

Sugden's team preferences

An inquiry into the nature of team-directed reasoning

Jakob Schönhuber

1. Introduction

What is the relation between the *I* and the *we* in social interactions? As a famous saying goes, 'there is no *I* in team'. Successful football players, for instance, are those that know how to cooperate and to act in a way that is beneficial to their team. Similarly, in everyday life, we often coordinate our behaviour without any problems. For example, when we enter a bus while someone else must leave, we usually manage the space smoothly.

Cases of coordination and cooperation, like the ones mentioned above, pose explanatory challenges. Given that it is always individuals who make the decisions in these cases, how can the abovementioned examples be explained? Cooperation problems demand that one forgoes individual incentives to defect in order for a group goal to materialise. We can imagine the football player mentioned above to be in a situation where he can try to score or to pass the ball on to another player who then might score. He knows that chances for his team to score are much higher if he cooperates. Yet, individually, the player has an incentive to deviate. For him, individually, the expected outcome is much higher if he tries to score. This is because if he makes the goal, he will get all the recognition. In case of failure, however, the player might still be able to blame bad luck. In situations like this, how can we explain cooperative behaviour, which we arguably still observe (Sugden, 2000)?

Similarly, coordination problems invoke some explanatory challenges. Even if there are no individual incentives to deviate from coordination, different ways to reach the same outcome might still exist. Imagine, for instance, when you have to enter a bus while another person needs to leave. Which side of the door should you occupy? Should you wait or just enter? For a smooth entrance, it is crucial that both of you react in the same way. For example, you take the right-hand side of the door, while the other person takes the left-hand side. But it is equally permissible that you choose the left-hand side and the other person the right-hand side. Note that there is nothing inherent in the decision situation which tells an individual what to do (Grüne-Yanoff and Lehtinen, 2010). Yet, our everyday experience shows that we manage to solve these problems without particular difficulties and without talking to one another.

In this paper, I shall address one specific account which tries to offer a solution to these types of problems: Sugden's (2000) theory of 'team reasoning'. Sugden proposes to think of coordination and cooperation problems as challenges that a group of people must solve *together*. Therefore, a person does not think about what she should do *individually*, constrained by how the other agents might act. But a person considers what she can do *together* with others, as a group. For a group to exist, all agents must acknowledge is that they must solve a problem together with others and they have to believe that the other agents engage in this perspective too. If this is in place, Sugden's team reasoning straightforwardly solves cooperation and coordination problems because agents can act on a group perspective.

Sugden (2000) calls thinking in terms of the group perspective 'team-directed reasoning'. The term 'directed' in team-directed reasoning hints at the special perspective Sugden takes. For Sugden, it is always individuals who engage in reasoning. Team-directed reasoning identifies what is best for the group, all things considered. This can also be called the group- or team goal. Team-directed reasoning

thus suggests the strategy that achieves the team goal. For example, in a football game, the goal is to win the match. So, everyone acting in line with the group goal should choose the strategy that maximises the team's chances to win. To achieve this, Sugden maintains that team-directed reasoning is a psychological process that individuals engage in for themselves. Furthermore, it can be maintained that team-directed reasoning is a methodologically individualistic account because it is reducible to individuals engaging in that perspective. Hence, an agent's decision to take the group perspective does not depend on what the other agents do. This allows to represent team-directed reasoning by rational choice language³⁴. Team preferences are the preferences an agent identifies by engaging in team-directed reasoning.

