Fabian Schuppert

Comment on Hans Bernhard Schmid. Coordination, Cooperation and the Origin of Normative Expectations

Abstract: This comment suggests that coordination and cooperation are very different things, as the former simply is a device for problem-solving, while the latter relies on the existence of some shared intentionality. Similarly there exist different origins for the normative expectations an agent might form. Hence the comment argues that Schmid's taxonomy of action types, though helpful, needs to be extended and revised.

1. Introduction

In his article on 'The Idiocy of Strategic Reasoning' Hans Bernhard Schmid offers an extremely rich and detailed account of the origins of human cooperation. Schmid (2011, 52) argues for a taxonomy of social action types, differentiating sharply between mere strategic action, which is always directed towards an agent's own goals, and consensual action, which presupposes some form of shared intentionality. The principal claim of Schmid's argument (35) is that strategic reasoners are unable to engage even in minimally complex forms of cooperation, and that humans are only able to do so because they have different forms of practical reasoning at their disposal. This means that agents are only able to cooperate if they are able to see their shared goals and to form stable normative expectations regarding the actions of others and the environment. In other words, cooperation requires shared intentionality and agents have to be able to *count on each other* (rather than merely calculate with each other).

There is much to admire in Schmid's paper and while it's richness would warrant a whole range of comments, I will focus in this comment on two problems: the first concerns Schmid's example of the driver causing an accident (41), which Schmid intends to be a proof for his claim that strategic reasoners can not cooperate. While Schmid's critique of rational choice theory and the inadequacy of strategic reasoning for grounding cooperation seems accurate, I question the validity and appropriateness of his example, arguing that simple coordination problems call for a different example (*section 2*). The second issue tries to connect Schmid's argument more closely to the theme of the workshop, from which

the paper originated, namely the question of the origin of normative principles in forms of social cooperation. I will consider different forms of coordination and cooperation, trying to see how they map onto Schmid's taxonomy and whether this classification can help us in establishing under which circumstances norms of fairness and reciprocity attain moral relevance (*section 3*).[1]

2. Strategic Rationality, Equilibria and Coordination Problems

The example Schmid gives in order to prove that rational choice theory, which is based on the idea of strategic reasoning, fails to explain mutually beneficial coordination and cooperation is the case of a driver who with oncoming traffic on a standard two-way road decides to swerve into the lane of the oncoming vehicle, claiming that, not knowing what the other driver would do not only the action of keeping straight, but also the action of swerving were perfectly rational (41).

Schmid claims that the example points out a familiar problem of standard rational choice accounts, namely that if within a coordination game two equlibria exist for reaching a mutually beneficial outcome, with one being much better (e.g. easier, higher pay-off, etc.) than the other, in cases in which each agent is facing the dilemma of strategically interdependent expectations (i.e. if agent A assumes agent B chooses X, X is the rational choice for A, but if A assumes that B chooses Y, Y is the rational choice for A, and vice versa) rational choice theory cannot explain why choosing the better equilibrium would be more rational, or more logical, than defecting and going for the alternative choice (42). In the case of Schmid's example this means that the driver who swerved into oncoming traffic acted fully rational as long as he expected the other driver to swerve, too. This surely seems counter-intuitive to most of us, since keeping straight seemed for both drivers the best course of action.

According to Schmid, then, this example shows that even simplest coordination tasks are impossible to achieve if one adheres to individual strategic reasoning only. The reason for this is that under strategic interdependence of expectations agents cannot form stable and reliable expectations of each other's actions (43). Hence, Schmid argues, we can only solve the coordination dilemma if we see that both drivers share an intention (e.g. passing each other without accident) and form corresponding normative expectations concerning the other's behaviour, which means that we no longer speak of strategic actions but consensual ones.[2]

While, as suggested above, I fully agree with Schmid's overall argument, namely that some form of shared intentionality seems to be a necessary aspect of cooperation, there seem to be a problem with Schmid's example; it is neither

---

[1] It is important to note that Schmid does not distinguish adequately between coordination and cooperation, which seems to be an important difference, though.

