission of moral knowledge.¹

Although initially intuitive, many if not most authors within the didactics of philosophy and the philosophy of education will likely reject the above idea. One main reason for this reluctance to regard the transmission of moral knowledge as an important goal of moral education is the worry that – given widespread philosophical disagreement about moral questions – such transmission is simply impossible: if all or most moral questions are subject to widespread and robust disagreement, there won't be much transferable moral knowledge to begin with, so that any attempt to directly transmit specific moral beliefs would simply amount to indoctrination (Hand 2018, 1).

In this paper, I would like to discuss this *challenge of disagreement* in more detail. More specifically, I will defend two claims. First, I will argue that the challenge of disagreement is not an *effective* challenge, because it undermines the possibility of knowledge transfer only with respect to a limited set of moral propositions. Second, I will argue that the challenge of disagreement is not a *specific* challenge for knowledge transmission accounts of moral education, because the epistemically destructive effects of moral disagreement also pose a challenge for other prominent accounts of moral education. I will proceed as follows. In Section 2, I will start by laying out the empirical and philosophical assumptions behind the challenge of disagreement, conceived as a specific epistemic challenge for moral knowledge transmission, in more detail. Having thus provided a specific interpretation of the epistemic dynamics behind the challenge of disagreement, I will proceed by discussing its didactical implications in more detail. In Section 3, I will critically discuss the *efficacy* of the challenge of disagreement, arguing that it doesn't undermine the possibility of moral knowledge transmission in general, but only with respect to specific first-order moral propositions that are sufficiently controversial among professional ethicists. In Section 4, I will discuss the *accuracy* of the challenge of disagreement, arguing that it also undermines other dimensions of moral education that were originally proposed as promising alternatives to moral knowledge transmission, such as teaching moral reasoning skills or enabling students to participate in moral discourse. Section 5 will sum up the main results.

If convincing, my arguments show that knowledge transmission accounts of moral education are not just better than widely assumed, but actually do better than their main competitors in dealing with the educational implications of moral disagreement. However, I want to stay explicitly neutral on the question of whether knowledge transmission accounts are superior *overall* to alternative accounts of moral education – or whether they are even feasible at all. Besides the problem of disagreement, there are a number of further challenges that knowledge transmission accounts of moral education must face. Most importantly, these accounts seem to rest on a number of controversial epistemological assumptions – e.g. assumptions about the existence of moral experts or the possibility of moral testimony. While I am personally optimistic that many of the various objections lurking in this context can be successfully met, any detailed discussion of them would clearly exceed the scope of this paper. What's more, even if we assume that knowledge transmission accounts of moral education are *epistemically* defensible, there will still be political and

¹ In this paper, I will use the terms ‚aim‘ and ‚goal‘ interchangeably.

pedagogical challenges to the idea that moral knowledge should be directly transmitted in educational contexts (for an overview, see e.g. Giesinger 2021). Accordingly, all I want to argue for in the following is the comparatively narrow claim that the specific phenomenon of moral disagreement is not so much a problem for knowledge transmission accounts as it is for competing accounts of moral education. While clearly being comparatively narrow, this claim still has bite: in fact, the challenge of disagreement is often seen as one of the most pressing problems for moral education in the relevant literature (Hand 2018, p. 1). Given that, a satisfactory solution of this specific challenge would already lead to a significant shift in the dialectical landscape.

Before I proceed, however, and in order to avoid misunderstandings, I would like to make an important conceptual clarification. Especially for experienced educators, the idea that „knowledge transmission “ should be a goal of any kind of education might already raise some concerns. To better understand this idea, it is helpful to distinguish between an epistemological, a psychological and a didactical reading of “knowledge transmission”. According to the epistemological reading, this notion only refers to a specific *epistemic* relation between the input beliefs and the output beliefs of educational transfer processes. According to the psychological reading, the transmission of knowledge also includes the transmission of *psychological* features that concern the cognitive embeddedness of beliefs. According to the didactical reading, “transmission of knowledge “ refers to top-down educational settings where instructors simply impart knowledge to their students instead of enabling them to acquire the relevant insights on their own. While the educational importance of knowledge transmission in the epistemological sense has been widely acknowledged (see e.g. Goldman 2006, p.11)—in fact, it is epistemologically uncontroversial that most of what we know we only know on the basis of other people’s knowledge (see e.g. Goldberg 2010)—knowledge transmission in the psychological and didactical sense is usually judged to be pedagogically illegitimate or even impossible.² In this paper, whenever I speak of „knowledge transmission “, I speak of knowledge transmission in the epistemological sense. Aiming at the transmission of knowledge in this sense in educational contexts is both compatible with psychological findings to the effect that the acquisition of knowledge is rather a result of individual construction than of passive perception and with the didactical demand that students should acquire insights on their own (Hand 2018, 38).

2 The Epistemic Dynamics behind the Challenge of Disagreement

At first glance, the idea that the pervasiveness of moral disagreement directly speaks against the possibility of educational transfers of moral knowledge seems pretty straightforward. However, upon closer inspection, it is at least not obvious *why* and *how* the controversiality of moral issues is supposed to undermine educational transfer processes. After all, one might argue, people disagree about almost everything – and still, at least in domains like mathematics, science, history or geography, educational transfer of substantive, domain-specific insights seems to be perfectly legitimate. So why should the mere fact that moral questions are notoriously controversial pose a specific problem for the educational transmission of moral knowledge?

² This is one of the main motivations behind so-called constructivist educational environments, see e.g. O’Donnell 2012.

To answer this question, I would like to provide a deeper analysis of the challenge of disagreement in this section. In order to get a better grasp of the philosophical details behind this challenge, it is helpful to cite one of its proponents in a little more detail. So, to begin with, consider the following passage from Kirsten Meyer (Meyer 2011, p. 229 f., my translation):

Let's suppose that there are in fact specific moral principles whose correctness we are able to grasp and upon which we can act. Even against the background of this assumption, there remains an important problem for moral education. After all, it is still not clear at all *which* moral principles are correct – and accordingly, which principles should be taught in educational settings. One simple reason why moral education can't be aimed at the transmission of moral principles is that such principles are controversial within academic philosophy. If moral education was about transferring knowledge of correct moral principles, one would first need to know which exact principles these should be. [...] Which principles should be taught in philosophy classes in order to enable students to act morally? Should we advocate for a kantian or a utilitarian position? Do we have to wait for the underlying philosophical disputes to be settled before we can answer this question? [...] When we try to identify the proper goals of moral education [...], we encounter [...] several open questions. A satisfying answer to these questions would require a settlement of century-old philosophical disputes. And such a settlement is not to be expected.

Against the background of this passage, the dialectical profile of the challenge of disagreement already becomes clearer. Most importantly, it becomes clear that the challenge of disagreement is a specific objection against knowledge transmission accounts of moral education that already rests on a number of specific assumptions: first of all, the basic idea behind this challenge is not that it is *in principle* impossible to transfer moral knowledge in educational settings – e.g. because there aren't any moral principles, or because such principles aren't truth-apt. Rather, the idea is that in light of widespread moral disagreement, it isn't clear *which* moral principles should be transferred in educational settings. In this way, the challenge of disagreement already presupposes the possibility of moral knowledge. And in fact, this shouldn't come as a surprise: for instance, let's suppose that a fundamentally pessimistic view on the possibility of moral knowledge – e.g. moral nihilism or moral non-cognitivism – was correct. In such a case, it seems that whether moral questions are in fact controversial or not would be completely *irrelevant* to the possibility of moral knowledge transfer, because even if everyone agreed on the same moral verdicts, we still wouldn't be in a position to transfer any moral knowledge in educational settings. The challenge of disagreement is supposed to show that crucial features of moral knowledge transfers are effectively undermined by the fact that moral questions are notoriously controversial. But if there is nothing to be undermined in the first place, this idea seems like a nonstarter.