I shall argue that (a), Sugden's (2000) theory of team reasoning does not sufficiently account for individual, self-interested temptations that potentially undermine an agent's ability to act on team preferences. I will give an example of cooperation where all of Sugden's conditions are fulfilled. That is, agents engage in team-directed reasoning and take themselves to be members of a team. I will show that in such a situation it is still perfectly legitimate for an agent to act on their self-interest. (b) based on (a), I propose that Sugden's theory of team-directed reasoning should be interpreted as an agent recognising what, all things considered, would be best for the group. However, I will also point out that this mode of thinking is not exclusive. I will show that it is legitimate to assume that an agent has access to both their individual- and their group preferences in a decision situation. Based on this characterisation, I will argue that Sugden's theory does not solve cooperation- and coordination problems. This is because if it is acceptable that agents engage in both team-directed reasoning and individual reasoning, and that the presence of a team in itself does not give reason to act according to the team goal, it is not clear why following team preferences should solve coordination- and cooperation problems. The point that I will exploit is that there is a difference between engaging in team-directed reasoning and acting accordingly. To act in line with the team goal, I will argue, it is necessary that every member can trust in the other agents acting according to team-directed reasoning. But since that is not given, I conclude that team-directed reasoning does not offer a solution to coordination and cooperation problems.

I shall proceed as follows: In Section (1) I present a problem of coordination and how Sugden's theory of team-directed reasoning would address it. I continue by arguing in favour of claim (a) in Section 2. That is, I will argue that Sugden does not sufficiently address the conflict between individual and group interests. Section 3 defends claim (b) and argues that team-directed reasoning does not solve coordination- and cooperation problems. Section 4 concludes.

³⁴ The tools of rational choice theory are preferences. Preferences are tuples comparisons of alternative states of the world. If these tuples fulfil the conditions of transitivity and completeness (that is, they include everything that matters for an agent in a particular state of the world), they can be represented numerically as a function. This function is called utility function (Savage, 1971; Bradley, 2014). In the following I will also use the word payoff function to describe the same concept. A second, important point is that rational choice theory can be seen as an idealisation of reasoning (Risjord, 2014, 106–107). Therefore, whenever I use the terms preferences, or reasoning/ thinking interchangeably, I refer to the concept of idealisation. This holds for both team-directed reasoning and team preferences as well as individual reasoning and individual preferences.

2. Coordination and Sugden's theory of team-directed reasoning

In this section, I present Sugden's (2000) theory of team reasoning. I begin with a coordination problem and then explain how Sugden's theory would solve the problem by employing team-directed reasoning.

At my university, every morning, many students need to go to class. The best way to reach the classrooms is via an escalator which directly goes to the third floor. Some of the students are in a hurry to arrive on time and need to move fast. I denote them as players of type- f (). Other students have more time and walk slowly. These students are identified as players of type- s (). Students of both types need to coordinate the space on the escalators. Unfortunately, talking to one another is not an option since all students are listening to music on their headphones³⁵. Therefore, the challenge that the students face is to interact with one another and to respond in the best possible way to players of the other type. Consider now Figure 1³⁶. For simplicity, all players are summarised under their types, that is, (p_f) and (p_s).

		P_s	
		Right	Left
P_f	Left	11, 11	0, 0
	Right	0, 0	10, 10

Figure 1: Coordination on escalators under strategic reasoning

If both do not coordinate, that is, use the same or both sides of the escalator at the same time, they end up with utility 0. Note that the exact number of utilities an agent gets if coordination fails does not matter here. What is important is that both players are considerably worse off compared to both coordinative equilibria. In the case of those students who are in a hurry, it is straightforward why they have a low utility if coordination fails: they will miss their class. But those students who have more time will be worse off as well. It arguably bothers one if other people constantly want to pass by and complain. Therefore, it would be best for both to coordinate. Hence, players of type- s would need to pick one side of the escalator and players of type- f the other side. In principle, which side is used by which players does not matter. However, we can imagine that the agents are used to walk on the right-hand side and to surpass someone on their left-hand side. Assuming that agents have a preference to coordinate according to this convention rather than the opposite way, it can be stated that if use the left-hand side and if use the right-hand side of the escalator, they achieve the best possible outcome to this problem.

³⁵ Unlike this case, in many coordination games it is not possible for players to communicate, but they still have to coordinate. Therefore, referring to language is not a general solution.