[2] Schmid (47) defines, following Weber 1981, 186, consensual action as "action where agents have normative expectations concerning each other's behavior, without any previous explicit agreement between them".

clear that Schmid's example is analogous to the cases presented by philosophers like Bacharach (2006) and Sugden (1996) (meaning that it fails to describe a situation with two equally rational, but clearly qualitatively different strategy equlibria),[3] nor that Schmid's example actually presents a coordination problem. In order to show how simple coordination problems work and why shared intentionality is necessary for resolving coordination problems, I will thus propose an alternative example.

The problem with Schmid's example is twofold: firstly, in its description the example is already normatively loaded, which influences the way we should asses the agents' behaviour; secondly, the example does not in itself pose a coordination problem, which means that when the example is put into a game-theoretic metric each agent has actually three and not two options. Let me briefly explain both these problems before offering a more straightforward example.

The first problem is that Schmid explicitly assumes that both drivers are aware of traffic rules and that swerving might even be an action worth penalizing. Construing the example in this way, though, means that the driver who swerves actually has to violate a normatively valid rule, he is explicitly aware of, in order to even consider swerving a possible option. This surely seems odd, if the point is to claim that an agent who is a strategic reasoner has in the described example two strategy equilibria at his disposal. This is only true if we were to forget about the driver's normative commitment to the existing traffic rules. If one of the two actions is permissible according to the normative framework the driver adheres to and the other is not, it is simply implausible to claim that both strategies are equally rational. The way Schmid describes the example it is not analogous to the standard examples by Bacharach (see Fn. 4). The easy way out to avoid this problem is to describe the situation in a different way, so that the situation is not normatively loaded. However, even if we do so, we still face the second problem.

The second problem is simply that Schmid's example does not even present a coordination problem as such, since if neither driver acts the cars would just pass each other, as it is a two-way road and the cars do not approach each other head-on. In other words, Schmid's claim that the swerving driver had only two options, namely, to choose to go straight, or to swerve, is not quite accurate, since the driver might not even consider the situation to be one in which he needs to choose either option, since not-choosing would not result in a collision. This means, there are in fact three options. If two persons walk past each other on a sidewalk nobody would seriously claim that this necessitated on either side the decision to keep straight, since no coordination problem seemed to exist in the first place. There seems to be no good reasons that sitting in a car would change this.

However, as I mentioned earlier, the point Schmid actually tries to make, to wit, that coordination seems to require some form of shared intentionality is a

---

[3] Bacharach 2006, chapter 1, gives several examples for intuitively clear decisions for which actually strategy equlibria exist, for which rational choice theory has no satisfying answer, such as picking a card, running a single in cricket, or deciding together on where to go for lunch.

valid and very important one, so that I suggest a slightly different example to make this point salient.

Consider a single lane road, which is open to traffic from both directions, as you find them, for instance, on the West coast of Ireland. Chance has it that two cars approach each other on this road head-on, with both cars driving in the middle of the road, though if one really wanted to, that is, by driving onto the gravel and dirt next to the road, one would manage to squeeze two cars past each other.

In this situation, there is no best option since it makes no difference to which side driver A swerves, as long driver B swerves to the other side. It is, thus, no example for showing that strategic reasoners are unable to always identify the best course of action, but it is an example for showing that unless both drivers realize that they share the goal of making it past each other, that is, unless they can identify a shared goal instead of focusing on their own goal(s), they will not be able to make it past each other as quickly as possible. Only if both drivers can assume that the other driver is able to see the coordination problem and the shared intentional solution, will they be able to act on the basis of their expectations. Because there is no single best solution to this problem, this second example is also a paradigmatic case for standard coordination for mutual benefit, since in a wide range of coordination cases, it is neither clear nor necessarily predetermined which action will lead to the best possible outcome.

The single lane road example is a collective action problem, in which we do not have a group, or some form of strong bond between the two agents, but simply two individuals with straightforward goals (making it without collision and as quickly as possible to their final destination) who are forced to coordinate their actions. The form of shared intentionality necessary for solving these simple coordination problems is a very ephemeral and simple one, which is a point we have to keep in mind for the second problem in Schmid's text, which I want to focus on now.

