

*Moral advice and joint agency**

Eric Wiland

University of Missouri-St. Louis

1. There are many alleged problems with trusting another person's moral testimony. For one thing, those testifying about moral matters sometimes have their own agenda, and it is difficult to know whether they have your best interests at heart (Driver 2006). Even if you can solve that problem, you face another: if you yourself don't know what you should do, and so are on the lookout of helpful testimony, it is difficult to figure out whose testimony is worth trusting, especially problematic given the fact that the people willing to give moral testimony rarely completely agree with one another (Cholbi 2007).

Let's here set most of these problems aside, and suppose that you have every reason to think you've identified the perfect moral testifier. Call her Sophie. You tell Sophie your situation and your problem. You tell her what matters to you. She and you agree about the nonmoral facts. You and Sophie try to work out what the various goods and bads and reasons at stake are. It's a welter of practical data to juggle, and you can't figure out what to do. But Sophie can, and she tells you that you ought to V. You yourself can see the reasons that favor V-ing and the reasons against V-ing, as well as the reasons for and against your other action-options. But since you don't see how it all adds up, you'll just have to trust Sophie's moral testimony. You just trust her that you ought to V.

Is everything ok? According to several recent arguments, no. Even if you accept Sophie's moral testimony, problems remain. Perhaps the most prominent complaint is that trusting Sophie's moral testimony still fails to deliver understanding why V-ing is the thing to do. You then may know *what* to do, but you won't completely understand *why* you should do it. And understanding what to do is important, or intrinsically valuable, or something the absence of which is highly regrettable. So while trusting moral testimony may be how to make the best of a bad situation, it is far from ideal. It's much better to understand for yourself why you should V (Hills 2009).

Is this really so? Here I'll concede for the sake of the present argument both that one cannot get moral understanding from trusting moral testimony, and that there's something intrinsically good about having moral understanding.

* I thank to participants at the 2016 Bled Conference on Ethics, the 2016 Rocky Mountain Ethics Conference and the 2017 Arizona Workshop in Normative Ethics for comments on earlier versions of these ideas. In particular, I thank Kathryn Lindeman for many conversations about these ideas.

Other writers complain that this lack of understanding has a *further* unappealing consequence. When you V on Sophie's moral testimony, you do the right thing in the sense that you do better by V-ing than you would by opting for any alternative to V-ing. You do *what* is right. But it's argued that you don't do as well as you would do if instead you had V'd from your own understanding *why* V-ing is the thing to do. A person who Vs because he himself understands why V-ing is the thing to do (call this person Otto) *does* better than you do when you V on Sophie's testimony. The value of Otto's action is greater than the value of yours, this despite the fact that you both V. This difference in understanding, intrinsically valuable or not, thus leads to a further difference in the value of the *action* undertaken (Nickel 2001, McGrath 2011).

But why think this? Recall Kant's discussion of the shopkeepers. Two shopkeepers each refuse to overcharge their customers, both thus doing what's right. One does it to preserve his shop's good reputation, thereby doing what's good for his own long-term happiness. The second does it simply *because* it is the right thing to do. The moral value of the second shopkeeper's action, Kant argues, is greater than that of the first, this even though in some sense they both *do* the same thing. The moral worth of an action depends not (only) on *what* is done, but (also) *why* it is done. Doing the right thing for the right reason is better than doing the right thing merely to stay out of trouble. This thought is highly intuitive.

Come back now to the example of you taking Sophie's moral testimony that you ought to V. When you V, you do the right thing. But you won't do the right thing *for the right reason*, this because you yourself don't understand why V-ing *is* the right thing. If you *did* understand this, you wouldn't need to trust Sophie in the first place. Just as the first shopkeeper does the right thing, but not for the right reason, you too do the right thing, but not for the right reason. And since you don't do the right thing for the right reason, then it appears your action lacks full *moral worth*. It is *less* worthy than it would have been if you instead were to so act from the right reasons, reasons whose rightness you yourself grasp. Thus, trusting another's moral testimony is putatively never the best way to act; at most, it is a regrettable way to cope with one's own imperfections as well as one can. Call this, then, *the moral worth objection* to moral testimony. The moral worth objection centers on the idea V-ing as a result of trusting moral testimony isn't as good as V-ing for the reasons that make V-ing right, reasons you grasp yourself. There are several rejoinders to this objection to trusting moral testimony, some of which I offer elsewhere.

Here I launch a brand new response to the moral worth objection. I'm granting here that the moral worth objection highlights an unfortunate aspect of trusting moral *testimony*. But moral *advice*, I'll argue, differs from moral testimony, differs from it in a way that enables a defender of moral advice to parry the moral worth objection. Trusting moral advice thus has dialectical benefits not had by trusting moral testimony. Or so I'll argue.

The basic idea I defend is that an advisor and an advisee can together constitute a *joint agent*, and that this joint agent's action can have moral worth. So while the advisee himself might not do the right thing for the right reason (this because all alone he lacks the right reason), and while the advisor herself might not do the right thing for the right reason (this because all alone she does not do the right thing), they *together* do the right for the right reason. And so *both* Otto's action and this joint agent's action have moral worth. Acting on the basis of moral advice thus can have moral worth.