Furthermore, the challenge of disagreement explicitly presupposes that moral questions are controversial among professional philosophers and that it is exactly this kind of disagreement that undermines the possibility of moral knowledge transfer. In fact, it is important to note that in the above passage, Meyer isn't just talking about any old moral disagreement, but rather specifically refers to moral disagreements *within academic philosophy*, e.g. between kantians and utilitarians. But what's so special about moral disagreements between academic philosophers? The idea seems to be here that – just as educational knowledge transfer in other domains – educational transfer of moral knowledge should be informed by the relevant experts. However, given that even the experts disagree about moral questions, we are in no position to decide which moral principles should be taught

in educational settings. In what follows, I would like to discuss these assumptions behind the challenge of disagreement in a little more detail in order to provide a more fine-grained understanding of the *prima facie* plausible idea that widespread moral disagreement poses a fundamental problem for educational transfer of moral knowledge.

Let 's start with the assumption that it is in principle possible to acquire moral knowledge. As we have seen, this assumption is implicitly presupposed by the challenge of disagreement, because if it was impossible to acquire any moral knowledge in the first place, it would be completely irrelevant whether moral questions are in fact controversial or not. Given how extensively pessimistic views on the possibility of moral knowledge have been discussed in the philosophical literature, this already seems like a substantive and potentially problematic metaethical commitment. However, there are several things to be said to soften the worry that the challenge of disagreement can simply be rejected by appealing to some form of metaethical pessimism. First of all, the assumption that we are in principle in a position to acquire moral knowledge doesn 't just underlie the challenge of disagreement, but also much of our everyday moral practice: at least from a pretheoretical perspective, it seems overwhelmingly plausible that we *can* and actually *do* have a lot of moral knowledge. For example, most people would agree that I *know* that torturing puppies out of boredom is morally wrong, or that I 'm not morally obliged to compliment random strangers on their outfits. Given this, the fact that the challenge of disagreement questions the legitimacy of moral knowledge transfer on the basis of critical considerations concerning the *scope*—and not the mere *possibility*—of moral knowledge seems to be a strength and not a weakness.

Furthermore, and more importantly, the assumption that it is in principle possible to acquire moral knowledge is obviously already presupposed by the very target of the challenge of disagreement – i.e. by knowledge transmission accounts of moral education. Proponents of these accounts cannot defend themselves against the challenge of disagreement on the basis of pessimistic views on the possibility of moral knowledge without undermining their own position. Given this, any critique of the challenge of disagreement that points to the optimistic presuppositions that this challenge implies with respect to the possibility of moral knowledge seems to be at least dialectically ineffective.

Notably, the same is true with respect to the second assumption behind the challenge of disagreement that has been identified above – the assumption that moral questions are controversial among professional philosophers and that it is exactly this kind of disagreement that undermines the possibility of moral knowledge transfer. One rather obvious problem with this assumption is that it seems to presuppose that academic philosophers enjoy some kind of *moral expertise* and are thus in a privileged position to inform educational transfer processes. This presupposition will probably strike many people as hopelessly elitist (Driver 2006) – but again, it is plausibly already presupposed by knowledge transmission accounts of moral education.³ For as already illustrated at the outset of this paper, the very idea behind these accounts seems to be that there is nothing special about moral education: just as in other domains, educational processes in the moral domain should be aimed at the transmission of domain-specific knowledge that is provided by the relevant academic

³ In a way, the identified assumptions concerning the possibility of moral knowledge and moral expertise can also be seen as *concessions* to knowledge transmission accounts of moral education. However, I explicitly want to stay neutral with respect to the question of whether actual proponents of the challenge of disagreement accept these assumptions (i) because they believe them to be true or (ii) because they only want to make specific concessions to knowledge transmission accounts of moral education in order to argue that these accounts ultimately fail *even against the background of these concessions*.

community. Given this, proponents of these accounts cannot defend themselves against the challenge of disagreement by rejecting its presuppositions concerning the possibility of moral expertise without undermining their own position.

So let's suppose that professional philosophers can plausibly be considered as moral experts. Is it true that moral questions are controversial among professional philosophers? And how exactly are these disagreements supposed to undermine the possibility of educational knowledge transfer? The first of these questions is clearly an empirical one. However, it is surprisingly hard to answer. One obvious difficulty is to identify the relevant group of people in the first place – what does it even mean to be a philosopher? And is every philosopher an expert on any philosophical question?⁴ Furthermore, many widespread views on the amount of controversy within academic philosophy will likely be distorted by naive misconceptions and prejudices which portray philosophy as a messy and hopelessly unoriented endeavor.

Despite all these difficulties and biases, there has been one notable attempt to get a clearer picture of the amount and distribution of disagreement among professional philosophers that has received a lot of attention within the philosophical community. In their much-discussed PhilPapers Surveys, David Chalmers and David Bourget regularly survey professional philosophers in order to help uncover their views on key philosophical questions. The results of this project allow for a more justified estimation of the degree of moral disagreement among experts. For example, considering only the answers of respondents who are regular faculty members in BA-granting philosophy departments, and who named “normative ethics” as their area of specialization, the 2020 PhilPapers Survey revealed a fundamental disagreement about which ethical theory is correct (for the following numbers, see Bourget and Chalmers 2021). While 27.27% accept or lean toward deontology, 20.45% accept or lean toward consequentialism, 17.42% accept or lean toward virtue ethics and 0.76% are agnostic or undecided. In fact, it looks like the degree of controversy has even increased over the last few years. For, in the preceding PhilPapers Survey of 2009, 35.25% accepted or leaned toward deontology, 23.02% accepted or leaned toward consequentialism, and 12.23% accepted or leaned toward virtue ethics (Bourget and Chalmers 2014). While it is also important to keep in mind that (i) conflicting ethical theories don't necessarily yield conflicting verdicts with respect to all or even most concrete moral questions and (ii) proponents of the same ethical theory regularly end up with conflicting applications of this theory, these numbers clearly justify the assumption that a lot of concrete moral issues will be highly controversial among the relevant experts.