³⁶ Standard games can be represented in this form: The rows represent the strategies of while the columns represent the strategies of . Moreover, the first number in the matrix represents the payoff of the first player, while the second number the payoff of the second player (Grüne-Yanoff and Lehtinen, 2010).

However, analysing the above situation from a game-theoretic perspective, this game has no unique solution. As Sugden (2000) points out, payoff dominance – that is, choosing the strategy that leads to the overall best outcome in terms of payoffs – gives strategic players no apparent reason to choose a particular strategy. This is because rational players choose their strategy taking into consideration their expected outcome³⁷ conditional on what they think the other player will choose. Thus, both will choose [left] when they expect the other player to choose [right] and vice versa. Since no player knows what the other player will choose, their thinking will be a regress. Thus, there is no unique solution to this coordination problem.

This particular result seems counterintuitive. In fact, at my university I can observe students that coordinate on the optimal equilibrium. Similarly, our experience of everyday life tells us that we can solve these problems without any particular difficulties. Sugden's (2000, 183) account of team-directed reasoning builds on this tension between our experience and what strategic game theory would suggest and proposes a solution. Sugden suggests changing the perspective. He argues that instead of asking 'What should *I* do?', the question 'What should *we* do?' offers a way out. In other words, Sugden suggests conceiving of these problems differently. Agents should not be seen as acting in accordance with their individual utility functions, but rather as being able to take the group perspective into account. That way, rational agents can overcome coordination failures and choose what is best for the team.

From a conceptual perspective, it works exactly the same way, whether agents engage in team-directed reasoning or in individual self-interested reasoning³⁸. The difference is what the players choose to care about. In team-directed reasoning, it is the team's utility function, and in individual reasoning, it is their individual utility function. Thus, team-directed reasoning can formally be characterised as a preference order that is in line with the group goal. Hence, team-directed reasoning with respect to some game of form G generates team preferences which can be represented by a utility function $t(\cdot)$ (Sugden, 2000).

It is important to see that Sugden (2000) characterises team-directed reasoning as not being conditional on what other members of the team think. Team-directed reasoning is an individual account of reasoning directed towards the goal of the group. And the goal of the group is simply taken from the problem setup. To see how team-directed reasoning is able to solve coordination problems, consider Figure 2. Let me assume for a moment that this is a situation where team-directed reasoning applies. In the following, I will explain why this example can be represented as such.

³⁷ The outcome as characterized by preferences, conditional on what the other player chooses.

³⁸ To clarify, individual preferences might be other regarding as well. It is perfectly fine for an individual, in their 'self-interested' preferences to care about other people. The difference between individual and group preferences is that the group preferences are identified by thinking about specific coordination and cooperation problems which a 'purely' individualistic perspective is not able to overcome (Sugden, 2000).

		P_s	
		Right	Left
P_f	Left	11	0
	Right	0	10

Figure 2: Coordination on escalators under team directed reasoning

Figure 2 represents coordination on an escalator according to Sugden's (2000) account of team-directed reasoning. The difference compared to Figure 1 is that an agent who applies team-directed reasoning considers the group outcomes rather than their individual preferences. Thus, the strategic element of game theory, where all agents choose their strategy under consideration of how the other agents might act, is not present anymore: An agent who employs team-directed reasoning can straightforwardly choose according to team preferences $t(\cdot)$. Hence, all players employing team-directed reasoning and belonging to type- f will take the left-hand side of the escalator, players belonging to type- s will act respectively. Consequently, team-directed reasoning induces coordination on [left, right].

To guarantee that agents employ this type of thinking, they have to take themselves as members of a team. Sugden (2000, 195) argues that for this to be in place, team confidence is needed. Team confidence is the confidence of an individual agent in a system of first and higher-order beliefs. That is beliefs that the other members engage in team-directed reasoning too. Further, beliefs that the other members believe, that the individual themselves follows team-directed reasoning. In the following, my argument will employ that there might be a difference between team-directed reasoning and Sugden's conditions being in place and acting upon it. I will highlight that Sugden never rules out individual interests and temptations completely and that this poses a challenge to his account.