I'll defend this claim by arguing 1) that highly informal duos can exhibit joint agency, 2) that joint agents can be constituted by individuals whose contributions are highly idiosyncratic, 3) that a commander and a commandee can exhibit joint agency, 4) that an adviser and an advisee can likewise exhibit joint agency, and finally 5) that these actions can be morally evaluated and have moral worth.

2. Here I'm arguing that under certain circumstances (to be partially specified below) when an advisee is taking an adviser's advice, they are acting together. They together do something, in virtue of a) the adviser's advising the advisee and b) the advisee's trusting and taking her advice. In this section, I want to explicate the more general notion of joint agency with which I'm operating.

Some philosophical discussions of collective agency focus upon the doings of highly formalized groups: corporations, states, courts, teams, clubs, and the like. These groups have three distinctive features. First, they often have names, names that enable one to refer to them as singular entities. Consider, for example, Apple, Bulgaria, the Lion's Club, and the Miami Heat. All of these collectives do things, and sentences that describe what they do can be put in the singular, e.g. "Apple *is* (not are) hiring a new CFO", "Bulgaria *is* importing more oil", and so on.

Second, these collectives typically persist through long stretches of time. They don't pop in and out of existence.

A third feature distinctive of these groups is that typically they act on the basis of fairly stable and knowable decision procedures. When Apple hires someone, its hiring is the upshot of a predictable array of actions of individuals in the corporation: someone requests the position; someone else approves the request; someone else writes an advertisement...and so on. When courts issue verdicts, they do so as the upshot of a rather simple decision procedure, typically majority rule (List and Pettit 2011).

When an adviser and an advisee together exhibit joint agency, none of these features are typically present. An adviser and an advisee typically do not form a duo that we can refer to with a name, as we can with Microsoft or Haverford College. They might be able to be referred to only with plural pronouns ('they', 'them') or the conjunction of their personal names with the adverb 'together'. Further, the joint agent that they constitute typically does not persist through long stretches of time. Many instances of joint agency comprising an adviser and an advisee are one-off actions: I might take your advice once and only once. Finally, joint agents comprising an adviser and an advisee will rarely if ever act according to a stable and knowable decision procedure. Sometimes I'll take your advice and sometimes I won't, but there may be no pattern that distinguishes the one case from the other, nor that characterizes how we go about it in those cases in which I do take your advice.

Given these several differences between standard cases of collective agency and what an adviser and advisee do, why think that an adviser and an advisee *do indeed* exhibit joint agency? More to the point, aren't *these* differences grounds for concluding that there is *no* joint agency in the case where an advisee takes his adviser's advice?

No. Named persistent formal groups are but one kind of joint agent. Margaret Gilbert, who has explored the nature of joint agency as much as anyone, frequently explains its features by discussing the pedestrian example of two people walking together. Two people who are walking together exhibit joint agency: there is something that they together are doing that is (arguably) not reducible to the sum of what each pedestrian is doing. There is a difference between two people merely walking side by side and two people walking together. The adverb 'together', then, signifies a different mode of agency: joint agency. On this view, two people walking together pool their wills in a way such that each has claims on the other to do their part to realize their walking together (Gilbert 1990).

The details of her account don't concern us much here, but perhaps I should articulate her basic insight in a more general key. There are cases in which two people (X and Y) are each doing something (V-ing) such that each of the following is true:

X is V-ing
Y is V-ing
X and Y are V-ing,

and yet it is false that X and Y are V-ing together. This will be most obvious in cases where the type of action in question *cannot* be something that a joint agent does (e.g. soloing). But it is also often true of many ordinary verbs: reading, typing, bathing, etc. Although it is possible for two people to, say, read together, usually two people who are reading at the same time in the same locale are not reading together. They just happen to

be doing the same thing as the other person in the same general vicinity. Genuine joint agency, by contrast, is typically captured by the adverb ‘together’, and includes not only actions that *must* be joint (e.g. tangoing, conversing), but also many others that are often so (e.g., singing, dining, talking, building).

All I hope to establish right now is that cases of taking advice are not disqualified from exhibiting joint agency merely on the grounds that the adviser and the advisee do not form a named group, that they don’t persist through time, nor that they lack a stable decision procedure by which they act. Two people who are walking together exhibit none of these characteristics either. Thus, if something as informal and transient as two people walking together can exhibit joint agency, perhaps an adviser and an advisee can too.

3. But of course there remain significant differences between the cases of joint agency that Gilbert investigates, such as walking together, and the case of an adviser and advisee. So we still may be tempted to think that advisers and advisees don’t cooperate as members of a joint agent.

There are, in fact, at least three ways the case of advice differs from Gilbert’s case of walking together. First, in the case of walking together, each member of the joint agent does roughly the same thing: walking. Second, each member of the joint agent does *something*, whereas it may seem that an adviser herself doesn’t really *do* anything at all (other than advising). Third, both of the walkers are walking at the same time, while in the advice case the adviser completes her advising before the advisee does *his* part. I’ll argue, however, that these differences do not disqualify the advising case from counting as a case of joint agency.