Why is such a high degree of controversy among experts a serious problem for moral knowledge transmission? In what follows, I would like to provide a specific interpretation of the epistemic dynamics behind the idea that widespread expert disagreement undermines the possibility of educational knowledge transfer. This interpretation is directly informed by the recent debate about the epistemic implications of disagreement within

⁴ In fact, it seems that when looking for moral expertise, one shouldn't exclusively focus on professional philosophers. If we assume that moral expertise is ultimately grounded in specific argumentative skills and knowledge of ethical theories (see e.g. Singer 1982), being a professional philosopher – or even a professional ethicist – is neither necessary nor sufficient to count as a moral expert. Given this, my focus on professional philosophers in this paper is merely based on pragmatic considerations: First of all, if moral expertise is closely connected to specific argumentative skills and knowledge of ethical theories, being a professional philosopher will at least be a strong indication of moral expertise. Second, the distribution of moral views among professional philosophers has been subject to empirical research, which allows for a somewhat reliable estimation of the amount of moral disagreement among professional philosophers.

social epistemology. One important insight from this debate is that the epistemic implications of a situation of disagreement crucially depend on the respective levels of competence of the disagreeing parties (see e.g. Kelly 2005, p. 168). In this context, one specific constellation of cases that have received a lot of attention within the epistemological literature are disagreement situations where the disagreeing parties are *epistemic peers*, which means that the disagreeing parties have access to the same (or equally good) evidence and are equally competent in assessing this evidence (Christensen 2009, pp. 756–757).⁵ One reason why such cases have received so much attention within the epistemological literature is that they apparently have profound epistemic implications. More specifically, peer disagreements seem to make it *epistemically irrational* to hold on to one's beliefs. For, whenever a person finds herself in disagreement with someone she considers her epistemic peer, she is rationally required to revise her original position. To illustrate this point, many authors refer to the so-called Restaurant Case which was originally developed by Christensen (2007). The case goes as follows (ibid., p. 193):

Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20% tip, and we further agree to split the whole cost evenly [...]. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each. How should I react, upon learning of her belief? [...] Let us suppose that my friend and I have a long history of eating out together and dividing the check in our heads, and that we've been equally successful in our arithmetic efforts: the vast majority of times, we agree; but when we disagree, she's right as often as I am. So for the sort of epistemic endeavor under consideration, we are clearly peers. Suppose further that there is no special reason to think one of us particularly dull or sharp this evening—neither is especially tired or energetic, and neither has had significantly more wine or coffee.

In this case, many authors agree that the protagonist of the case is rationally required to give up her belief that the share is \$43 each. Although there is some considerable disagreement about the relevant epistemic mechanics behind the rationally required response in such a case, and its exact doxastic nature, there is wide agreement that the described situation requires a *state of substantive uncertainty* (for a discussion, see e.g. Christensen 2007; Grundmann 2019; Kelly 2011). While this verdict and its generalization over all cases of peer disagreement can – and actually has been (see e.g. Kelly 2005, 2011) – disputed, I will assume a conciliationist view in what follows, according to which situations of recognized

⁵ It is important to note that other authors have proposed alternative definitions of epistemic peerhood. For example, Thomas Kelly has suggested that two persons are epistemic peers when they are (i) equally familiar with the relevant evidence and (ii) equals with respect to general intellectual virtues (Kelly 2005, 174). Alternatively, it has been suggested that two persons are epistemic peers when they are equally likely to be mistaken (Elga 2007, 499; Wedgwood 2010, 236). However, given that these definitions are more liberal than Christensen's conception of epistemic peerhood, the following considerations are neutral with respect to this controversy.

peer disagreement undermine the epistemic status of the conflicting beliefs.⁶ If we assume a conciliationist view on the significance of disagreement and combine this view with the above assumption about the degree of controversy among professional ethicists, it is easy to see how moral disagreements directly lead to a problem for the possibility of moral knowledge transfer. In situations where the relevant experts agree, disagreement with laypeople doesn't pose a problem for such transfers, because it doesn't undermine the experts' knowledge, which is the ultimate source of the transfer processes. However, this is not the case in situations where the relevant experts disagree *with each other*.⁷ For, if even the experts are (or should be) in a state of substantive uncertainty about the questions at issue, they are in no position to inform educational transfer processes.⁸

⁶ While I will assume a conciliationist view on the significance of disagreement, I would explicitly like to stay neutral with respect to any deeper explanation of how exactly recognized disagreements undermine the epistemic status of our beliefs. One widespread view in this context is that disagreement situations provide defeaters for our beliefs – however, the exact nature of these defeaters is controversially discussed (see e.g. Constantin 2021, ch. 6). In any case, it seems plausible that disagreement situations have epistemically destructive effects because they confront us with our epistemic limitations – or as David Christensen puts it (Christensen 2007, 187 f.): ‘*We all live out our lives in states of epistemic imperfection. [...] One fairly common situation that may present opportunities for improvement is that of discovering that another person's belief on a given topic differs markedly from one's own.*’ In this respect, conciliationism might be seen as a non-ideal epistemological theory: It tells us how to adequately deal with our epistemic limitations.

⁷ For a detailed discussion of the epistemic implications of disagreements with epistemic *inferiors*, see Priest 2016.

While the conditions for epistemic peerhood are sometimes criticized as being unrealistically strict, it seems that these conditions are plausibly often met in academic contexts. Just consider the striking degree of epistemic homogeneity between professional philosophers – they usually have equal academic credentials and a comparable level of competence. Furthermore, they are equally familiar with the relevant literature and acquainted with the same set of arguments. Given this, the kind of disagreement situations that professional ethicists find themselves in seem sufficiently similar to artificial situations like the Restaurant Case.

⁸ One objection that immediately suggests itself at this point is that the above considerations fail to appreciate the uniqueness of moral disagreements. To my mind, this objection seems ad hoc – why should moral disagreements behave epistemically different than other kinds of disagreement? To establish such a fundamental difference, one would first need to identify some special feature of moral disagreements that could properly explain their alleged epistemic innocence. For example, it has been argued that some of our beliefs – including our political, moral and religious beliefs – are at least partially based on private evidence (see e.g. Feldman 2007; Rosen 2001; van Inwagen 1996). However, even if that were the case, conflicting beliefs of other people that are based on private evidence would still have epistemically destructive effects, because the mere fact that there is evidence that supports a conflicting belief already speaks against the epistemic quality of one's own belief, even if that evidence is inaccessible (Feldman 2007). According to another suggestion, the epistemic peculiarity of moral disagreements lies in the fact that these disagreements are grounded in differing, but equally reasonable evaluations of the available evidence (for a critical discussion of this claim, see e.g. Ballantyne and Coffman 2011; Brueckner and Bundy 2012; Schoenfield 2014). However, even if we accepted this suggestion, it would at best show that conflicting moral beliefs of other people don't speak against the reasonableness of our own beliefs. But as long as we assume that only one side of a moral disagreement can be correct, they would still speak against the truth of our own beliefs and are thus epistemically destructive (Balg 2021a, pp. 73 ff.). According to the last suggestion that I would like to discuss here, the epistemic peculiarity of moral disagreements lies in the fact that the truth values of moral beliefs are relative to personal standards, so that conflicting moral beliefs of different persons can be true at the same time (see e.g. Kölbel 2004). While such a relativistic view would indeed speak against the epistemic toxicity of diverging moral beliefs, it is also highly controversial and likely to be rejected by the majority of professional philosophers. One reason for this reluctance is the worry that moral relativism doesn't even allow for the bare existence of moral disagreement (see e.g. Lasersohn 2005; Wright 2001). So it seems that establishing a fundamental epistemic difference between moral and non-moral disagreements is no easy task. Although there are some features that have been suggested as possible epistemic peculiarities of moral disagreements in the literature, these features would in no way suffice to show that there are genuine moral disagreements with epistemic peers that are still epistemically innocuous.

