NATURALISING THE GROUP MIND:

THE COGNITIVE LIFE OF SMALL GROUPS AND TEAMS

Kellie Williamson

For the Degree of Master of Philosophy

Department of Cognitive Science

Macquarie University

Submitted 29th February 2012

Author Statement

This work has been produced for the degree of Master of Philosophy, at Macquarie University. It has not been submitted to any other university or institution for a higher degree. The content of this thesis is wholly my own contribution, unless indicated otherwise, in which case references and explicit acknowledgements are provided.

_____

Kellie Williamson

Acknowledgements

My immense thanks and appreciation to Professor John Sutton for his continual insight, feedback and encouragement during the creation of this thesis. His enthusiasm for and expertise in truly interdisciplinary research meant the whole process was hugely exciting and intellectually rewarding. Thanks to Dr Rochelle Cox for her patience and careful, astute instruction in how to design novel empirical research, how to make it interesting and how to keep it scientific. Thanks to Gareth Hipwell for his sharp editorial eye, and to Dr Nicolas Bullot for valuable feedback on earlier drafts. Finally, thanks to Associate Professor Amanda Barnier and the Collective Memory Team in the Department of Cognitive Science at Macquarie University, for theird helpful and encouraging feedback on this research.

Table of Contents

Abstract

Groups of people can do extraordinary things. A team of chefs is able to satisfy hundreds of guests, night after night, impressing them with the quality of that which they create and the efficiency and professionalism with which it is created. A Premier League football team astounds fans across the world with captivating on-field performances. Teams of paramedics operate under some of the most extreme temporal and emotional pressures confronted by humans, and are able to do so successfully. While there are countless examples of high-performing groups of people, there are likely just as many unsuccessful groups; not all football teams dazzle spectators, especially in amateur competitions throughout the world, and even at the elite level the best groups can have an 'off day.' Just as there are differences in abilities, skills and performance between individuals, there are differences between groups. Groups of people vary not only in the extent or continuity of their successes, but also across numerous other dimensions. Some groups are much less formal than professional teams of chefs or athletes. These less formal groups are the everyday groups we find ourselves in: book clubs, parent groups, craft groups, temporary committees, families and so on. What all of these groups, both formal and informal, have in common, is that on many occasions they involve multiple people undertaking cognitively demanding tasks together. They're remembering, planning, acting and deciding together. The cognitions and actions of each group member are coordinated such that the group is able to attempt, and hopefully complete, the relevant task and, if all goes well, do so successfully. From a scientific and philosophical perspective, the existence of such groups raises a number of interesting questions. In particular, what, if anything, makes genuine group behaviour different from aggregated individuals' behaviour, such as the behaviour of a mob or herd? And, particularly compelling, what are the processes or factors that enable multiple individuals to act as a group, especially successfully?

One well-known way of characterising group behaviour, in both the social sciences and in the popular media, is to describe this behaviour in terms of a 'group mind'. A specific

theory of the group mind is hard to pin down, but it has been associated with the surrendering of individuals' own intentionality to the group's will, a problematic metaphysics of downward causation and 'crowd' or 'herd' mentality - not the kind of skilful group performance one expects from the kinds of groups described above. Moving away from this characterisation of group behaviour, in this thesis I construct a newly robust theory of group-mind-cognition that avoids the problematic aspects typically associated with group mind claims. The motivation for this is straightforward. If we are to understand and explain how cognition unfolds in our everyday lives, an acceptable and shared goal among the cognitive sciences, then we need to understand the nature of groups, given their prevalence in our everyday life. I argue that the most useful way to do this is to characterise groups of the kind described above, as cognitive systems or information processors in their own right, under the right conditions. Throughout this thesis, a general account of group cognition is developed, which is applicable to a variety of groups across a variety of domains.

Central to this account is the notion of 'cognitive interdependence' between the members of the group. Members are cognitively interdependent if they mutually influence one another's cognition, not merely by way of a stimulus and response relation, but through transforming one another's cognition and affecting one another's cognitive processing. By highlighting cognitive interdependence, two important steps can be made: firstly, the metaphysical claim that group cognition is a real phenomenon, subjectable to scientific and philosophical inquiry, is secured, and secondly, it helps to identify the kinds of factors and processes that facilitate successful group performance. Through identifying different forms of cognitive interdependence, this account emphasises the ways that cognitive processes can be distributed across interdependent group members. So as to generalise across a variety of different groups performing many different tasks, this account also explores the dynamic nature of cognitive interactions between group members, whereby members mutually influence one anothers' cognitions and actions, seemingly automatically, thereby enabling

swift, improvisatory responses by the group to novel, unpredictable situations. On this view, group cognition is characterised as both distributed and dynamic.

As an exemplar of group cognition, I explore the various ways in which members of a sports team are cognitively interdependent. By analysing sports team performance it is possible to extract some general principles that apply to groups in other domains, particularly the ways in which multiple individuals can become cognitively interdependent. This account of group cognition seeks to overcome the silences or gaps between several research fields, empirical and conceptual, each of which is independently concerned with aspects of group behaviour, social cognition and skilled action. In particular, the account integrates work from social ontology in philosophy, cognitive psychology, sports psychology, philosophy of science and organisational psychology, creating a conceptually hybrid account but also a hybrid methodology for investigating, describing, explaining and ultimately understanding the mechanisms of group cognition.