First, consider the fact that in Gilbert’s example of two people walking together, each person does roughly the same thing. They are not merely each individually walking, but they are walking together. More generally, when two people are V-ing together, each person is V-ing: e.g., they are eating together, and each is eating. That’s part of how the way the adverbial phrase ‘V-ing together’ seems to work. And so it looks like it’s a necessary or at least central case of joint agency.

But we need to expand our diet of examples. Two people who are playing a duet are clearly acting together and are exhibiting joint agency no less than our pair of walkers. But this musical duo may comprise one violinist and one pianist. Thus the two of them are not doing ‘the same thing’, as anyone pianist who has goofed around on a violin knows. Although they are playing a duet together, neither is individually playing a duet, nor is either doing anything like the same thing that the other is.

Of course, each is playing a musical instrument, and so in *this* sense they are doing roughly the same thing as the other. But many other clear examples of joint agency don't lend themselves to even this sort of analysis. Suppose you and I are hosting a dinner party together. You might invite the guests, while I might construct the shopping list; and then you buy the groceries, during which I clean the kitchen; after which you cook the food, during which I take a nap; and then I re-clean the kitchen after our guests leave, while you deservedly take a bath and read a book. This rather ordinary example of joint agency may comprise many different individual activities no two of which are even roughly of the same kind. Many joint agents divide labor via specialization, as well-meaning egalitarian parents eventually realize. And so the mere fact that an adviser and an advisee do different things does not disqualify the advice case from counting as a genuine case of joint agency.

4. Let's examine the second worry listed above: in cases of joint agency each individual *does* something; but in cases of advice, the adviser doesn't really do anything. She merely directs another's activity, but she herself isn't any further involved in what transpires.

To address this, we need to understand the different ways the activity (or the labor) of a joint agent may be divided among the individuals who constitute it. In the case of our walkers, each individual does the same thing as the other individual of the duo: each walks. At a more general level, something similar holds of our musicians. But not all joint agents resemble these. Here are two examples.

Consider an orchestra. The St. Louis Symphony Orchestra comprises dozens of people. There are several violinists, several cellists, several percussionists, and many other musicians as well. There are differences among them, but they all play music. But there is also the conductor, David Robertson, who plays no musical instrument in the orchestra. The conductor instead *directs* the musical performance, and the musicians take the conductor's direction about things like tempo, entry, phrasing, and so on. The conductor doesn't make a sound, and yet is clearly one of the individuals who together constitute the Orchestra. The conductor counts as such in virtue of directing the other members of the Orchestra. It would be silly to insist that while the various *musicians* constitute a joint agent, every joint agent that they do constitute excludes Robertson on the grounds that *he* does not play an instrument. Instead, the labor of an orchestra is very broadly divided. Orchestras include conductors, and the activity of the orchestra depends upon the activity of the conductor, even though his activity is 'merely' directive in nature.

We see something similar in the case of teams. The Golden State Warriors has various individuals on the team: Stephen Curry, Draymond Green, and so on. Each of these is a player. But the coach Steve Kerr is also a member of the Golden State Warriors. Yet the

coach only directs (commands, advises, encourages, chastises) the players. It would be silly to insist that while the various players together constitute a joint agent, every joint agent that they do constitute excludes Kerr on the grounds that *he* does not play in the game. The labor of a team is very broadly divided. Teams include coaches, and the activity of the team depends upon the activity of the coach.

So, orchestras and teams are examples of agents that a) act, b) have members who direct the individual actions of other members, and thus c) whose directing members contribute to some of the things that the collective itself does. Those who direct others can qualify as members of joint agents. And so the fact that advisers ‘merely’ direct others’ activities does not thereby disqualify them from being members of the joint agents whose activities they steer.

5. In conversation I have found that many people experience the third worry mentioned earlier: the fact that what an adviser does typically completely *precedes* what her advisee then does—*this* fact makes it difficult for us to grasp the unity of what they putatively do together. When Robertson directs the musicians in the orchestra, he is directing them *while* they are playing. When Kerr coaches the players on his team, he is coaching them *while* they are playing. By contrast, once I advise you to do something, what I do is over; and if you then take my advice, you begin what you do only after I’ve already done my bit. So, given this temporal gap, why might we be entitled to say in the advice case that *we* do anything at all?

The problem can appear even more extreme if we imagine cases where the gap between advice and execution of advice is very long indeed. For example, suppose that when you were young your grandfather advised you not to retire when you hit 65, but to work until you no longer can. (He is regretting his own choice.) You don’t immediately decide to take his advice, but you file it away for future consideration. Several decades later as you approach 65, you remember his advice, and decide to take it, working until you can no more. Now, it seems rather hard to believe that the two of you are doing anything together whatsoever! Your grandfather is not even alive when you take his advice, and it just isn’t true that the two of you together retired at age 65. (He alone did not.) The long temporal gap between his advice and your taking it makes it difficult if not impossible to spot any joint agency in the works here. But this example is just an exaggerated version of the standard case of advice: advice always entirely precedes the taking of the advice. So if the grandfather case of advice is not an instance of joint agency, it appears that no case of advice is an instance of joint agency.