I take the epistemic dynamics that has been delineated above to be the most plausible interpretation of the assumption that widespread expert disagreement undermines the possibility of moral knowledge transfer in educational settings. This interpretation doesn't only neatly explain how the controversiality of moral questions might undermine the possibility of such knowledge transfers, but also sheds some light on the significance of *agreement* for moral education. Given the epistemic dynamics analyzed above, it gets clear that the absence of (sufficiently robust) disagreement among moral experts is a necessary condition for moral knowledge transfer. However, that doesn't necessarily mean that (sufficiently robust) agreement among moral experts is necessary for knowledge transfer. To see this, we just have to consider situations in which widespread agreement among experts is absent because only a very small fraction of the expert community has formed any beliefs on a given question, because this question hasn't even been considered by most of their colleagues. If we assume that the beliefs of these few experts constitute knowledge, this status will in no way be undermined by the mere fact that other experts don't agree. Furthermore, moral agreement is also not sufficient for moral knowledge transfer to be legitimate or even possible. For example, if we have good reasons to believe that a specific group of experts is fundamentally corrupt or misguided, the mere fact that they agree with each other should be of no relevance. In this way, the above interpretation doesn't just help to better understand the challenge of disagreement, but also to locate its dialectical profile more accurately.

As we have seen in this section, the challenge of disagreement is a specific challenge for knowledge transmission accounts of moral education that rests on a number of specific empirical and philosophical assumptions. Some of these assumptions are more controversial than others, and some also have to be presupposed by proponents of knowledge transmission accounts. In what follows, I would not like to discuss these assumptions in more detail. Instead, I will simply accept them in order to make the challenge of disagreement as strong as possible. For as we will see, even if we accept all the assumptions that underlie the challenge of disagreement, it will still fail to pose a devastating problem for knowledge transmission accounts of moral education.

3 The Efficacy of the Challenge of Disagreement

In the last section, we saw why moral disagreement among professional ethicists poses a fundamental challenge to the possibility of moral knowledge transfer. That is, if professional ethicists are rationally required to adopt a stance of uncertainty with respect to moral questions, then they are in no position to present their moral views as knowledge to laypeople. In this section, I would like to critically discuss the *efficacy* of this challenge in more detail. Obviously, moral disagreement among professional ethicists has profound implications for the prospects of moral knowledge transfer. Yet, at the same time, it seems that it doesn't undermine *all* transfer of moral knowledge, but only *with respect to certain propositions*. As we saw in the last section, the challenge of disagreement as it stands ultimately rests on the observation that professional ethicists disagree about which moral theory is correct (Meyer 2011, pp. 229–230). Given this, professional ethicists will also disagree about any first-order moral question to which their competing moral theories imply conflicting answers, which in turn will undermine any possibility of knowledge transfer with respect to these questions. However, as I will argue in this section, there will also be a lot

of moral questions which are subject to substantive agreement – and with respect to these questions, the challenge of disagreement remains ineffective.

At first glance, to someone who is familiar with the relevant literature, it might seem that this is something that has already been done. For example, in his paper ‘Towards a Theory of Moral Education’, Michael Hand, a prominent proponent of knowledge transmission accounts of moral education, writes (Hand 2014, p. 528):

[...] Disagreement about the justificatory status of many moral standards [...] looks set to be a salient feature of our moral landscape for the foreseeable future. But it would [be] premature to infer from this state of affairs that robustly justified moral standards are unavailable. From the fact that some moral standards have an uncertain justificatory status, it does not follow that all do; and from the fact that some arguments for moral subscription are dubious, it does not follow that all are. Perhaps, somewhere in the melee of controversial moral standards and arguments, there are at least some standards on which all are agreed and to which subscription is demonstrably justified.

This, I think, is just how things are. Within the reasonable plurality of moral standards there is an identifiable subset to which more or less everyone subscribes and for which the reasons to subscribe are compelling. There is a very broad consensus in society on some basic moral prohibitions (on stealing, cheating, causing harm, etc.) and prescriptions (to treat others fairly, help those in need, keep one’s promises, etc.). And there is a familiar rational justification for those basic moral standards whose cogency is hard to dispute. It is, briefly, that moral standards are justified when their currency in society serves to ameliorate the ever-present risk in human social groups of breakdowns in cooperation and outbreaks of conflict.

To my mind, there are several problems with the above passage. First of all, it seems that Hand is talking about the wrong kind of justification. In order to constitute knowledge – and thus in order to qualify as eligible input for educational transfer processes – moral beliefs need to enjoy a sufficient degree of *epistemic* justification. In contrast, Hand seems to talk about *pragmatic* justification: his idea is that we have pragmatic reasons to accept certain standards because they help us to sustain social cooperation – and not that we have epistemic reasons that speak for the truth of these standards.⁹

But even putting this worry to the side, it still seems that Hand is also looking for the wrong kind of agreement: just as in other domains, educational knowledge transfer in the moral domain should be, given that it is feasible at all, informed by the relevant experts – and not by the general public. Given this, the kind of agreement that is necessary for educational knowledge transfer to be epistemically legitimate is agreement among the relevant

⁹ In fact, in some of his other work, Hand seems to explicitly endorse a radical pluralism with respect to the kind of justification that moral standards can enjoy in order to qualify as eligible input for educational transfer processes. For example, in his book “A Theory of Moral Education”, Hand writes (Hand 2018, p. 69):

[Recognising] that the problem-of-sociality justification is sound is quite compatible with believing that some other putative justification for subscribing to moral standards is also sound. It is no part of the foregoing argument that this is the only kind of good reason we could have for [...] subscription to standards. It may turn out to be the case that basic moral standards are justified both by their capacity to avert conflict and sustain cooperation in human social groups and, say, by their divine authorization, or by an imperative of pure practical reason [...]. Different justifications for the same standards can sit quite happily alongside each other; acceptance of one does not necessitate rejection of all others.

experts – because as we have seen, it is exactly this kind of agreement whose absence potentially undermines the epistemic status of expert beliefs.

To be clear, this should obviously be a rather significant distinction. It is easily conceivable that there are moral questions which are almost completely uncontroversial among the general public, but highly controversial among moral experts – and vice versa. Furthermore, there might also be moral questions which are sufficiently uncontroversial among the relevant experts, but which haven't even been considered by the general public. So at least against the background of the specific interpretation of the challenge of disagreement that has been developed in the last section, it seems that the question that needs to be answered in order to assess the efficacy of this challenge is not whether there are any moral principles that are sufficiently uncontroversial among the general public to be pragmatically justified, but whether there are any moral principles that are sufficiently uncontroversial among the relevant experts to be epistemically justified. In what follows, I would like to make some steps towards answering this question.