To begin, Chapter One introduces the areas of research that are relevant to explaining group cognition, but which have not previously been integrated. The points of contact between these formerly separate areas of research are sketched, foreshadowing future opportunities for fruitful research on group cognition. As the study of group cognition *per se* is in its infancy, some of the basic, foundational conceptual work still needs to be done, so Chapter One also lays out the terminological terrain, explaining the key concepts as informed by the theoretical claims developed across subsequent chapters.

Chapter Two is a metaphysical account of group cognition as a real phenomenon. This chapter lays out a plausible metaphysics for group cognition, informed by work on mechanism and emergence in the philosophy of science. This metaphysics underpins the account of group cognition developed across the thesis, and is used in this chapter, in conjunction with select work from social ontology, to distinguish between cases of genuine group cognition and aggregated individual cognition. From the integration of this work, the

notion of cognitive interdependence emerges.

Across Chapters Three and Four, the focus shifts from a general account of group cognition, to sports teams as an exemplar of this phenomenon. Both chapters explore the ways in which members of sports teams are cognitively interdependent, highlighting variability between teams and across different sports. The kind of sports teams discussed in Chapter Three are those with a shared history of playing and training together, emphasising the enduring forms of cognitive interdependence. A combination of existing qualitative research and original qualitative research is used to explore such interdependence. In contrast, Chapter Four is concerned with the situated and embodied aspects of sports team cognition, emphasising the way in which features of a teams' environment affords certain actions. Evidence of this phenomenon is drawn from existing lab-based studies of individual elite athletes from sports psychology, and alignment studies from cognitive psychology. Chapters Three and Four are evidence of the various ways in which multiple individuals can be cognitively interdependent.

Finally, in Chapter Five key principles are extracted from the sports team exemplar and applied to other kinds of groups, thereby generalising this account of group cognition, while also acknowledging the existence of variability across groups and domains. Theoretical and methodological implications for a philosophy of team and small group cognition are also discussed in detail.

Chapter One

The Conceptual and Empirical Space for Group Cognition

1. Introduction

Groups of people can do extraordinary things. A team of chefs is able to satisfy hundreds of guests, night after night, impressing them with the quality of that which they create and the efficiency and professionalism with which it is created. A Premier League football team astounds fans across the world with captivating on-field performances. Teams of paramedics operate under some of the most extreme temporal and emotional pressures confronted by humans, and are able to do so successfully. While there are countless examples of high-performing groups of people, there are likely just as many unsuccessful groups; not all football teams dazzle spectators, especially in amateur competitions, and even at the elite level the best groups can have an 'off day.' Just as there are differences in abilities, skills and performance between individuals, there are differences between groups. Groups of people vary not only in the extent or continuity of their successes, but also across numerous other dimensions including longevity, formality and composition. For instance, some groups are much less formal than professional teams of chefs or athletes. These less formal groups are the everyday groups we find ourselves in: book clubs, parent groups, craft groups, temporary committees, families and so on. What all of these groups, both formal and informal, have in common is that on many occasions they involve multiple people undertaking cognitively demanding tasks together. People are remembering, planning, acting and deciding together. The cognitions and actions of group members are coordinated such that the group is able to attempt, maybe even complete the relevant task, and if all goes well, do so successfully. People are remembering, planning, acting and deciding together.

From a scientific and philosophical perspective, the existence of such groups raises a number of interesting questions. In particular, what, if anything, makes genuine group

behaviour different from aggregated individuals' behaviour such as the behaviour of a mob or herd? And, particularly compelling, what are the processes or factors that enable multiple individuals to act as a group, coordinating their actions and cognitions, especially when doing so successfully? To answer these questions, this thesis characterizes groups of people, under the right conditions, as cognitive systems in their own right. According to this view, in the right circumstances, multiple individuals interact in such a way that they form and behave as a cognitive system. Some groups make decisions, form beliefs, act in light of these beliefs, remember together and solve cognitively demanding problems together, as a dynamic and distributed cognitive system. The characterisation of groups as cognitive systems and the theory of group cognition developed herein is a naturalized account of the 'group mind', informed by contemporary cognitive science. The relationship between group mind and group cognition will be addressed more directly in Section Six of this chapter.

While it would be difficult to deny that human beings complete a variety of tasks through co-operation with others as a group, it is something else altogether to claim that the group is itself a cognitive system. The basic idea is that in the right circumstances groups of people undertaking a task together display intelligent behaviour that is best attributed to the group as a whole. Roughly, intelligent behaviour is understood as involving a variety of capacities such as learning, decision-making, memory, flexibility and the capacity to change and update behaviour in light of changes in the environment. As part of a scientifically tractable and philosophically plausible account of group cognition it is argued herein that truly cognitive groups possess the kinds of capacities that distinguish intelligent behaviour.

This thesis has three key aims: firstly, to provide a way of distinguishing cases of real group cognition from cases of aggregated individual cognition; secondly, to explain how some aspects of group behaviour and some group properties can be cognitive; and thirdly to identify the defining features and processes of group cognition, particularly those that allow

multiple individuals to coordinate their behaviour and collaborate successfully. Chapters One and Two are devoted to the first aim, providing a way of distinguishing group cognition from aggregate individual cognition by drawing on existing work in social ontology and refining it with a scientifically informed account of emergence. Distinguishing between real and accidental groups in this way has the added advantage of warding off any objections to the effect of the claim that group cognition is a non-starter because it can be reduced to individual cognition. The first two chapters are essentially concerned with showing that group cognition is a real phenomenon. Chapter Two emphasises the importance of cognitive interdependence as a yardstick for identifying cases of real group cognition. Group members are cognitively interdependent if they mutually influence each other's cognition, not merely in terms of a stimulus and response relation, but through transforming each other's cognition and affecting each other's cognitive processing.