My response to this worry has several parts, most of which are conciliatory. First, I want to point out that the (alleged by me) unity of advice and its execution bears structural

similarity to the (again, alleged by me) unity of an agent's intention and its execution. When you fully and wholeheartedly intend to do something that you can do, it's often appropriate for you to describe yourself as already doing the thing you intend, e.g. 'I'm writing a dissertation on Leibniz'. Making the decision to do something can make it true that you are already beginning to be doing it. What this fact shows about advice is unclear, but, since intentions always precede their executions, it should at least make us cautious about dismissing cases of advice as instances of joint agency merely on temporal grounds.

Second, one complicating factor of the grandfather case is that there is a long stretch of time between when your grandfather advises you to do something, and when the time arrives at which it would be appropriate to implement his advice. This large temporal gap threatens to undermine the alleged unity of his advice and its execution. I am not sure what to say about cases like this. Thus I will restrict my conclusions to cases of advice where there is not a prolonged temporal gap between advice and its time of execution.

Third, and similarly, another complicating factor of the grandfather case is that in it you don't even *decide* to take his advice until long after he has died. So, I will restrict my conclusions to cases of advice where an advisee decides to take his adviser's advice soon after that advice is delivered.

Putting these last two points together, I will restrict my conclusions to cases in which an adviser advises an advisee to do something rather immediately, and stays around to see whether the advisee takes this advice, ready to help if further guidance is called for. Once we see what to say about cases like these, then later we can relax these simplifications to see what to say about other cases of advice. Lest this seem *ad hoc*, recall that in the simplest cases of successful *testimony*, the hearer usually accepts what the speaker says soon after the speaker says it, rather than wait around a while before deciding to believe what she says. So for now, then, let's focus on cases of advice that work the same way.

6. We've seen that despite the fact that it may initially seem peculiar to regard execution-of-advice cases as instances of joint agency, we have not — at least as of yet — identified any features of these cases that are not also present in some *other* standard cases of joint agency. Joint agents can be highly temporary, unnamed, and operate via a division of labor wherein some member takes a merely directive role.

But in order to argue that some cases of moral advice escape the moral worth objection, I also need to argue the putative actions of joint agents performed by adviser-and-advisee are indeed morally evaluable. If the only morally evaluable actions on the scene are those of each individual, then my argument will go nowhere. What the adviser does alone can

be morally evaluated. When the advisee then does alone can be morally evaluated as well. But why should we be confident that anything *they together* do can be morally evaluated? Are the actions of adviser-and-advisee really the sort of thing that can possess moral worth?

I will argue that they are, and I will do so by analogy. The criminal law provides us with a fruitful resource for understanding the nature of joint agency and *responsibility* in cases where one individual of a pair plays a merely directive role. When one person directs another to commit a crime, the law typically regards *each* of them as a perpetrator of that very offense, and each is criminally liable and blameworthy for the offense. The best justification for so regarding things, I'll argue, is because the two are acting *together*. *Their* action violates the law. (After establishing this, I will argue that morality operates as the law does.)

Let's illustrate with a simple fictional example. Suppose Barney robs Moe's bar. Unless there is something unusual about the circumstances of the crime or Barney's mental state, Barney is criminally liable for robbery.

Now suppose further that Homer *ordered* Barney to rob Moe, and that Barney executed Homer's command. Most systems of criminal law treat Homer not as a mere accessory to the crime of robbery, but as a second *principal* to the robbery, and so also guilty of robbery. That is, according to the law, *Homer* robbed Moe.¹ Of course, Barney robbed Moe too. They each did. And Homer and Barney can each be convicted of robbery, facing equally severe punishments.² There is something about ordering a crime that draws the orderer into the very heart of the crime committed. Those who order underlings to commit crime are liable for the crimes thereby committed and are punished at least as severely as the underlings who carry out the order. But how should we account for the nature of the agency involved here? What is it about Homer's command that makes *him* a robber?

Answer: *Homer* robbed Moe's bar only in virtue of the fact that *Homer and Barney together* robbed Moe's bar, and *this* is so because Barney's obeying Homer's order suffices (in normal circumstances) to make it true that Homer and Barney together robbed Moe's

¹ Anscombe (1981) similarly argued that Truman killed Japanese civilians, this in virtue of his ordering his soldiers to bomb them with the further intention to kill them.

² Things may be even worse for the orderer when the orderer occupies a role in a persisting organization. See Eldar 2010, 183.

bar.³ In other words, when Homer effectively orders Barney to rob Moe's bar, they together rob Moe's bar, and thus each of them thereby individually robs Moe's bar.⁴

Here's some evidence for this answer: if Homer orders Barney to rob Moe's bar, but Barney disregards the order, then Homer is *not* criminally liable for robbing Moe's bar. That is, *merely* ordering Barney to rob Moe's bar is not enough to make Homer criminally liable for robbing Moe's bar. (In some jurisdictions, Homer may still be guilty of some lesser crime.) Homer is criminally liable for robbing Moe's bar (in virtue of ordering Barney to rob it) only if Barney indeed robs it. Thus when Homer *is* criminally liable for robbing Moe's bar, this is not *merely* because Homer ordered Barney to do so, for that was not sufficient to trigger criminal liability for robbery in the disregarded-order case. What, then, explains why Homer *is* criminally liable in the executed-order case? Only the fact that they *together* robbed Moe's bar adequately explains Homer's criminal liability. The fact that they together robbed it makes it just to punish both of them for robbing it. Nothing less would do.