3. Coordination, Cooperation and Normative Principles

The workshop at which Schmid originally presented his work focused on the idea of cooperation and the question of whether cooperative practices, such as working together, trigger a certain kind of (binding) normative principles, such as duties of fairness, or adequate treatment. This is a question that lies at the very heart of many theories of justice, which claim that cooperation, reciprocal interaction, and membership in an association can ground a whole range of normative concepts, including associative obligations, duties of fair play, or claims to support and assistance.

However, as existing debates in moral and political philosophy show, it is, according to many theories, not any kind of shared activity, or interaction, which triggers such obligations and principles. If we take a look at Schmid's taxonomy of social action types, then, the question is whether the distinction he makes can help us in figuring out what sort of practices (if any) generate norms of fairness,

and how we can distinguish between different kinds of relations. Unfortunately, though, Schmid's taxonomy seems of little help for our endeavour, as his taxonomy is based on communicative differences in action types, not conceptual ones.

According to Schmid, shared intentionality goes hand in hand with the existence of (more or less) warranted normative expectations. That is to say that Schmid's taxonomy does not differentiate between different kinds of normative expectations and different forms of shared intentionality, since the difference between the action types he classifies is the communicative form in which the normative expectations are expressed. For consensual actions, agents do not need to verbally communicate with each other (as in the single lane example above), since all that matters is that agents A's expectation that B also has a shared interest in passing each other is warranted and a reason for action for both drivers. Communal action, meanwhile, requires verbal communication, but no expressive agreement between the cooperating parties, since that is the defining feature of associational actions. In Schmid's taxonomy, then, it is the form of communicative commitment which defines the action type, a view which seems problematic when it comes to distinguishing between different forms of social interaction and their *normative content*. Schmid's taxonomy is seemingly primarily about assessing the degree of certainty (or reliability) of an agent's normative expectations, meaning that agent A has even better reason to count on agent B doing her share if agent B did not only say that she will do X, but also sign a paper saying that she will do X. While this might be true (at least in many everyday cases) this is an entirely different question from the one moral and political theorists try to answer, namely, if certain forms of cooperation come with more normative strings attached than others.

As Philip Pettit and others have pointed out, though, there seems to be a significant difference between solving collective action problems by way of coordination, simple short-term cooperation and full cooperation as a group agent. The problem with Schmid's taxonomy is twofold, then: firstly, the taxonomy does not differentiate between mere coordination and cooperation proper; secondly, the taxonomy also fails to distinguish between various kinds of cooperation and their respective normative content. Let me briefly explain both these problems with the help of a few simple examples.

The first example concerns the difference between coordination and cooperation. Take our earlier example of two cars meeting on a single lane road: here we have two agents with no shared goals (apart from passing each other without accident), no shared history, no shared identity, whatsoever. The two agents simply face a coordination problem as they have to find a way to pass each other on this narrow road. This means their shared goal is shared only in a very weak instrumental sense, since both have this goal simply because they cannot fulfil their original individual goals (e.g. for agent A to make it to place P, and for agent B to make it town T) without coordinating their actions at this particular place and time. Coordination in this example seems to involve little more than actually passing each other without accident.

Compare this simple coordination case with the case of two people baking a cake together. Even though baking a cake seems like a very trivial thing to do, this situation seems normatively much more loaded as several questions arise almost inevitably: who pays for the ingredients, how is the cake to be divided, and is the workload going to be shared equally? The answer to this question seems to depend on the particular case, although many philosophers seem to think that the standard distribution between cooperating parties should be equal, meaning that both the costs and the benefits of a cooperative project are shared equally by all cooperating parties.

While baking a cake might actually involve some form of explicit agreement between the cooperating partners (and thus might still fall into Schmid's scheme), there are forms of cooperation, which do not. Take the example of two people living in the same country: while the two people might not know each other and maybe even never meet, many philosophers, including Rawls (1971), would claim that a society is a form of cooperation and that this cooperative joint venture comes with a series of normative rights and obligations, even if a member never explicitly committed to living in that particular country. In other words, cooperation does not seem to necessarily depend on a particular form of communicative confirmation, but rather on the existence of certain shared properties (as in the case of two people living in the same country), or non-instrumentally shared goals (as in the case of baking a cake together).