Building on the characterisation of group cognition as an emergent phenomenon which is laid out in the first two chapters, Chapters Three and Four are a case study of sports teams as an exemplar of group cognition. These chapters address the third aim of the thesis by identifying the processes and factors that facilitate successful group performance. Across both chapters, the various ways in which members of a sports team are cognitively interdependent are highlighted. Chapter Three deals specifically with enduring but updatable forms of cognitive interdependence and Chapter Four focuses on the kind of situation-specific embodied and situated forms of cognitive interdependence that allow groups of people to respond rapidly and effectively to a changing task environment. The account that arises across the sports case study highlights the distributed and dynamic aspects of group cognition. The dynamic component of this account focuses on the many ways that co-actors or group members mutually shape one another's cognitive processing, quickly and seemingly automatically. The distributed component emphasises the specific kind of knowledge representations that are spread across group members, becoming refined and updated over

time, and accessed in ways idiosyncratic to the group itself. This dynamic-distributed account of group cognition captures both the synchronic and diachronic life of cognitive groups. By analysing sports team performance it is possible to extract general principles that apply to groups in other domains. This is the task for Chapter Five, where key principles are extracted from the sports team exemplar and applied to other kinds of groups, thereby generalising the account of group cognition, while also acknowledging the existence of variability across groups and domains. The theoretical and methodological implications of the dynamic-distributed account are also described in detail, as a philosophy of team and small group cognition. The five chapters comprising this thesis culminate in an account of group cognition that is fundamentally naturalistic in its methodology and commitments, making it both scientifically tractable and philosophically plausible.

Using a combination of conceptual analysis, existing empirical research and original empirical research, this thesis explores the features and processes that enable a group to complete cognitively demanding tasks, like remembering and deciding, and to adapt and change its behaviour in response to an unpredictable and changing environment. On the basis of this research, groups of people should come to be seen as a new and fruitful unit of analysis for the cognitive sciences, with the potential not only to reveal previously unexplored cognitive phenomena, but also to assist scientists and philosophers in refining what we currently know about individual cognition. Perhaps the most interesting implication of characterising groups as cognitive systems is that it brings us closer to understanding cognition as it happens in everyday life or in the 'wild.' A large part of human behaviour is conducted in the company of others or in co-operation with others, and many tasks can only be completed through the co-ordination of multiple individuals' actions and cognitions. So by understanding not just what each individual contributes to the collaborative effort but also how the various contributions mesh together, how it is that group members shape and constrain each others' cognitive processing, we come to know more about cognition as it

happens in the real world, namely as dynamic and distributed group cognition.

2. Finding Group Cognition

Commonly, in our everyday language, we unthinkingly and naturally refer to groups as thinkers or cognizers. We talk about political parties making a decision, families planning for the future and sports teams being disppointed with their performance. Given the prevalence of these kinds of mental state attributions to groups, we can question whether or not there is any truth to them. Does our everyday language capture something real about the world? Do groups have minds of their own?

If possible, momentarily suspend any skepticism you may have and accept for argument's sake that it might be possible to count groups among the various candidates for cognitive systems (the notion of group cognition remains to be more fully developed and defended, but let's just see what intuitive merit it has). Where might we expect to find cases of group cognition and what sorts of phenomena does the term 'group cognition' catch? The following three examples provide a starting point for delineating group cognition from individual cognition, and help to highlight the sorts of things that an account of group cognition must begin to explain if such a phenomenon is possible in principle.

The Deft Coordination of a Soccer Team

Consider the way in which the members of a soccer team must co-ordinate their actions to defend against a charging opposition. Defenders must locate themselves in space in a way that is responsive not only to the opponent attackers, but also to the location of the rest of their own team: standing too close, too flat or too far apart can play into the opposition's hands. The team defending must 'decide' who acts, and when: which defender should approach the oncoming attacker who has the ball? Who should mark the opponents that are positioned to receive the ball? And who should be marking space, anticipating the direction in which the attacker will be forced to play or where a teammate will look once they steal the ball away? Collectively, the team defending must make these decisions within the spatial and temporal

constraints imposed by the opposition. Some of this play will of course depend on each individual's own skills and expertise, but it will also depend on how the teammates interact and share information so as to maximize defensive resources.

Often this sort of defensive play arises after the opposition has thwarted the team's attacking efforts. This leaves very little time to respond and organise their actions. The defending team needs to monitor and respond to the attacker with the ball while the remaining defenders spread themselves between attacking players, the goal, and potentially open spaces. This arrangement has to be sorted in a matter of seconds with little time to provide verbal instructions to every teammate as to where they should run and how they should respond. Instead, it is much more complex. Players must anticipate, on the fly, where certain teammates are likely to move next, and must adjust their own positioning accordingly. We can speculate that this co-ordination is likely to be achieved through memory, a mix of low-level perceptual information exchange between players, the use of short, sharp verbal cues unique to the team, and expectations as to how teammates will act, based on some combination of their current behaviour and teammate familiarity. All of these factors shape and constrain what the players will do. The more efficient these modes of co-ordination and awareness, both implicit and explicit, the more successful a team is likely to be (think of the difference between the highly polished performance of professional club teams as compared to the often dull, and less polished performance of national teams.) This combination of individual level behaviour and group level co-ordination is central to the account of group cognition developed herein. Unlike a bunch of individuals playing together spontaneously in a park, the team with a history of playing and training together is likely to be much more interdependent, and better able to co-ordinate their actions.