Confidence in these beliefs is an empirical concept (Sugden, 2000, 196). This means confidence is based on an agent's observations about how other agents behave. If one sees that other agents coordinate in accordance with team-directed reasoning, this suffices to induce confidence in the team's existence. But this does not rule out stronger forms of confidence. Mutual agreement, for instance, would be such a stronger form of team confidence.

Finally, for a team to exist, all members of the team (with) have to engage in team-directed reasoning and have to take themselves to be members of a team as well. If these two conditions are fulfilled, team agency is in place (Sugden, 2000).

Concerning the escalator example, these points clarify why team reasoning applies. Sugden (2000) never states that a group identity is necessary, that agents need to know one another, or that someone ought to act according to team preferences. For team agency to be in place, it suffices that everyone, individually, observes coordination in line with team-directed reasoning. Moreover, every agent is aware of the coordination problem and knows how to solve it: by engaging in team-directed reasoning.

To conclude this section, it is worth to pause for a moment and appreciate what Sugden (2000) has achieved. Sugden suggested to think of coordination problems, like the one on the escalator, as challenges that a group of people has to overcome together. This reframing of decision problems solves a challenge to game theory. Game theory is not able to explain how we, in our everyday lives, manage to coordinate. Because, by simply maximising one's own utility function under consideration of what other agents might do, many interactions do not have straightforward solutions. By assuming team-directed reasoning, however, a theoretical answer is given to these problems³⁹.

In the next section, I will point out that Sugden (2000) might be too confident in his account. In fact, I will show that there are legitimate cases where self-interest overrides team preferences even though Sugden's conditions of team agency are in place.

3. The challenge from self-interest

Let me now continue with an example which fulfils all of Sugden's (2000) conditions of team agency where, however, behaviour in accordance with the team's goal is not given. Based on this example I will then move on to argue in favour of claim (a), namely that team reasoning as it is understood by Sugden is not able to accommodate the conflict between individual and group interests.

Consider again the case where students must coordinate on escalators. Given Sugden's (2000) conditions it can easily be explained why this is a case where team agency is in place. As highlighted above, all agents face a coordination problem about the space on the escalator and this coordination problem cannot be solved by strategic thinking. But there is a solution which is easily visible for everyone. Namely, for all players $\{$, if you are a player of type- f , choose [left] and if you are a player of type- s , your strategy is [right]. Confidence in this type of coordination is given because every agent observes that the other players are acting accordingly. Thus, team agency is in place. To make this step clearer, consider that every agent faces a choice between [left] and [right]. Note that this is a situation where the agent must interact with other players. By observing others, the system of first and higher-order beliefs emerges. Every player believes that the other players engage in team-directed reasoning too. Thus, coordination should take place in accordance with team preferences.

This is the general setup. But now consider one of the students, Ann, in detail. Let me assume that Ann is aware of the common goal of all $\{$, on the escalator to facilitate movement. Furthermore, by engaging in team-directed reasoning, Ann identifies [left, right] as a solution to the coordination problem. Also, Ann knows that she must coordinate. Therefore, she takes herself to be a member of $\{$. That is, she confidently believes that all other players engage in team-directed reasoning and that they expect her to do so too. Therefore, Ann behaves according to team-directed reasoning, at least most of the time. Sometimes Ann uses the escalator while chatting with John. In such cases, Ann, even though a player of type- s , prefers to stand next to John on the left-hand side of the escalator. Note, Ann is aware of the team goal in that case too. But that does not stop her from standing next to John. This is because Ann values her self-interest at that moment more than acting on team preferences. Note that regarding

³⁹ With regards to cooperation problems, Sugden (2000) claims that his account works analogously. Team-directed reasoning not only identifies solutions for coordination problems, but also for cooperation problems.

Sugden's conditions of team agency nothing changes: Ann still engages in team-directed reasoning and is confident that she is part of the group $\{, \}$. Nevertheless, Ann chooses to stand on the left-hand side of the escalator. That is, Ann chooses to put her individual self-interest above the group preferences.