The problem with cooperation is simply that it comes in many different shapes and forms, and in order to understand under which circumstances cooperation goes hand in hand with which normative principles we need to delve deeper into the issue than Schmid's taxonomy allows us to. In order to highlight the complexity of the relationship between cooperative practices, the reliability of normative expectations and the normative weight of different forms of cooperation, let us briefly compare of few instances of coordination and cooperation.

The convention of driving on a certain side of the road is a convention stemming from a mere coordination problem. The underlying reasoning is simple: if not all people drive according to the same rules, driving will be incredibly dangerous and unsafe. Hence, traffic rules and laws are created and a driving culture of driving on a particular side of the road established. This leads to a situation in which the actions of all participating drivers are neatly *coordinated*, so that every driver can have reliable normative expectations concerning the other drivers' behaviour on the road. However, there seems to be no real cooperation, no non-instrumentally shared goal, and only very little concern for principles and duties of reciprocity, fairness and equal respect. All participating agents are individual agents and there seems to be no collective agent (e.g. the drivers) to speak of.

The case of baking a cake together is different. Here we seem to have a form of voluntary, non-instrumentally motivated *cooperation*. Moreover, the cooperative practice probably is based on some form of expressive agreement, such as 'okay, let's bake a cake together'. However, the duration of the cooperation is very limited (since baking a cake normally does not take very long) and even though we think that norms of fairness exist it is not very likely that agent A would

morally blame agent B for eating a little bit more than half of the cake (i.e. the fruits of their cooperative joint venture) as long as A herself gets enough, too. The normative expectations of both agents, meanwhile, are warranted and stable, based on the knowledge that they share a common goal. However, it is doubtful that we can speak in this case of a collective agent (e.g. the cake-bakers) above and beyond the two agents A and B.

Other than in the first two cases, though, when we deal with a society, or a political party, or a government, all of which seem to be cases of complex multi-agent cooperation, we regularly speak of these entities as collective agents. Hence, it is the Swiss, the Greens, and the Government whom we hold responsible for this or that, even though we might have a hard time declaring which goals exactly the individual members of these entities share. This is a very important point, especially since one might even be able to claim that the normative expectations the individual members have of each other are less stable and warranted than in society-wide coordination cases (our first example) and voluntary small-scale cooperation (our second example). This poses an interesting challenge for Schmid's taxonomy, as on the one hand, then, societies and governments seem to be special, complex kinds of uniquely human cooperative enterprises, while on the other hand the normative expectations the participating individual agents have seem much more likely to be disappointed. However, the stability of the agents' normative expectations has no bearing on the cooperation's normative bindingness, since societies seem to include rather stringent norms of fairness and justice.

These three cases, thus, show that while Schmid is perfectly right in claiming that coordination and cooperation both require some form of shared intentionality, his taxonomy does not allow us to adequately distinguish between the varying kinds of coordinative and cooperative practices human societies generate. That is to say that for moral and political philosophers the crucial question of which normative principles human cooperation triggers remains unanswered. The given taxonomy offers us a viable starting point, but it cannot provide any criteria for distinguishing between cases in which two agents share only a very trivial goal for a short period of time (as in the case of a stranger helping an agent to carry a pram up a flight of stairs) and cases in which a multitude of agents cooperate over a long period of time in order to secure for each other a good and meaningful life (as in the case of a collective, or commune). Hence, Schmid's taxonomy unfortunately cannot help us in determining which forms of cooperative practices come with greater normative weight than others, and why.

Bibliography

Bacharach, M. (2006), *Beyond Individual Choice. Teams and Frames in Game Theory*, Princeton

Rawls, J. (1971), *A Theory of Justice*, Cambridge/MA

Schmid, H. B. (2011), The Idiocy of Strategic Reasoning, in: *this issue*

Sugden, R. (1996), Rational Coordination, in: Farina, F./F. Hahn/S. Vanucci (eds.), *Ethics, Rationality and Economic Behavior*, Oxford, 244–262

Weber, M. (1981), Some Categories of Interpretive Sociology (translated by E. Graber), in: *The Sociological Quarterly* 22(2), 151–180