As well as fast, on the fly performances, there are also longer, more-enduring features of sports team performance. Despite the idiosyncrasies of specific teams or clubs, sports teams generally seem to have some system of explicit agreement in place that ensures team

members see through the particular commitments that arise as a member of a team, such as attendance at matches and training. This aspect of team action may be best explained in terms of individuals' intentions and desires, and the ways in which each individual's intentions are shaped by others' intentions or desires, creating some sort of interconnected, mental state structure that spans the group. Explanations of group behaviour in terms of these more enduring mental structures can then be supplemented by explanations of the lower-level, implicit processes that also guide the team's behaviour.

Other enduring, historical aspects of team behaviour might include particular and idiosyncratic patterns of practice that arise and are repeated at training, and the gradual spread of systems of verbal and gestural cueing between players, both of which have some bearing not only on team success, but also on the specific way in which the team performs, even in moments of on the fly play. In a sense, it is possible that the history of the team is present in the specific ways in which the team responds to a changing, often chaotic environment. This of course remains to be more solidly argued, but at first glance we can imagine that a team of players go through processes of adapting to each other through training, and even through socializing outside of matches and training, which leads to changes in the kinds of play that the team executes, and how they execute them. From cases like the soccer team, a specific set of questions arises for an account of group cognition: what is the difference between a team that has a history of playing and training together and a group that just meets in a park and decides to have a kick around? How best can we understand these differences? And why might the former be a cognitive unit, but not the latter? What is the nature of the enduring mental structures, and are they akin to intentions? And what is the nature of the more implicit, non-verbal and perceptual processes of information sharing? Why should we think of the team as a cognitive unit?

The Smooth Running of a Legal Practice

In a different vein, consider the possibility that a small office of professionals, say, a law firm,

is a cognitive system. This group is in some ways very different from a soccer team, as their primary mode of information exchange will be spoken and written words, and their actions will often rely on detailed pre-planning and deliberation. We can imagine though that this planning and deliberation is done not by each of the individuals in isolation, but is instead a matter of specific kinds of co-operation and co-ordination. A law practice, with four solicitors and three administrative staff, might operate via weekly meetings where matters are distributed among the solicitors and other tasks among the remaining staff. The process of dividing up the cases might be based on what each staff member knows of the expertise of the other staff members, where the cases each solicitor gets depends upon what others know of their expertise. And then in dealing with the matters, the solicitors and support staff might exchange ideas and suggestions as to how best to proceed with a given matter. Smaller meetings and appointments might be made between sub-groups within the law practice, or more impromptu information exchanges such as calling out from one room to someone in an adjoining room for advice on what to do, both of which are examples of the kind of interdependence between staff members that is at play within the firm.

The successful functioning of the law firm is not the result of each person's efforts simply added together. Instead, it is achieved through their interactions, through some members relying on other members for specific pieces of information and knowledge, and doing so routinely. Law practices are also likely to have their own idiosyncratic systems of rules and norms for dealing with particular matters. There might be rules about having junior staff members show their work to more senior members, or requirements that senior members are cc'd into all emails. Like the soccer team, the sharing of information between members is fundamental to the success of the law practice. However, unlike the soccer team, the law firm's transmission of information is going to leave more of a paper trail, relying on more static styles of communication including emails, dictations, other forms of verbal communications and, of course, various written forms of communication including letters,

forms and, very importantly, post-it notes. Despite these differences, there are some important similarities between the soccer team and the law practice. For one, they both rely on multiple individuals fulfilling specific roles, not in isolation but in collaboration with each other. While no one would doubt that soccer teams and law practices do what they do through collaboration, it is more controversial to claim that in doing so the team and the practice act as a cognitive system. Developing and defending this latter view is the main aim of this thesis. An initial hint, though, comes in the ways in which members of the soccer team and the law practice depend *vitally* on others for the successful completion of their own task. We could speculate that each individuals' performance would be significantly different were they acting alone or with another group. Given this interdependence, many of the processes undertaken by the team, either deciding to execute a particular play in the case of the soccer team, or remembering a client's history in the case of the law firm, are spread across multiple individuals. These tasks are cognitive, just as they are for individual cognition -they require decision-making skills, working memory, long-term memory, attention and so on, and yet these processes are distributed throughout the team or office.

Sharing Memories

In contrast to both the soccer example and the law practice, a family may be considered a much less formal group of interacting agents. It is not unusual for a family to share experiences together, which at a later point in time are best remembered through impromptu family reminiscing. This might happen for instance on the way to a grandparent's place, or on the way to a holiday destination that had previously been visited, or maybe in even more mundane situations such as over dinner or around the television. In these situations different family members contribute different aspects to the memory, someone might remember when it was, someone else might remember who was there, and someone else might remember that a poignant anecdote, now commonly spoken about in the family, first surfaced on this particular holiday. There are also more systematic ways in which a family might remember

together, including the use of a communal calendar that helps regulate who needs to be where and when. Particular family members might also be responsible for remembering specific kinds of information such as phone numbers, or the names of distant relations, knowledge of which can be invoked by other family members' requests. The family differs from the soccer team and the law practice in that in some senses it seems to be more inherently informal, it is less common that members of a family have signed a contract or applied to be part of a family. Many of the norms or rules that guide how family members undertake activities are often unspoken, and the activities themselves can sometimes be leisure-based. Of course, the are other senses in which a family might be more formal in so far as it has a more formal legal status than, say, a soccer team. This kind of variation and difficulty in identifying and quantifying variation in groups is a recurring theme across this thesis. Despite such variation, for the family, the soccer team and the law practice, the tasks being completed all have a cognitive flavour to them - they require cognitive capacities like memory, planning and decision-making, which seem to be carried out by multiple individuals who are bound together by their history of interactions.