3.1 (Meta-)Ethical Agreement

One first obvious source of moral expert agreement are (i) first-order moral questions with regard to which the competing moral theories yield the same result and (ii) meta-ethical questions. With respect to these questions, the possibility of moral knowledge transmission seems to be untouched by the challenge of disagreement. And in fact, the 2020 PhilPapers Survey seems to support this assumption. For example, of the 227 meta-ethicists who participated in this study, 77.53% accept or lean toward cognitivism, and 65.35% accept or lean toward moral realism. With respect to the latter issue, the degree of controversy has indeed significantly decreased in comparison to 2009 (Bourget and Chalmers 2021). Given the prevalence of student statements like “When it comes to moral questions, there is no true or false” and “There are no moral facts”, it seems that these results will be of some didactical significance.¹⁰ Furthermore, there also seems to be some striking agreement with respect to first-order moral questions. For example, 86.11% of professional ethicists¹¹ agree that abortion is generally permissible, 74.13%¹² agree that capital punishment is impermissible, and 74.26% agree that human genetic engineering is permissible.¹³ And again, given how controversial these issues usually are among students, this result should be of some didactical significance. To be fair, the above results in no way suffice to establish any significant number of moral insights that could then be readily taken as the proper input of knowledge transfer processes. One first problem is that the available empirical evidence is still pretty meager: what we would need is a substantial body of empirical data that enables us to get a clearer picture of the scope and degree of moral agreement among the relevant experts. A second problem is rather philosophical in nature. That is, even if we knew exactly what moral philosophers agree on, there would still be the question of how much agreement is needed to allow for the possibility of knowledge transmission. Or, to

¹⁰ Non-cognitivist and anti-realist statements like the above can be seen as typical expressions of a well-known phenomenon that is discussed in the didactical literature under the label “student relativism” (for an overview, see Pfister 2019).

¹¹ $N = 144$.

¹² $N = 143$.

¹³ $N = 136$. All these numbers are based on the answers of all respondents who named “applied ethics” as their area of specialization.

put it differently: what degree of expert disagreement suffices to undermine the possibility of knowledge transmission? While any answer to this question depends on a variety of complicated issues, epistemologists do agree that numbers matter (Grundmann 2013).¹⁴ In a situation where only a small group of experts deviate from the majority view, this level of disagreement won't suffice to undermine the possibility of knowledge transmission. At the same time, it is not easy to determine any specific threshold above which disagreement starts to have sufficiently destructive epistemic implications. Solving this philosophical problem would also be of great help for estimating the exact scope of the challenge of disagreement for moral knowledge transfer. Nevertheless, the above considerations do show that it would be premature to simply dismiss any attempt to legitimize moral knowledge transfer in educational settings by vaguely pointing to the alleged pervasiveness of moral disagreement. While professional ethicists certainly disagree about many moral questions, there will also be a considerable amount of moral agreement.

3.2 Agreement on Second-Order Principles

What I would like to argue for in the remainder of this section is that even in cases where professional ethicists disagree, the transmission of moral knowledge is still possible. To see how, let's start with a simple thought experiment developed by Matheson (2021, p. 15):

Trina is travelling for work. When she travels for work, her work covers her travel costs upon receiving the receipts. On this trip, Trina has a particularly bad travel experience. Her travel experience is so bad that the airlines refunds the price of her return trip. Upon returning home, Trina is thinking about whether she should still submit the original receipt to her work since the refund was intended as compensation for a bad experience. Trina talks things over with her good friend Lesley. Trina and Lesley disagree about what is morally permissible here even though they agree about all the non-moral facts relevant to the issue.

What should Trina do in this scenario? Intuitively, it seems that under the described circumstances, Trina shouldn't submit the receipt to her work. Furthermore, it seems that

¹⁴ One reason why this question can't be answered on a general level is that the epistemic significance of expert agreement depends (i) on the exact number of the agreeing experts and (ii) on the degree to which these experts came to their respective verdicts independently: according to Condorcet's jury theorem, an increasing group of people who (i) have a individual reliability that is higher than 50% and (ii) independently came to the same conclusion approach a cumulative reliability of 100%. In light of this, it gets clear why identifying the point at which disagreement becomes epistemically irrelevant is not only complicated, but needs to be done on a case-by-case basis.

Furthermore, it is important at this point to carefully distinguish between different epistemic contexts. While it seems epistemically legitimate to ignore sufficiently small groups of disagreeing experts *in educational contexts*, it would clearly be illegitimate to do so *in academic contexts* – after all, a small chance will always remain for the majority view to turn out false (for a detailed discussion of such a distinction, see e.g. Balg 2021b, 64 ff.). For example, it seems completely unproblematic to educate children about the risks and dangers of global warming while at the same time ignoring critical views that contradict the scientific consensus of climate change. However, it would also be misguided to prematurely dismiss such critical voices within academic research – after all, they could surprisingly turn out to be true. One important implication of this view is that educational programs that are explicitly directed at the transmission of knowledge will sometimes end up transferring false beliefs. At the same time, this needn't be a problematic result – for example, it is not surprising at all that students learn different things in science classes today than they did 200 years ago.

Trina is also in a position to *know* that she shouldn't do this and that she would therefore be blameworthy if she decided to submit the receipt anyway. At the same time, Trina *doesn't know* whether submitting the receipt would be morally permissible. To describe this peculiar case in a slightly improper way, we could say that although Trina doesn't know what to do, she knows what to do *given that she doesn't know what to do*. How can we put this more precisely? At this point, it is helpful to distinguish between different levels of decision-guiding norms. With respect to the relevant first-order norms, Trina is in a state of substantial uncertainty. Given her disagreement with her friend Lesley, she doesn't know whether submitting the receipt would be a morally objectionable form of fraud, which would *directly* give her a first-order reason to not submit the receipt. However, it still seems that Trina is in a position to know perfectly well what she should do. In light of this, it looks like there is a different set of norms – second-order norms – that guide decisions under moral uncertainty. And in the above case, it is easily conceivable that Trina knows about the second-order norms guiding her decision while not knowing about the relevant first-order norms.

More specifically, the second-order moral norm that Matheson identifies on the basis of the above case is the following (see Matheson

MORAL CAUTION (MC): Having considered the moral status of doing action A in context C, if (i) subject S (epistemically) should believe or suspend judgment that doing A in C is a serious moral wrong, while (ii) S knows that refraining from doing A in C is not morally wrong, then S (morally) should not do A in C.

One immediate objection at this point is that this norm doesn't tell us how to *deal with*, but only how to *avoid* states of substantive moral uncertainty. To counter this objection, it will help to look at cases of *complete moral uncertainty*, i.e. cases where the moral status of every available option is unclear. Consider, then, the following case from MacAskill et al. (2020, p.16):

Susan is a doctor, who faces two sick individuals, Anne and Charlotte. Anne is a human patient, whereas Charlotte is a chimpanzee. They both suffer from the same condition and are about to die. Susan has a vial of a drug that can help. If she administers all of the drug to Anne, Anne will survive but with disability, at half the level of welfare she'd have if healthy.

Now consider the following: if Susan decides instead to give all of the drug to Charlotte, then Charlotte will survive with a slight disability at three quarters of the welfare that she'd have if she were healthy. Alternatively, if Susan splits the drug between the two, then they will both survive at a little less than 50% of the welfare that they'd have if they were healthy. Given this specification, we can now consider the rest of the case (MacAskill et al. 2020, p. 17):

Susan is certain that the way to aggregate welfare is simply to sum it up, but is unsure about the value of the welfare of non-human animals. She thinks it is equally likely that chimpanzees' welfare has no moral value and that chimpanzees' welfare has the same moral value as human welfare. As she must act now, there is no way that she can improve her epistemic state with respect to the relative value of humans and chimpanzees.

Her three options, then, are as follows:

A: Give all of the drug to Anne.

B: Split the drug.