Of course, we are and should be reluctant to punish individuals merely because they are members of groups other members of which have acted badly. The practice of regarding an individual as vicariously liable for what others do has an ugly history. But those who order others to act badly are not *vicariously* liable for what others do. They are more directly involved in the wrongdoing that they order. They are an integral part of the actions performed by their underlings or subjects. Just as David Robertson gets credit when the orchestra he conducts performs well, and Steve Kerr gets credit when the team he coaches plays well, so too one who orders another to break the law is appropriately blamed and held responsible for the crime that they together have committed.

7. At this point, the reader might grant that while two people can act together in virtue of one effectively *ordering* the other, the same is not true of one who merely effectively *advises* another. Even if it's true that orderers are caught up in joint agency, it does not follow that something similar is true of advisers.

To make the case that advising is indeed sufficiently similar to ordering in the relevant respects, I'll do two things. In this section, I'll (finally!) distinguish advice from testimony,

³ "It is a principle of law that if several persons act *together* in pursuance of a common intent, every act in furtherance of such intent by each of them is, in law, done by all." Macklin, Murphy, and Others' Case (1838) 2 Lewin CC 225, *emphasis* mine.

⁴ I understand an *effective order* roughly to be an order that is obeyed at least in part because the action ordered was so ordered. Orders that are disregarded are thus not effective. Later, I will understand the term *effective advice* analogously.

including moral testimony. In the next section, I'll review how the criminal law tends to treat those who effectively advise others to commit crimes.

So first, here is a very brief sketch about some differences between moral testimony and advice. One might think that moral testimony and advice pretty much amount to the same thing. But this is not so. As a first approximation, testimony is a kind of communication about *how things are*, and so a speaker who offers testimony does so with the intention that the hearer *believe* something he might not otherwise believe. By contrast, advice is a kind of communication about *what to do*, and so a speaker who offers advice does so with the intention that the hearer *do* something he might not otherwise do.

The difference between *trusting* testimony and *trusting* advice can be understood in much the same way. A hearer trusts testimony only if he thereby *believes* the content of the testimony. So if you trust someone's moral testimony, then (roughly) you come to believe a moral proposition you hadn't believed before. Suppose someone tells you that patience is a virtue, or that you should donate more money to the poor. If you trust her testimony, you'll *believe* what she tells you: that patience is a virtue, or that you should donate more money to the poor.

Trusting advice is different. Advice is advice to *do* something, not to believe something. So to trust someone's advice is (roughly) to do what is advised. If someone advises you to be patient, or to donate more money to the poor, then you take her advice only if you *are* patient, or if you *donate* more money to the poor. Trusting advice requires action. And while trusting moral testimony often has practical consequences too, it need not. It is possible to trust another's moral testimony, and yet not act any differently. You might trust another's testimony that you should donate more money to the poor, but never actually donate more money. (Perhaps the overwhelming feeling of guilt about one's heretofore stinginess immediately induces a fatal heart attack.)

It is possible to offer both moral testimony and advice by the very same speech act, e.g. "It would be right for you to tell the truth", said in a context where, say, both speaker and hearer understand that they are both indefeasible moral judgment internalists. Such a speaker at once directs the hearer both to believe something and to do something, and the hearer understands this and takes it in this spirit. But moral testimony and advice can also come apart, e.g. "It would be just for you to turn yourself in to the police", said in a context where the speaker and the hearer mutually understand that they are both immoralists. This may count as a bit of moral testimony, but it sure isn't any kind of advice.

So advice is rather different from moral testimony. But is it at all like commands or orders? Both advice and commands are specific forms of directives. If Sophie commands

you to V, she is directing you to V. And if she instead advises you to V, she is likewise directing you to V. The same goes for other directives: recommendations, requests, and pleas (Searle 1976).

But there are differences. Here's one. If Homer commands Barney to rob Moe's bar, Homer in no way presupposes the thought that it was true all along that Barney should rob Moe's bar, and Homer is merely pointing this out to Barney. Rather, Homer seems to be *making it true* that Barney should rob Moe's bar, this in virtue of his command. He's radically changing what Barney has reason to do.

Advice, however, typically operates differently. If Homer sincerely and felicitously advises Barney to rob Moe, then Homer thinks that robbing Moe is what Barney *already* has reason to do, although perhaps Barney does not recognize this yet, and so Homer merely intends to point this out to Barney. It might seem to make more sense to regard Homer as a robber when he (tries to) make it true that Barney has reason to rob Moe than it seems when he merely intends to point out to Barney that robbing Moe is what he has reason to do. Homer seems more integrally involved in the robbery in the command case than in the advising case. After all, in the advice case, Barney might have arrived at the view that he ought to rob Moe all on his own without Homer's directive, especially if that really is what he has most reason to do. In the case of command, this seems much less likely. This is why it may seem that Homer-and-Barney compose a joint agent in the command case but not in the advice case.