The three cases discussed here are primers for the account of group cognition developed in the following chapters. They gesture at the kind of cases that might qualify as group cognition, and where 'in the wild' we might find group cognition. They also help to identify some of the questions that need to be addressed as part of a positive account of group cognition, questions such as what, if anything, about these social, group interactions is cognitive? Are any of the mental processes distributed across the group as well as being realized by a single individual? How is that a group, like a soccer team, is able to achieve the goals it sets out to achieve? And importantly, what makes a group of people a cognitive system, rather than just an accidental bunch of people? From these kinds of questions, and surfacing from the disparate areas of research canvassed across this account of group cognition, three main questions or puzzles need to be addressed if this account is to be

scientifically and philosophical plausible.

What each of the three examples sketched above has in common is that they are all relatively small, tight-knit groups. Smaller groups, or groups of comparable size, will be the focus of the account of group cognition developed in this thesis. This is not a principled choice, but a pragmatic one. Smaller groups are easier to deal with in terms of the quantity of data that they generate and, because of this, most of the existing research uses smaller groups as their objects of study. Given that this is not a principled choice, it is possible, and indeed desirable, that this account of group cognition can be extended to larger groups. The kind of emergence outlined in Chapter Two can be applied to larger scales in an effort to distinguish between real group behaviour, and aggregate individual behaviour. And the dynamic-distributed account should also extend to larger groups. The kind of larger groups that might be successfully captured by an extended version of this account would be institutions like political parties or nations of people, and perhaps even internet-based communities, towns, unions and so on. Aspects of the account may also extend to animal groups, including herds, hives and schools. As with larger human groups, animal groups are not the focus on this account of group cognition for pragmatic reasons and would ideally be accommodated by this account.

3. The Whys and Hows of Group Cognition Research

From the above examples, and surfacing from the disparate areas of research canvassed across this account of group cognition, three main questions or problems need to be addressed if this account is to be scientifically and philosophically plausible. Firstly, we need a way of securing the reality of group cognition: how is group cognition distinct from aggregated individual cognition? This is addressed explicitly in Chapter Two, by adopting a mechanistic characterisation of emergence. Secondly, how do groups of people achieve successful collaboration? What are the cognitive factors and processes that drive successful group

action? This questioned is answered in Chapters Three and Four. And thirdly, but relatedly, how does such collaborative behaviour operate in on the fly situations, or in cases where the group must respond to an unpredictable, changing environment? This problem is specifically addresses in Chapter Four as it applies to sports teams, and then generalized to other groups in Chapter Five. A plausible account of group cognition should be able to address the three questions and accommodate the diversity of cases of group behaviour as represented by the above examples.

The first of these problems, the metaphysical problem, appears to be unusual in the cognitive sciences. For example, most scientists and philosophers seem comfortable with the idea that a person's memory exists, or that a person's visual perception exists. On the whole, much of the theoretical and experimental work on memory and vision is aimed at explaining how memory and vision work, and not at trying to prove that they exist in the first place. Group cognition, on the other hand, is strange in that it we first need to prove it exists before exploring how it works. So the first problem one must overcome is to show that group cognition is real, that there really is something cognitive that exists above and beyond individual cognition, that the whole really is more than the sum of the parts. And if this is to be consistent with mainstream cognitive science, then it must be done in such a way that is metaphysically acceptable, and not mystical in any way.

It is worth noting that while the problem of the reality of group cognitive properties is the natural starting point, it is also a possibility that the reality of group properties more generally needs to be proven. For the purposes of the current project, the reality of groups as ontological entities will be taken for granted. One point that needs to be made, however, is that groups of people likely possess or display different kinds of properties, including cognitive and non-cognitive properties. For example, a group of four people pushing a broken down car will display both non-cognitive properties, such as strength and fitness, and, as this thesis proposes, cognitive properties such as information-processing and shared intentions.

The non-cognitive properties possessed by a group can themselves be complex phenomena, and may be attributable to the group as a whole rather than to a single individual. Many non-cognitive properties, like strength and fitness, are going to be important for the group's successful completion of cognitively demanding tasks. But just as cognitive science separates cognitive properties such as information-processing or mental representation from non-cognitive properties in order to explain individual cognition, so too does this account of group cognition, in the interests of making focused explanatory progress. A fuller, more complete picture of group behaviour would address the relationship between various cognitive and non-cognitive properties, but this is beyond the scope of the current project which is centrally concerned with the only recently proposed explanandum of group cognition. There is more to say about the reality of group cognitive properties in Chapter Two, where the notion of the group being more than the sum of its parts is given a precise explanation. The important thing for now is to begin to motivate the kind of approach taken in this account of group cognition.