C: Give all of the drug to Charlotte.

Susan's decision situation can be represented in the following table:

	Anne's welfare	Charlotte's welfare
Option A	50%	0%
Option B	45%	45%
Option C	0%	75%

The crucial difference between this case and Trina's is that in the above case, *every* option that is relevant to the decision situation is subject to substantive moral uncertainty. More specifically, given that Susan is unsure about the moral value of the welfare of non-human animals, she is not in a position to determine the relative choiceworthiness of any of the available options. To see this, suppose that chimpanzees' welfare has no moral value at all. Against the background of this assumption, A is the best option and C is the worst option. However, if we suppose that chimpanzees' welfare has the same moral value as human welfare, B will be the best option and A will be the worst option. So given that Susan is unsure about the moral value of chimpanzees' welfare, she is in a state of substantive moral uncertainty with respect to each of the options available to her.

Let's call cases like the above cases of *complete* moral uncertainty. The important point in our present context is that even in cases of complete moral uncertainty, there seem to be specific second-order norms that guide our decisions. For example, in the above case, it seems intuitive that it would be morally reckless for Susan not to choose option B. For given her uncertainty about the moral status of non-human animals, she would risk severe wrongdoing by choosing either option A or option C. However, Susan doesn't know whether B *really is* the best option – if it turned out that chimpanzees' welfare didn't have any moral value, then A would be the best option. So if we suppose that Susan knows that she should choose option B, then her knowledge wouldn't consist in first-order knowledge about what the morally best option is, but rather in second-order knowledge about what she should do given that she doesn't know what the morally best option is – i.e. knowledge about second-order norms that guide our decisions under moral uncertainty.

While these first-order and second-order norms are epistemically independent in the way just described, they are still both *moral* norms. One just has to consider the kind of blame that the protagonists in our cases would deserve if they violated the second-order norms guiding their decisions: in such a case, we would say that they have made a *moral* mistake and that they accordingly deserve *moral* blame. Furthermore, while the specific hypothetical scenarios underlying the above cases are comparatively artificial, they also seem to have obvious implications for more realistic situations. For example, take Susan's uncertainty with respect to chimpanzee welfare: just as Susan, many people are uncertain about the moral status that they should assign to non-human animals. And just as in Susan's case, this uncertainty will have direct implications for their moral decisions. To see this, we can just apply the Moral Caution Principle to the situation that many people find themselves in with respect to the question of whether they should continue to consume animal products or adopt a vegan lifestyle. Given the increasing public awareness of the horrific conditions under which many animal products are produced, many of these people will be unsure whether continued consumption of these products would be morally permissible or not. At the same time, many people will plausibly assume that a vegan lifestyle is at least

morally permissible. If we apply the Moral Caution Principle to this situation, it seems that it would directly require veganism. What these considerations show is that second-order moral norms, while usually being developed and discussed against the background of highly artificial and idealized counterfactual scenarios, are directly applicable to real-life contexts. In fact, Matheson himself has applied his principle to the question of whether eating meat is morally permissible and a number of further concrete moral problems like abortion and charitable giving (Matheson).

So as far as the above considerations are convincing, there are substantive moral norms guiding our decisions that are epistemically accessible to us even in cases of moral uncertainty. This result should clearly have significant implications for the prospects of knowledge transmission accounts to adequately deal with the challenge of moral disagreement. The obvious idea at this point is that even in cases where the experts are uncertain about what is morally right, they can still tell us what we should do given that they are uncertain about what is morally right. Or to put it a little more formally: even in cases where the relevant first-order moral norms are subject to substantive expert disagreement, experts can still agree about the corresponding second-order norms that would therefore constitute the proper input of educational transfer processes. Given that these second-order norms, albeit their apparent abstractness, are directly applicable to many real-life situations, this result should be highly relevant to educational practice.

In light of this, there is a direct possibility of transmitting substantive moral insights even with respect to questions that are controversial among professional ethicists. This possibility consists in the transmission of knowledge about second-order moral norms that guide our decisions under moral uncertainty. One obvious objection at this point is that while the above strategy may indeed point to a specific *theoretical* possibility of transferring moral knowledge in the face of moral expert disagreement, it is still doomed to fail, because the second-order norms that it relies upon will again be subject to persistent expert disagreement. But to my mind, such skepticism is at least premature. First of all, it would have to be substantiated by concrete empirical evidence – and given how new the philosophical debate about moral uncertainty is, such evidence won't be readily available. However, simply resorting to pessimistic platitudes about the inevitable controversiality of all philosophical questions won't be enough, since as we have seen, there seems to be a surprising number of philosophical questions with respect to which the relevant experts agree. Given this, a more realistic prediction would be that at least some second-order moral norms will be controversial. If we assume that, then the proposed strategy will be inapplicable in some cases of first-order expert disagreement – namely in those cases in which no fitting second-order principle is available or in which the relevant second-order principles are subject to persistent expert disagreement. Nevertheless, it will still be applicable in some – and perhaps many – cases of first-order expert disagreement, and thereby effectively help to further mitigate the challenge of disagreement for moral knowledge transmission.

4 The Accuracy of the Challenge of Disagreement

In the last section, I have argued that the challenge of disagreement is not an *effective* challenge against knowledge transmission accounts of moral education: although it effectively undermines the possibility of knowledge transmission with respect to first-order moral propositions that are in fact controversial among moral experts, the transmission

of knowledge about sufficiently uncontroversial first-order norms and about second-order norms that guide our decisions under moral uncertainty remains untouched. In this section, I would like to discuss the *accuracy* of the challenge of disagreement: are we really dealing with a *specific* challenge to knowledge transmission accounts of moral education, or does the underlying problem run deeper? In this context, I would like to focus on so-called skill- and virtue-based accounts, which have been particularly influential in the literature. One core idea behind these accounts is that if – given how controversial moral issues are within academic philosophy – students can't just rely on moral experts in gaining moral insights, then they will have to develop these insights on their own. And to do this, they have to be provided with specific skills like critical reasoning skills (Musschenga) and debating skills (Meyer 2011), but also emotional skills (Slote 2009) and general intellectual virtues like open-mindedness or tolerance (Haydon 2003). Michael Hand summarizes this idea as follows (Hand 2018, p. 11):

A [...] standard response to the problem of [...] disagreement is the suggestion that we educate children *about* morality rather than *in* it. On this view we should make children aware of a broad range of moral codes and justificatory arguments, encourage them to subject those codes and arguments to critical scrutiny, and invite them to subscribe to whichever code they take to enjoy the strongest argumentative support. [...] Our job as educators is to cultivate moral autonomy by enabling children to make their own independent judgements on the content and justification of morality.

In this passage, Hand rightly calls these competing accounts explicitly a ‚response‘ to the challenge of disagreement. To appreciate the dialectical significance of this point, it is helpful to elaborate on it in a little more detail. If alternative accounts of moral education are to be read as a response to a specific objection against knowledge transmission accounts, this will already constitute a significant concession. More specifically, if proponents of skill- and virtue-based accounts have developed their theories as a reaction to specific problems of knowledge transmission accounts, this will indicate that they would in principle be willing to accept knowledge transmission accounts – given that those problems can be solved satisfactorily. And in fact, while many proponents of skill- and virtue-based accounts of moral education will plausibly also promote the development of intellectual skills and virtues as a value in itself, at least one important idea behind these accounts seems to be the following: just as educational measures in other domains, moral education would ideally – among other things – also aim at the transmission of substantive domain-specific knowledge. However, this is not how things are. In face of some regrettable epistemic peculiarities of the moral domain, we are in no position to simply pass moral knowledge on to future generations. Given this, we have no other choice than to enable students to make their own, independent moral judgements.