I will concede that in *some* cases of advice, it is hard to understand the adviser as a constituent of a joint agent which itself is acting in the way advised, as the criminal law understands those who order others to act. This is especially so when the advisee was already leaning in the direction of the advice anyway. Sometimes an adviser seeks advice merely to garner a second confirming opinion.

But remember that the moral worth problem arises in cases where the advisee does *not* grasp what he should do, and is *not* already so inclined to do the thing he is later advised to do. We are to imagine the advisee as highly undecided, but ready to do what Sophie advises. Without Sophie, he does not act, at least not in the way he does when he follows her advice. So when such an advisee *trusts* advice in this way, the adviser supplies a *necessary* and critical element of the act. Although there are surely some interesting differences between advice and command, they are alike in that both cases the director and the directed can play an ineliminable role in what they together do.

In fact, this sheds light upon why many people may be reluctant to *offer* advice in the first place. If you offer mistaken advice, and your advisee trusts you, you haven't merely *said* something false or bad. You would also be on the hook for what has been done. You

often may rightly be blamed for *that*. It thus is usually safer (although not necessarily thereby all-things-considered better) to dodge this responsibility by keeping mum.

It is true that those who order others typically wield some power over those they order, threatening to make things worse for the hearer if the hearer disobeys the order, even if they threat is only implicit. Advice, by contrast, is devoid of the threat of additional sanctions for disregarding advice. An adviser won't punish you for ignoring her advice.⁵

But the fact that orders and advice differ from each other in *this* way fails to support the claim that an adviser and an advisee can constitute a joint agent. If anything, this fact may seem to cut in the other direction: if a commander orders an underling to V under the threat of a severe penalty — a penalty in the absence of which the underling would not V — then we might view the commander as the *only* true agent on the scene, while the underling is little more than a puppet whose strings are being pulled. If I command you to steal a candy bar from a convenience store, and you know that I will slaughter your family if you disobey my command, then it does seem that *you* aren't really exercising your rational agency at all, and so that *we together* aren't really doing anything either. Only I am. In any case, the fact that an adviser does not threaten her advisee with a punishment for disregarding her advice—*this* fact does not obliterate what would otherwise be joint agency.

It *does* matter that the adviser is really giving full-blooded advice, and that the advisee take this advice in the same spirit. So, if someone says to you “Well, if I were you, I would V, but do what you want”, this is not full-blooded advice, at least not as I understand it. Or, if someone says “V-ing would be a great idea, but don't take my word for it”, this also is not full-blooded advice. Someone advises you only if she invites you to *rely* on her, much as someone offers you testimony only if she invites you to rely on her (Hinchman 2005). Reporting the contents of your beliefs by itself is not testimony, and predicting what you would do were you in the hearer's circumstances by itself is not advice. Advising someone requires that you intend to *direct* them.

Advising someone (as I'm understanding it) also involves more than merely discussing the hearer's practical situation, rehearsing arguments for and against various options, and aiming to persuade the hearer to see for himself that a specific practical option is the one to opt for. These various activities are often precursors to the need for advice. They are the materials of joint deliberation. But, as with testimony, advice comes in when persuasion falls short. People who are open to trusting advice are those who remain unpersuaded by the reasons and the arguments to which they have been exposed.

⁵ But we should be careful not to overstate this difference: acting against another's advice opens you up to the 'I-told-you-so' sanction.

This partially explains why in most cases a good adviser will not aim to lead another person to act *against* the hearer's pre-existing intentions and plans, at least not by means of advice alone. If I know that you intend *not* to V, I probably won't (merely) advise you to V. I may instead only explain why I see things differently, giving reasons and arguments whose conclusion supports V-ing. If these reasons and arguments don't dislodge your intention not to V, I'll then refrain from inviting you to rely on me. Instead, typically I'll advise you to V only if I think that you are feeling somewhat *uncertain* what to do, that you are open to my advice, and that you think that I am someone whose advice is worth trusting.⁶

So full-blooded advising is directive, most naturally expressed in the imperatival mood, albeit politely. And in this way, advice is much like orders and commands.

8. With this understanding of advice, let's look at complicit liability in the criminal law a bit more. Surprisingly, in most jurisdictions the criminal law tends *not* to distinguish liability for commanding a crime and liability for advising a crime. Consider, for example, the following section of the United States Code:

US Code, Title 18 Part 1 Chapter 1 section 2:

*(a) Whoever commits an offense against the United States or aids, abets, **counsels**, commands, induces or procures its commission, is punishable as a principal. (**emphasis** mine).*

Not only does U.S. law punish those who *command* others to commit a criminal offense as guilty of committing that very offense, but it also similarly punishes those who *counsel* others to commit a criminal offense. The standard of complicit liability is the same for each. Neither is merely guilty of some lesser crime.

The U.S. is not peculiar: many other systems of criminal law do not treat advising differently from other ways of directing others to do things. In Roman law, someone who merely persuaded slaves to run away from one master to a better one was himself liable of theft.⁷ In Rome "the advice to commit an offence was considered an offence itself if it was

⁶ There can be a two-step process: 1) the reasons and arguments in favor of V-ing I offer you might move you to give up your intention not to V, leaving you in a state of uncertainty about whether to V, a state in which 2) you are open to my advice about whether to V, and so I *then* might invite you to rely on me about what to do.