The first motivation for setting out to answer the three key questions is a rather optimistic one. Based on casual observation, many activities are undertaken by a group of people and can only be completed by a group, as is the case with a soccer team. By attempting to account for the group's cognition, it is possible to better characterise cognition as it happens in everyday life. After all, it is the group that is undertaking the task and in our everyday observations and descriptions of the task it seems quite natural to refer to the group as the entity completing the cognitively demanding task. So if nothing else, we may as well look to see if we can attribute performance to the group as a whole in a robust, non-mystical way. Moreover, in everyday language we readily attribute mental or cognitive properties to groups of people. We talk of political parties believing, of sports teams being disappointed and of bands deciding to break up. We can therefore ask whether or not there is any truth to these attributions. That is, are groups of people cognitive?

The second motivation stems from that fact that this account of group cognition does

not exist in isolation. There are other areas of research that flirt with the notion that groups of people are information processors, or cognitive units, or intelligent entities. So this particular account of group cognition is intended to re-invigorate studies of group behaviour such as those from organisational psychology that offer very broad, abstract descriptions of group information-processing. These studies uncover a lot in their own right but are not without their own limitations, so it is worth seeing what can be gained by introducing a more cognitive flavour into existing theoretical frameworks.

Thirdly, by characterizing group cognition in the right way, it may be possible to start to understand how and why a group or team with a history together may be able to outperform a group or team of high performing individuals without a shared history. In other words, what is special about an expert team when compared to a team of experts? (Eccles 2004).

The guiding methodology for the dynamic-distributed account is to bring together the most promising areas of theoretical and empirical research that have, until now, been relatively isolated from other. Taken separately, each area of research takes a step towards showing that group cognition is real and distinct from individual cognition. However, when the strengths of these areas are brought together and the limitations are overcome, the reality of group cognition as distinct from individual cognition is much clearer and stronger. In particular, I draw on research from cognitive psychology, organisational psychology, sports science and social ontology. Given just how common joint activities are for human beings, it is no surprise that they are the object of inquiry for several different academic disciplines, including philosophy, psychology, sociology and economics. It is only a relatively recent proposal in the philosophy of mind and on the fringes of cognitive science that groups are a unit of analysis for intelligent behaviour (Hutchins 1995a, 1995b; Giere & Moffatt 2003; Kirsh 2009; Sutton 2010; Theiner 2010a, 2010b, Huebner 2008, 2011). As a result, the current state of research is still rather speculative. This provides the opportunity to bring together a variety of different existing research in an attempt to answer questions pertaining to

group cognition and behaviour. Unlike individual cognition, which has enjoyed over half a century of research since the origin of modern cognitive science, a great deal of groundwork needs to be laid just to show that group cognition really does exist and is a phenomenon out there in the world, ready to be studied and explained using a combination of conceptual and empirical means. This is a fortunate state of play, making it possible to draw on seemingly disparate fields.

4. In What Sense is Group Cognition Cognitive?

While it remains to be shown that under the right conditions groups of people are distributed and dynamic cognitive systems (this is the task of all subsequent chapters), it will be useful to sketch the characterisation of cognition that is operating in the background of the distributed-dynamic account of group cognition. Here, we can sketch the relationship between mind and cognition, both in general terms and as it applies to groups of people. The primary explanandum of this thesis is the intelligent behavior of groups. Therefore, the terms mind, cognition and intelligence are used fairly interchangeably and liberally to denote intelligent behaviour.

In keeping with contemporary cognitive science, cognition can broadly be described as information-processing whereby the behavior of the entity or systems is goal-directed and driven by the transformation of representations whose content shapes the systems' behavior (Dennett 1996; Clark 2005; O'Brien & Opie 2006; Boden 2006; Bermudez 2010). While there are a great number of conceptual issues surrounding the exact nature of mind and how it relates to the concept of cognition, such as the problem of explaining qualia or phenomenal states, and the possibility of explaining consciousness, the approach to cognition and mind advanced as part of this account of group cognition is to bracket off these more peripheral debates. Instead, I adopt a form of methodological optimism and operate with a concept of mind as cognition, driven by goal-direct information-processing, leaving the other debates for

another time. What, then, does this characterisation entail?

The unifying view of cognitive scientists and philosophers is that thinkers are physical systems whose behaviour patterns are reason respecting (Clark 1999, 2001). Thinkers act on the basis of beliefs, thoughts and goals. Mental states of this kind guide the thinker's behaviour and direct the thinker's information-processing. The motivation for describing cognition in terms of computation and information-processing guided by intentional states like beliefs and goals, arises from the inadequacy of certain kinds of purely physical explanations in explaining human behaviour. There is a sense in which intelligent behaviour is constant despite the underlying physical structures and processes involved in the behaviour. As Pylyshyn argues, if a person witnesses an accident and rushes to call the emergency services, an explanation of the person's behaviour in physical terms does not capture what is important about that person's behaviour. We can imagine that this kind of behaviour would be found time and again, in all cultures and parts of the world regardless of any underlying physical differences in the phone device or number dialed (Pylyshyn 1986; Clark 2005). This example shows that the behavior can be the same despite the underlying physical differences. Therefore, we can abstract away from the physical features of the behavior and explain intelligent behavior in terms of the role played by states such as beliefs, desires and intentions.

Considerations such as these have informed the functionalist characterisation of mind and cognition, whereby mental or cognitive states are identified and explained by the causal role they play in the system's behaviour, not exclusively in terms of the underlying physical processes (Fodor 1986; Clark 2008) . Such states can be realized by a variety of different substrates. Cognitive properties are identified in terms of their functions, operations and causal powers, rather than in terms of their biological or physical constitution. This is the first suggestion that cognitive states may be realized by groups of people, rather than just a single person: cognizers could be realized in teams or in the nation of China. If the belief or

intention or goal causes the behaviour of the group, and is not attributable to the group in terms of an aggregation of individuals' cognitive states, then the group is a candidate for a cognitive system. The sports teams and groups described in this thesis, whose behaviour is the result of the interdependence between members giving rise to beliefs, intentions and decisions shared across multiple members, are likely candidates for a cognitive system on a functionalist view of cognition.