In what follows, I would like to argue that this line of thought is ultimately unconvincing, because skill- and virtue-based accounts of moral education are also heavily affected by the epistemically destructive implications of moral disagreement—and that the challenge of disagreement is therefore not a specific challenge for knowledge transmission accounts. Actually, on closer inspection, it is not at all clear how the acquisition of intellectual skills and virtues could serve as an adequate response to the challenge of disagreement. For instance, one just has to envision the situation that students will find themselves in *after* they have used all these skills and virtues to develop their own moral views. In fact, many of the beliefs that students will form will be controversial, both among professional philosophers and among their classmates. Whenever students take a stance on a controversial moral issue, they will inevitably find themselves disagreeing not only with their

superiors, but also with their peers. And given that the epistemically destructive effects of these disagreements are equal to or greater than those of moral disagreements among professional philosophers, this will undermine the epistemic status of their beliefs.¹⁵

At this point, it becomes clear why proponents of moral knowledge transmission and proponents of skill and virtue development are really in the same boat. One main motivation behind promoting the development of intellectual skills and virtues was that it allows students to arrive at well-formed moral beliefs that serve as a suitable basis for decisions and actions, even in cases where the relevant moral questions are controversial among experts. However, this is not the case. When it comes to moral questions that are controversial among experts, it is simply not possible for students to arrive at justified verdicts on their own, because the disagreeing experts undermine the justificatory status of any belief that the students could possibly form.

Given this, it seems that the challenge of disagreement is not a specific challenge to moral transmission accounts, but rather a general challenge to all accounts of moral education that aim at the development of moral insights. What's more, it seems that knowledge transmission accounts are even less affected by the challenge of disagreement than their direct competitors. For instance, take cases where moral experts do in fact *agree*: in such cases, students who are encouraged to make up their own minds will likely still find themselves disagreeing with each other, which will again undermine their freshly formed beliefs' justificatory status. In light of this, it looks like skill- or virtue-oriented accounts of moral education are actually more severely affected by the epistemically destructive implications of moral disagreement than knowledge transmission accounts.

That being said, it is important to stress that the above considerations are in no way meant as a general rejection of the idea that moral education should also aim at the development of various skills and virtues. Enabling students to make their own independent judgements is rightly widely regarded as a central educational goal, and should certainly be an integral part of their philosophical and moral education. However, this idea is in no way incompatible with knowledge transmission accounts of moral education. Passing moral insights on to future generations doesn't mean to simply tell students what's right and what's wrong. Any pedagogically respectable realization of knowledge transfers will encourage and assist students to autonomously engage with the arguments and considerations that actually support the views that are presented to them. In fact, one might even argue that the transmission of moral knowledge already implies teaching the processes through which experts arrive at their conclusions. Understood in this way, enabling students to identify, evaluate and formulate ethical arguments and to critically assess the plausibility of different ethical theories is not just compatible with, but an integral part of moral knowledge transmission.

¹⁵ For a discussion of the epistemically destructive effects of disagreements with one's epistemic superiors, see e.g. Constantin and Grundmann 2018.

In fact, it seems that besides academic contexts, the educational system is one of the few real-life contexts where the highly artificial conditions for epistemic peerhood are plausibly often met. At least until they enter the university level, students are usually the same age. Furthermore, they have been provided with the same body of arguments as a basis for classroom discussion and have had the same amount of time to evaluate these arguments. Lastly, at least in some countries, students are grouped together on the basis of their grades, which further increases their epistemic homogeneity. Given all these peculiarities of the educational system, disagreements between students often have a remarking similarity with highly idealized scenarios like the Restaurant case.

Throughout such a process, students will inevitably develop a variety of valuable skills and virtues. And this is a good thing: while there are plausibly some important moral questions with respect to which we can and should provide students with well-founded answers, simply *informing* students about these answers clearly won't do. In order to have any impact on their practical decision-making, the moral principles that are taught in educational contexts need to be properly *grasped* by the students. Furthermore, students need to reliably identify the specific principles that are relevant to a given situation and then be able to adequately apply these principles in consideration of concrete situational circumstances. Lastly, students also need to critically reflect on the scope and the limitations of the principles that they have been taught in order to identify situations in which any definite moral judgement would be premature.

To be able to do all this, students obviously need to be provided with certain skills and virtues. So in a way, skill- and virtue-based accounts of moral education got it right by stressing the importance of the development of skills and virtues for successful moral education. However, they got it wrong by assuming that students can simply use these skills and virtues in order to arrive at well-founded judgements on moral questions that are even controversial among experts. In light of these considerations, it also gets clear that moral knowledge transmission – rightly understood – has nothing to do with wrongful manipulation or indoctrination. Indoctrination, as it is standardly defined, involves the attempt “to impart beliefs [...] [to someone] in such a way that she comes to hold them non-rationally, on some other basis than the force of relevant evidence and argument” (Hand 2018, 6). Purposefully providing students with those arguments and considerations that have been judged by experts to be most compelling seems like the direct opposite of indoctrination, since it helps students to arrive at the epistemically best supported moral views they can possibly have.

5 Conclusion

In this paper, I have discussed the so-called challenge of disagreement as a specific epistemic challenge to the idea that moral education should aim at the transmission of moral knowledge. Insofar as the above considerations are convincing, the observation that many if not most moral questions are controversial fails to pose a critical epistemic challenge to knowledge transmission accounts of moral education. First, the claim that many if not most moral questions are subject to persistent disagreement already requires some important relativization. While almost all moral questions may be controversial among laypeople, only disagreements among moral experts plausibly have the potential to undermine the educational transfer of moral knowledge. And although moral experts disagree about many moral questions, they apparently also agree on a surprising number of important moral issues. Second, there seems to be an attractive and hitherto neglected possibility of transferring moral knowledge even in cases where the relevant first-order moral questions are subject to persistent expert disagreement. This possibility consists in the transmission of knowledge about second-order moral norms that guide our decisions under moral uncertainty, and which are epistemically independent of the underlying first-order norms that are subject to disagreement. Given these considerations, the challenge of disagreement is not an *effective* challenge to the transmission of moral knowledge. Furthermore, it is important to notice that those controversies that actually persist among moral experts threaten

to undermine not just the prospects of moral knowledge transmission, but all accounts of moral education that consider the attainment of moral insights to be an important goal of moral education. In this respect, the challenge of disagreement is also not a *specific* challenge for the transmission of moral knowledge.

In light of these results, it seems that knowledge transmission accounts of moral education are in a much better position than is usually expected to incorporate the fact that moral questions are notoriously controversial. As I stressed at the outset, this in no way means that knowledge transmission accounts are superior overall to alternative accounts of moral education. Indeed, there are a number of further important challenges that must be met in order to plausibilize the idea that moral knowledge can and should be transferred in educational settings. However, given the amount and severity of the moral problems that we currently face, we shouldn't give up hope of also gaining and passing moral knowledge too quickly.¹⁶

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations • The author did not receive support from any organization for the submitted work.