In other words, there is a difference between engaging in team-directed reasoning and acting according to it. This example highlights that Ann recognises both, the team's preferences and her individual interests. To clarify, the case where Ann talks to John is not the same as standard coordination on the escalator. This example transformed the payoffs of individual agents so that it became a cooperation problem. But, as emphasised, Sugden (2000) claims that his account is able to solve both types of problems. Therefore, in terms of intuitive plausibility, nothing should be objected. Ann engages in a situation which fulfils all of Sugden's conditions, but acts on her self-interest. The challenge with team agency lies in the fact that it does not necessarily give an agent reasons to act in accordance with it.

Sugden (2000) might respond to this critique by arguing that it genuinely misconceives the nature of team-directed reasoning. Sugden (2000, 182– 183) writes:

The idea is that, in relation to a specific decision problem, an individual may conceive of herself as a member of a group or team, and conceive of the decision problem, not as a problem for her but as a problem for the team [...] For someone who has framed a decision problem in this way, relevant advice about what to do has to be addressed, not to the individual alone, or even to each member of the team independently, but to each member of the team as a member of the team.

In other words, Sugden's (2000) theory builds on the idea that team-directed reasoning is about solving problems *together*. Therefore, is the idea that individual self-interest can play a role simply incommensurable. Consequently, individual preferences or strategic considerations do not play a role in team-directed reasoning, they are something different.

From the methodological perspective, this interpretation of Sugden's (2000) account implies that if an agent engages in team-directed reasoning and takes herself to be a member of a team, that agent only considers the team's utility function $t(\cdot)$ and not individual preferences. Thus, an agent will act according to $t(\cdot)$ in virtue of being a member of a team.

I consider this interpretation of team agency to be implausible. Being able to identify what *we* should do as a group, as distinct from what *I* should do individually, does not mean that the latter is ignored. An important reason for this is that group confidence is an empirical concept. Simply observing that other agents act in accordance with a group goal is unlikely to induce confidence in team agency which is strong enough to ignore self-interest (Hindriks, 2012). But in some cases, team agency might be stronger than being based on empirical observation. Soldiers, for instance, are trained to commit to their orders and their company even under risk of their life. But soldiers are trained to do so and know that they can trust their companions. Yet, even the military knows soldiers who desert. Arguably, many of them are afraid for their lives when they decide to disregard orders and not to act on the team goal.

My point here is not to argue that it is impossible that agents will act on team preferences rather than their individual ones. What I aim to show here is that it is plausible to assume that even if team agency is in place, individual interests might come first, which means that agents act on their individual preferences, rather than on the team preferences.

This point can be stated more strongly. It is not clear why someone who chooses to engage in team-directed reasoning should ignore their self-interest completely. Being able to identify what *we* should do

as a group, as distinct from what *I* should do individually, and choosing that the *we* perspectives matters does not mean that individual considerations will be ignored. Since Sugden (2000) takes his theory to be broadly applicable, I contend that it is legitimate to maintain that agents might sometimes act on team preferences and sometimes not. Note that I am not making any claims about how an agent chooses between individual or group interests. I believe that this is a crucial question too because it is not clear from a rational choice perspective which goal should be taken as prior. My argument is that Sugden simply ignores the possibility of self-interested motivations in cases of team-directed reasoning. And this, I argue, poses a challenge for his account to solve coordination and cooperation problems.

Let me restate my argument so far. Sugden's (2000) theory of team reasoning fails to account for Ann's decision not to coordinate in line with the group goal. The problem is that Sugden does not sufficiently accommodate for challenges between individual and group interests. In the next section, I will defend claim (b). Claim (b) states that it is important to account for this conflict between individual and group interests in situations where team-directed reasoning applies. But if this conflict is accounted for, Sugden's theory of team reasoning fails to offer a solution to coordination problems.