⁷ *Digesta* 47.2.36.

specific enough to persuade the perpetrator to commit the offence”.⁸ Similarly, the Talmud calls for punishing someone who incites another person to commit an offence, treating the inciter as a primary accomplice. This is notable because the Talmud more generally *prohibits* imposing criminal liability on accessories, this because they are merely indirect and secondary accomplices.⁹ Inciting others, however, is taken much more seriously: apparently, it is worse to encourage someone else to commit a crime than it is to aid him in committing it. Likewise, in German criminal law, strongly encouraging another to commit a crime is punished more harshly than is being an accessory to that crime.¹⁰ English law holds that “[w]hosoever shall aid, abet, counsel or procure the commission of any misdemeanour...shall be liable to be tried, indicted, and punished as a principal offender.”¹¹ And Article 31 of the Criminal Code of Bosnia and Herzegovina stipulates that one who advises another to perpetrate a criminal offence should be punished as if he himself perpetrated such an offence. Thus, *many* systems of criminal law regard those who advise and encourage others to commit crimes no differently than those who order others to commit crimes.

Note that in many of these cases, the adviser is not typically a member of a *formal* criminal organization: a person can be criminally liable for the crimes they encourage even when they are only counseling others to run away from their masters or to steal property. Nevertheless, the law treats such counselors or advisers much as it treats those who direct the actions of more formal organizations. You can be criminally liable for the crimes you have counseled others to commit even though you do not direct a formal organization. Rather, in virtue of a) your counseling him to commit a crime, and b) his taking your counsel and committing that crime, you are *thereby* a member of something at least akin to a joint agent.¹²

Many systems of law, then, criminalize the act of advising another to do something that violates the law, and many of *these* systems regard the adviser as also guilty of the very offense perpetrated, at least when their advice has been executed. Thus the criminal law tends to treat advisers as it does commanders, and the best explanation for regarding *either* sort of director as complicit is the same: it is because the director occupies a role in the joint agent whose actions violate the law.

⁸ Hallevy 2012, 4.

⁹ *Ibid.*, 7.

¹⁰ *Ibid.*, 9. See Stewart 2015, 540.

¹¹ Accessories and Abettors Act 1861, sections 1-7.

¹² See, for instance, Article 29 of the Criminal Code of Bosnia and Herzegovina.

7. I am proposing that at least sometimes when an advisee trusts and takes his adviser's advice, they act together. I admit that that *is* a bit weird, but I'm not making an even weirder claim. Although in *some* cases when collectives act, it can also make sense to attribute the action to the individual who directs the collective (e.g. coaches win when their teams win; dons kill when their families kill), this is not always so. Sometimes when an individual directs a collective, the action of the collective (e.g. overriding a veto) is *not* attributable to any individual in the collective, or this action (e.g. rebounding a basketball) is attributable only to an individual who occupies a more executive role, not to every individual in the collective. Likewise, while I think that the pair of adviser and advisee often does *something* together, I don't think that this action can always be attributed to the individual adviser too. Thus when Sophie advises you to do something, and you trust and take her advice, you-and-she do something together, but perhaps she alone only advises you. She need not be doing much else.

I think we can better understand what an adviser-and-advisee do (and don't do) by finally considering a particular example of what we might consider moral advice. Suppose you are caught up in a standard trolley case. You can save five people only by turning the trolley you are on, but turning the trolley would then kill one. Although aware of all the relevant facts, you remain highly uncertain about what to do. Fortunately, you have enough time to ask Sophie for advice — perhaps she, unable to use her own arms due to an injury, is sitting next to you — and she firmly advises you to turn the trolley. You take her advice. Let's stipulate that turning the trolley is in fact the right thing to do, and that Sophie grasps the reasons why it is the right thing to do.

Now, *who* saved the five people? You did. But also, the two of you together did. The two of you acted together, and you two together saved the five lives. There would be nothing infelicitous about a person whose life was thereby saved to thank *both* of you for saving her life.

Note how while it would make sense to say that the two of you together saved five lives, it would be a little odd to say that the two of you together turned the trolley, and it would be *very* odd to say that the two of you together moved your arm. 'Moving your arm' is, it seems, something *only* you do. But you-and-Sophie-together save five lives.

This reveals something about when it is and is not appropriate to attribute an action to a joint agent. In this example, you move your arm, and *thereby* steer the wheel, and *thereby* turn the trolley, and *thereby* save five lives. This pattern exemplifies what Anscombe characterized as "the A-D order" of intentional action: the various descriptions of what you do exhibit an instrumentally rational order. Each more basic description is performed for the reason that is described by a less basic description. The lesson seems to be that it is more difficult to correctly attribute the directed agent's more *basic* actions to a joint agent

than it is the less basic actions. The joint agent (you-and-Sophie) saved five lives, but only you moved your arm.

But *this* is a lesson a version of which we already know from studying the actions of more formally organized collective agents. For example, only individual legislators say ‘Aye’, but Congress *thereby* passes the resolution by means of their saying so. We cannot attribute the more basically described action(s) to the collective agent: *Congress* does not say ‘Aye’. But the less basically described action that happens in virtue of the individual legislators saying ‘Aye’ — viz., passing the resolution — can indeed be attributed to Congress. Just as the acts of collectives happen *through* the actions of individuals composing them, so too do the actions of joint agents happen through the more basic actions of individuals composing them.