- No funding was received to assist with the preparation of this manuscript.
- No funding was received for writing this article.
- No funds, grants, or other support was received.
- The author has no relevant financial or non-financial interests to disclose.
- The author has no competing interests to declare that are relevant to the content of this article.
- The author certifies that he has no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.
- The author has no financial or proprietary interests in any material discussed in this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Balg D (2021a) Live and let live. A critique of intellectual tolerance. Springer, Berlin
- Balg D (2021b) Toleranz. Was müssen wir aushalten? Metzler, Heidelberg
- Ballantyne N, Coffman E (2011) Uniqueness, evidence, and rationality. *Philosophers' Imprint* 11:1–13
- Bourget D, Chalmers DJ (2021) Philosophers on philosophy: The PhilPapers 2020 Survey. <https://survey2020.philpeople.org>. Accessed 24 Mar 2022
- Bourget D, Chalmers DJ (2014) What do philosophers believe? *Philos Stud* 170:465–500. <https://doi.org/10.1007/s11098-013-0259-7>
- Brueckner A, Bundy A (2012) On “Epistemic Permissiveness.” *Synthese* 188:165–177. <https://doi.org/10.1007/s11229-011-9921-9>
- Christensen D (2007) Epistemology of disagreement: the good news. *Philos Rev* 116:187–217. <https://doi.org/10.1215/00318108-2006-035>
- Christensen D (2009) Disagreement as evidence. The epistemology of controversy. *Philosophy Compass* 4:756–767. <https://doi.org/10.1111/j.1747-9991.2009.00237.x>
- Constantin J (2021) Epistemic defeat. De Gruyter, Berlin
- Constantin J, Grundmann T (2018) Epistemic authority. Preemption through source sensitive defeat. *Synthese* 196:199–213. <https://doi.org/10.1007/s11229-018-01953-x>
- ¹⁶ I would like to thank three anonymous reviewers for their very helpful comments on earlier versions of this paper.

- Driver J (2006) Autonomy and the asymmetry problem for moral expertise. *Philos Stud* 128:619–644. <https://doi.org/10.1007/s11098-004-7825-y>
- Elga A (2007) Reflection and disagreement. *Noûs* 41:478–502. <https://doi.org/10.1111/j.1468-0068.2007.00656.x>
- Feldman R (2007) Reasonable religious disagreements. In: Antony LM (ed) *Philosophers without Gods. Meditations on atheism and the secular life*. OUP, Oxford, pp 194–214
- Giesinger J (2021) Vermitteln und Mitteilen – Die Meinung der Lehrperson in der Diskussion kontroverser Themen. In: Drerup J, Zulaica y Mugica M, Yacek D (eds) *Dürfen Lehrer ihre Meinung sagen? Demokratische Bildung und die Kontroverse über Kontroversitätsgebote*. Kohlhammer, Stuttgart, pp 19–30
- Goldberg S (2010) *Relying on others*. OUP, Oxford
- Goldman AI (2006) Social epistemology, theory of evidence, and intelligent design. *Deciding what to teach*. *South J Philos* 44:1–22. <https://doi.org/10.1111/j.2041-6962.2006.tb00026.x>
- Grundmann T (2013) Doubts about philosophy? The alleged challenge from disagreement. In: Henning T, Schweikard D (eds) *Knowledge, virtue, and action. Putting epistemic virtues to work*. Routledge, London, pp 72–98
- Grundmann T (2019) How to respond rationally to peer disagreement. The preemption view. *Philos Issues* 29:129–142. <https://doi.org/10.1111/phils.12144>
- Hand M (2014) Towards a theory of moral education. *J Philos Educ* 48:519–532. <https://doi.org/10.1111/1467-9752.12116>
- Hand M (2018) *A theory of moral education*. Routledge, London
- Haydon G (2003) Moral education. In: Curren R (ed) *A companion to the philosophy of education*. Blackwell, Malden, pp 320–331
- Kelly T (2005) The epistemic significance of disagreement. In: Hawthorne J, Gendler T (eds) *Oxford studies in epistemology*, vol 1. OUP, Oxford, pp 167–196
- Kelly T (2011) Peer disagreement and higher-order evidence. In: Goldman AI, Whitcomb D (eds) *Social epistemology. Essential readings*. OUP, Oxford, pp 183–221
- Kölbel M (2004) Indexical relativism versus genuine relativism. *Int J Philos Stud* 12:297–313. <https://doi.org/10.1080/0967255042000243966>
- Lasersohn P (2005) Context dependence, disagreement, and predicates of personal taste. *Linguist Philos* 28:643–686. <https://doi.org/10.1007/s10988-005-0596-x>
- MacAskill W, Bykvist K, Ord T (2020) *Moral Uncertainty*. OUP, Oxford
- Matheson J (2021) Applying Moral caution in the face of disagreement. *Ethical Theory Moral Pract*. <https://doi.org/10.1007/s10677-021-10155-x>
- Meyer K (2011) Moralische Bildung im Philosophie- und Ethikunterricht. In: Limbourg M, Steins G (eds) *Sozialerziehung in der Schule*. VS Verlag für Sozialwissenschaften, Wiesbaden, pp 223–237
- Musschenga AW (2009) Moral intuitions, moral expertise and moral reasoning. *J Philos Educ* 43:597–613. <https://doi.org/10.1111/j.1467-9752.2009.00707.x>
- O'Donnell AM (2012) Constructivism. In: Harris KR, Graham S, Urdan T, McCormick CB, Sinatra GM, Sweller J (eds) *APA educational psychology handbook, Vol. 1. Theories, constructs, and critical issues*. American Psychological Association, Washington D.C., pp 61–84
- Pfister J (2019) Classification of strategies for dealing with student relativism and the epistemic conceptual change strategy. *Teach Philos* 42:221–246. <https://doi.org/10.5840/teachphil2019730107>
- Priest M (2016) Inferior disagreement. *Acta Analytica* 31:263–283. <https://doi.org/10.1007/s12136-015-0277-5>
- Rosen G (2001) Nominalism, naturalism, philosophical relativism. *Noûs* 35:69–91. <https://doi.org/10.1111/0029-4624.35.s15.4>
- Schoenfield M (2014) Permission to believe. *Noûs* 48:193–218. <https://doi.org/10.1111/nous.12006>
- Singer P (1982) How do we decide? *Hastings Cent Rep* 12:9–11. <https://doi.org/10.2307/3561822>
- Slote M (2009) Caring, empathy and moral education. In: Siegel H (ed) *The Oxford handbook of philosophy of education*. OUP, Oxford, pp 211–226
- van Inwagen P (1996) It is wrong, everywhere, always, for anyone, to believe anything upon insufficient evidence. In: Jordan J (ed) *Faith, freedom, and rationality. Philosophy of religion today*. Rowman & Littlefield, Lanham, pp 137–154
- Wedgwood R (2010) The moral evil demons. In: Feldman R, Warfield TA (eds) *disagreement*. OUP, Oxford, pp 216–246
- Wright C (2001) On being in a quandary. Relativism vagueness logical revisionism. *Mind* 110:45–98. <https://doi.org/10.1093/mind/110.437.45>