4. An alternative interpretation of team-directed reasoning

As emphasised above, agents acting in situations which are characterised as cooperation problems can face a decision between their self-interest and the interest of their group. In this section, I will do two things. First, I will show that it is acceptable to interpret Sugden's (2000) account in a way which accommodates this tension. I propose to think of an agent in decision situations where team-directed reasoning applies, as having access to both the group as well as their individual preferences. Second, accepting this interpretation of Sugden's theory of team reasoning, I will point out that team-directed reasoning does not offer a solution to coordination and cooperation problems anymore.

Team-directed reasoning can be understood as an individual recognising what, all things considered, would be best for the group, distinct from the question of what would be best for oneself. With this, I mean that both individual- and team-directed reasoning can exist in the same coordination problem and that it is legitimate to assume that an individual is aware of both. Thus, the question an agent must face is upon which preferences they should act.

Sugden characterises team-directed reasoning as a mode of thinking which recognises what, all things considered, one should do when acting as a member of a team. Sugden (2000, 195) writes:

My analysis prescribes what he [an agent] should rationally choose if he takes himself to be acting as a member of the team - that is, if he engages in team-directed reasoning.

But note that Sugden (2000) never rules out that an agent cannot consider their self-interest too. In fact, by characterising team-directed reasoning as a mode of reasoning and team agency by a series of conditions – which boils down to understanding what the group goal is as well as observing other agents engage in behaviour in accordance with the group preferences – Sugden does not rule out that individuals can also think what they, individually, prefer and what is best for themselves.

Maintaining that individuals can engage in both team-directed reasoning as well as self-interested considerations does not rule out that coordination can happen in line with the group goals. Though, what follows from this claim is that every agent must choose whether they want to act on their team preferences or their individual ones. Consider, for example, what it means to live together in a household.

Be it as a family or a shared flat, keeping the common spaces clean is a collective endeavour which, arguably, everyone agrees upon. But this does not rule out that, when it comes to one individual to clean, a person has a series of incentives to clean only superficially or even not at all. I think this example nicely points at the inherent tension between individual and group interests. This tension does not have to be based on a conscious reasoning process. What I take it to be is a primitive notion of an agent recognising what is best for oneself, as opposed to what is best for the team. Now, how an agent chooses in an actual situation is a different question and to answer it would require detailed accounts of cognitive decision making. But for me, to make my point, this is not necessary. I simply maintain that agents are, generally speaking, aware of this tension between individual and group interests. Agents are aware of this tension not only for themselves, but they also know that other agents face similar conflicts.

Interpreting team-directed reasoning as a mode of thinking that exists alongside self-interested considerations seems the plausible working hypothesis for rational agents. However, I will show that this interpretation ultimately undermines what team-directed reasoning tries to establish in the first place: A solution to coordination problems.

To see why, consider once again coordination at the escalator. This time we need to look at Carol who approaches the escalator in a hurry. Carol sees Ann and John in front of her approaching the escalator too. By engaging in team-directed reasoning Carol knows that for all $\{$, the following strategy-set applies: If you are a $\}$, choose [left], if you are a $\}$, choose [right]. This would be the optimal solution, which is also followed by other players, as Carol observes. All $\{$, on the escalator act in line with team preferences. Carol is also confident that Ann and John individually coordinate in line with team-directed reasoning. So far, Carol saw them every week taking the same escalator. However, in these cases, Carol observed Ann and John only individually, instead of walking together. Thus, based on Sugden's condition, Carol can assume that team agency is in place. But should Carol act on the team preferences? Carol knows based on team-directed reasoning what she should do, but she is also aware that individuals can act on their self-interest. Thus, observing Ann chatting with John in front of her causes Carol to seriously doubt whether Ann and John will follow the team preferences. Therefore, Carol identifies two possible strategies: She can either try to coordinate on team-preferences. In that case, Carol would take the left-hand side of the escalator and hope that the space in front of her is empty so that she can easily walk on the escalator. Or, Carol could also try to push forward on the right-hand side of the line to get to the escalator before Ann and John arrive. Carol must choose her strategy depending on what she thinks that Ann will do. If Ann acts on her self-interest and stands next to John on the escalator, Carol should defect too to be on time for class. However, if Ann would coordinate on team preferences, for Carol that would be the best choice as well.