Intuitive resistance to attributing the directed agent’s action to the joint agent, then, may stem from a misplaced focus on the more basic descriptions of the directed agent’s action. If instead we consider the less basic action description, the end for the sake of which the more basic actions are undertaken — in this case, *saving five lives* — then it will seem more natural to attribute *this* action to the joint agent comprising adviser and advisee. And this seems so no matter whether the adviser’s advice is explicitly *phrased* more basically (“Turn the wheel!”) or less basically (“Save the greater number!”). No matter which way an adviser happens to phrase her advice, it will be more apt to attribute to the joint agent the less basic description of the action (“Y’all saved five lives!”) rather than the more basic description of it (“Y’all turned the wheel!”)

When people work together to bring about some great good or bad, it will often be fitting to attribute this accomplishment to them collectively, and this is so even if some individual agents are playing an exclusively directive role. *They* saved the five people stuck on the tracks. (In Anscombe’s Truman case: *they* killed the civilians.) But there will always be many descriptions of actions performed by individuals that will not be attributable to the group.

Thus when you-and-Sophie together save five lives, *y’all* do what’s right.¹³ Furthermore, it seems that y’all do the right thing for the right reason. (Remember that I’m assuming that Sophie knows what’s what.) And so what y’all do has moral worth, just as Otto’s action above does, or as Sophie’s would if she alone had turned the trolley.

¹³ Unfortunately, formal English does not lexically distinguish between the singular and the plural second person pronoun, as many other languages do. Thus: *y’all*. Relatedly, *y’all’s* is the second person plural possessive adjective, c.f. *tus, tes, euer, vester*, ὑμέτερος.

Now recall that the moral worth objection to moral testimony says that when you act on the basis of trusted moral testimony, your action lacks moral worth, this because you are not doing the right thing for the right reason. After all, in such a case *you* don't grasp which reasons are the right reasons.

But we can finally see that moral advice is not so vulnerable to the moral worth objection. It is similarly true that the *individual* advisee does not do the right thing for the right reason, this because *he* does not grasp which reasons are the right reasons. But the joint agent comprising adviser and advisee both does the right thing and grasps which reasons are the right reasons. And so *their* action has moral worth. So, even if acting from moral *testimony* banishes moral worth from the scene, as the moral worth objection maintains, the same need not be true for those who act from moral *advice*.

Perhaps not all cases of full-blooded advice are like the trolley advice case, and I don't want to pretend that I have shown that there aren't also cases of advice that more closely resemble moral testimony. And I certainly haven't offered necessary and sufficient conditions for determining just when a case of trusting advice is a case of joint agency. I merely hope I here have opened up a line of inquiry.

I'll instead close with some undefended thoughts. When we think about the social nature of morality, we may be tempted to apply the tools of social epistemology straightforwardly to our topic. Social epistemology, like epistemology generally, concerns what to believe. And there *are* such things as moral beliefs. But morality is fundamentally practical in a way that most topics are not. A moral person is, in the first place, someone who acts well. Her believing the correct things about morality is important largely relative to this. When we consider how others ought to affect our own moral agency, perhaps we should focus instead upon how we can act *with* others, rather than (only) whether we should believe what others who seem to know report to us. Acting with others—acting together—is *massively* overlooked topic within normative moral philosophy. Why think that the only thing that matters morally is how *each* of us acts?

References

- Anscombe, G. E. M. 1981. Mr. Truman's Degree. The Collected Philosophical Papers of G. E. M. Anscombe, vol. III, Ethics, Religion and Politics. Oxford: Blackwell: 62-71.
- Cholbi, Michael. 2007. Moral expertise and the credentials problem. Ethical Theory and Moral Practice 10: 323-334.

- Driver, Julia. 2006. Autonomy and the asymmetry problem for moral expertise. *Philosophical Studies* 128: 619-644.
- Eldar, Shachar. 2010. Punishing Organized Crime Leaders for the Crimes of their Subordinates. *Criminal Law & Philosophy* 4.
- Gilbert, Margaret. 1990. Walking Together: A Paradigmatic Social Phenomenon. *Midwest Studies in Philosophy* 15: 1-14.
- Hallevey, Gabriel. 2012. *The matrix of derivative criminal liability*. Berlin: Springer-Verlag.
- Hills, Alison. 2009. Moral testimony and moral epistemology. *Ethics* 120: 94-127.
- Hinchman, Edward. 2005. Advising as inviting to trust. *Canadian Journal of Philosophy* 35: 355-386.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- McGrath, Sarah. 2011. Skepticism about Moral Expertise as a Puzzle for Moral Realism. *The Journal of Philosophy* 108: 111-137.
- Nickel, Philip. 2001. Moral testimony and its authority. *Ethical Theory and Moral Practice* 4: 253-266.
- Searle, John. 1976. A taxonomy of illocutionary acts. *Linguistic Agency University of Trier*: 355-368
- Stewart, James G. 2015. "Complicity" in *The Oxford Handbook of Criminal Law*, Oxford: Oxford University Press.