Many theorists build into their characterisation of cognition further requirements about the architecture of the computations performed by the cognitive system. The nature of this architecture is itself a contested topic. Again, in the interests of making progress in explaining group behaviour, we can bracket these debates and adopt a sufficiently broad characterisation of cognition to accommodate the different perspectives. For this, we can borrow from philosophers of cognitive science Gerard O'Brien and Jon Opie the characterisation of cognition as consisting in computations that:

> "…are causal processes that implicate one or more representing vehicles, such that their trajectory is shaped by the representational contents of those vehicles" (2006, p 32.)

In creating this characterisation, O'Brien and Opie (2006) have sought to capture a diverse range of accounts of cognition from classical, symbol manipulating accounts to connectionist models and conceptualizations of computation. What is important about this characterisation is the emphasis on representation. A cognitive system's behaviour is driven by the content of its representations.

This characterisation of cognition is also compatible with Georg Theiner's account of cognition, which forms part of his research on group cognition (2010a, 2010c; Theiner, Allen & Goldstone 2010b). Theiner proposes that cognition is best treated as a multifaceted concept consisting in a variety of capacities, one of which is the capacity for representation (2010a). Other capacities include selectively attending to the environment, information-processing and

adaptability to the environment. There are some theorists, however, who reject one or more of these capacities – this is further motivation for applying a broad characterisation of cognition to groups – and while these features are picked out and endorsed by many philosophers and cognitive scientists, there is no single, accepted understanding of cognition. The above capacities enable the cognizer to negotiate its environment and complete goal-directed tasks and drive the system's intelligent behaviour, including exercising intelligent behaviours such as remembering, planning, deciding and acting. How, then, might these capacities apply to groups of people?

Drawing on Theiner (2010a, 2010c), we can begin by identifying the intelligent capacities of a group, and emphasizing the way in which this behaviour emerges from the interdependencies between group members in genuine cognitive groups. We can start with decision-making in a soccer team. A not altogether uncommon move for a soccer team on the attack is to spread apart so as to spread the defence and create gaps to run the ball through at un-defendable angles. One way that this can be achieved is through verbal instructions between players. Or, it could be achieved through some kind of de-centralised deciding, where the team appears to self-organise, with no one person in particular deciding that the team should spread. In Chapters Three and Four, I explore the possible ways in which information can be shared throughout the team without explicit instruction. In these cases, the team's attackers may adjust their own body positioning relative to that of their co-attacker. Each player's bodily movement is detected by the other team members and in turn influences their own bodily movement, by detecting affordances for action in their fellow team members' movements. Here the decision to spread emerges from the coupling of the team members' actions and perceptions to each others' actions and perceptions. In this example, the team is not only deciding, but is also acting and planning their act, anticipating how to approach the defence. As we will see in Chapter Four, this can happen rapidly and automatically.

Moving on from identifying displays of intelligence by a group, we can now look for potential evidence of cognitive capacities. As a cognitive system, a group should rely on representations to generate its behaviour, where the contents of those representations cause the system's behaviour. Staying with sports teams as an exemplar of group behaviour, we can see that there are a variety of representational forms used by a team in order for them to decide, remember, plan and act. One obvious representational form is verbal instruction. Verbal instruction may be produced by the coach or training staff during training sessions or during play. They may consist of code words for certain set-plays or routines, or instructive phrases intended as reminders or cues to players to adjust their actions. There is also likely to be a host of verbal instructions exchanged between players on the field during play. Idiosyncratic code words for moves may be used to obscure the team's intentions from the opposition. Gestures may also be used to carry content about a player's intended actions, or what it is that another should be doing. But we have also seen that there are less obvious forms of representation that are relied upon by sports teams. Bodily movements, posture, the angle of limbs or patterns of limb movement all convey information that is meaningful to other team member and illicit specific actions from those who detect this information. Further, it is possible that written and artifactual representations, such as models of set plays, may also be used. Of course, the neural representations with each individual group members will also be efficacious. From this list of representational modes, it is evident that the kinds of representations that facilitate the group's behaviour are extremely hybrid. This suggests that, in order to study the processes that drive the group's behaviour, a mixed methodology will be required. Linguistic analyses of verbal and written communication would need to be supplemented with video coding and eye tracking of more subtle modes of representation. We can imagine that the same kind of representational diversity could be found in other groups. Even a committee meeting to elect a new member would achieve this end via spoken, written and gestural representational sharing.