Clearly, this is a case of strategic interaction as game theory models it. Team-directed reasoning interpreted as a mode of reasoning which exists alongside individual, self-interested reasoning does not offer a solution anymore. Note that what this example establishes is not that Ann's choice between individual or group interests in itself undermines the possibility of coordination and cooperation. It is the possibility that Ann can act on either mode which challenges Sugden's theory of team reasoning. Carol, who simply wants to coordinate, must take this point into consideration and her confidence in the team is undermined. Generally speaking, the problem is that confidence in other agents engaging in team-directed reasoning does by no means imply confidence that other agents act accordingly. But to trust in team-preferences, confidence that other agents act according to team-directed reasoning is necessary.

5. Conclusion

In this paper, I presented Sugden's theory of team reasoning. I argued that (A), Sugden's theory is not able to sufficiently account for challenges between individual self-interest and team goals. The point I made is that there is a difference between saying that individuals engage in team-directed reasoning and have to solve a coordination or cooperation problem together, and claiming that, because of engaging in team-directed reasoning, self-interest will be ignored. (B) Exploiting this argument, I proposed to think of team-directed reasoning as an agent being aware of the group goal, but that this way of thinking does not lead one to ignore their self-interest. However, I showed that on this account, team-directed reasoning does not solve coordination problems anymore. Hence, I conclude that team-directed reasoning does not offer a solution to coordination and cooperation problems.

Acknowledgements

I am grateful to Dr. Conrad Heilmann for supporting my submission with helpful comments and suggestions on earlier versions of this essay. Moreover would I also like to thank the anonymous referees of this journal. Their relentless effort helped me to significantly improve the quality of this essay.

Jakob Schönhuber obtained a B.Sc. in Economics from the Vienna University of Business and Economics. Currently is Jakob doing the Research Master's in Philosophy and Economics at Erasmus University Rotterdam. He is also working as a Research Assistant at the faculty of Philosophy for Dr. Conrad Heilmann. Jakob's research interests concern questions of rationality and cooperation in the context of game theory.

References

- Bicchieri, C. (2006). *The Grammar of Society. The Nature and Dynamics of Social Norms*. Cambridge University Press. Cambridge.
- Bradley, R. (2014). *Decision Theory: A Formal Philosophical Introduction*. In: Hansson, S. O. and Hendricks, V. F. (eds.) *Introduction to Formal Philosophy*. Springer Undergraduate Texts in Philosophy pp. 611 – 656.
- Colman, A. M. (2003). *Cooperation, psychological game theory, and limitations of rationality in social interaction*. Behavioral and Brain Sciences. 26. pp. 139 -198.
- Gilbert, M. (2008). *Social Convention Revisited*. Springer Science+Business Media B.V. DOI 10.1007/s11245-008-9032-5. pp.5–16.
- Gold, N. & Sugden, R. (2007). *Collective Intentions and Team Agency*. The Journal of Philosophy. 3. pp. 109 – 137.
- Grüne-Yanoff, T. & Lehtinen, A. (2010). *Philosophy of Game Theory*. In: Mäki, U. (ed) (2010). *Handbook of the Philosophy of Science*. Volume 13: Philosophy of Economics. Elsevier BV.
- Hindriks, F. (2012). *Team reasoning and group identification*. *Rationality and Society*. Sagepub. DOI: 10.1177/1043463111429274. pp. 198-220.
- Risjord, M. (2014). *Philosophy of Social Science. A Contemporary Introduction*. Routledge. New York.
- Savage, L. (1972). *The Foundations of Statistics*. Wiley. Second revised edition.
- Sugden, R. (2000). *Team Preferences*. *Economics and Philosophy*. Cambridge University Press. pp.175-204.