The capacity for representation has been subjected to considerable debate in the philosophy of cognitive science, relating mostly to the extended mind hypothesis, but may be equally relevant for a theory of group cognition. According to philosophers of mind Adams and Aizawa (2001, 2005, 2008), the 'mark of the cognitive' does not only entail the capacity for representation, but that such representations features non-derived content. This is content that is not derived from public convention, in contrast to linguistic representations whose content is socially determined. For Adams and Aizawa, cognition, as a matter of empirical fact, unfolds through the transformation of representation with intrinsic, non-derived content (2005). To get a foothold against this objection, we can invoke Andy Clark's response (2005), and tailor it to group behaviour. Clark, as a proponent of the extended mind hypothesis[1], responds to this objection by identifying cases of individual cognition that do not involve solely non-derived content, but rather a mix of derived and non-derived content. Clark asks us to imagine solving a problem, by calling to mind a Venn diagram. For Clark, the overlapping circles in our internally imagined diagram have their meaning by virtue of convention. And yet, this seems likely to be a fairly banal example of cognitive processing. An available response to this is that the overlap gets its meaning because of the way in which it is associated with neural underpinnings whose meaning is intrinsic and non-derived (Clark 2005). This need not be too big a problem for a theory of group cognition, nor for the extended mind hypothesis, as the processes that drive group cognition and that are representational are themselves realized in underlying neural processes in the brains of group members. The upshot of this example is that just as individual, internal cognition can involve a hybrid of derived and non-derived content, so too can group cognition. It can involve inner neural processing of group members, linguistic exchanges, gestures, bodily cues and so on. In some cases of sharing memories, it is also possible for groups to combine representations with

[1] The extended mind hypothesis holds that cognition can and does extend beyond the boundary of an individual's skin and skull to objects and maybe even other people in the world (Clark & Chalmers1998). It may be used to motivate claims of group cognition (see Theiner 2010a for example) but is not a central part of the account defended in this thesis.

non-derived content, such as various aspects of an autobiographical memory, to form an emergent representation of the autobiographical memory as a whole. This sort of reply to an Adams and Azaiwa-style objection at least clears the way for group cognition to be potentially plausible.

Having dealt with the capacity for representation in groups, the remaining cognitive capacities are fairly straightforward. The capacity to selectively attend to the environment can be observed in sports teams - for instance, in the way different team members are charged with monitoring and attending to different parts of the environment, which is perhaps akin to the way in which different sensory organs or receptors attend to different parts of an individual's environment. For example, in soccer the defense might be charged with monitoring the opponent attackers, at times passing on information to the midfielders about this part of the playing environment. Similarly, we can think hypothetically about the way in which a team is able to adapt to the environment. We can imagine that a team is adapting to its environment when they must abandon a pre-planned routine in favour of improvising in light of a new and unpredictable opposition. We will see evidence in Chapter Four of the kind of low level processes that would facilitate that adaptability. We could make the same case for a committee who act faster to meet a deadline than was originally anticipated. Finally, as we will see, implicit in many of the cases and empirical evidence described in Chapters Three and Four is the fact that teams process information. One component part, or player, acquires information about an aspect of the task environment and shares it with another, which in turn influences that second player's actions. The information can be shared deliberately, via verbal instruction, or may occur implicitly, via perceptual and attentional processes.

From this, we can see that it is at least theoretically possible that groups can possess cognitive properties. In Chapter Two, the metaphysical framework of mechanistic emergence will be introduced to show that under the right conditions groups can possess cognitive properties that cannot be accounted for in terms of aggregated individual cognition. While

there are many questions about the nature of cognition still rife with dispute, the account of group cognition proposed throughout this thesis is consistent with the core themes of contemporary operationalisations of mind and cognition. As our concepts of mind and cognition become further refined through empirical and conceptual research, so too will our concepts of group mind and cognition. For now, we have a solid starting point.

5. Research Foundations: Filling in the Gaps

The approach adopted in exploring group cognition across this thesis is essentially synthetic. It involves identifying useful conceptual and empirical research that is useful in its own right, but which takes on new explanatory life when integrated with research from other areas. In many cases, the aim is to synthesise the most promising aspects of different research areas and, using this synthesis, to addresses limitations or oversights made within each area. It is a careful balancing act, cherry-picking the most robust and useful research, while also being mindful of the limitations each area meets. The overall sentiment of this thesis is pluralist, and motivated by the possibility of rendering compatible seemingly disparate areas of research, with a view to more usefully capturing all the complexities of group cognition, as it unfolds in the world. The central areas of research drawn on are philosophy of science, social ontology, organisational psychology, cognitive psychology, and sports science and psychology. Bringing together these areas, although this may not be seamless, allows many gaps to be filled. For instance, social ontology, while providing thorough conceptual analysis of how the individual relates to the social, and the nature of joint or group action, has largely ignored empirical research on synchronization and joint action (Deborah Tollefsen (Tollefsen & Dale 2011) and Beth Preston (2012) are exceptions to this). Social ontology and philosophers of action have also under-theorised spontaneous, improvised or on the fly action (Preston 2012). Within social ontology, while many theorists refer to sports teams in passing, as an example of a joint or group action, to my knowledge sport has not been given a thorough treatment.

Yet surprisingly, there is a wealth of lab-based and qualitative empirical research on sports cognition in sports psychology. Sports science is not without its own oversights. The wealth of empirical research has typically taken individual sports performance as its object of study, rather than teams. This is surprising given the large number of sports that are team-based. Similarly, empirical joint action research in cognitive psychology has been limited to working with dyads and triads, at best. In light of these gaps, the account of group cognition developed in this thesis seeks not only to develop a robust explanation of intelligent and successful group behavior, but to suggest routes via which these distinct but relevant areas of research can meet.

The thesis is centrally concerned not only with explaining how successful collaboration is possible, but also with how fast paced, on the fly collaboration can operate in a changing task environment. To exemplify this kind of collaborative action, the thesis draws heavily on skilled, collaborative sports performance, as an exemplar of the more general phenomenon of group cognition. Many of the research areas relied on by this account do not explicitly address skilled action or spontaneous, improvisatory collaborative action, making a sport an ideal test case for these existing areas of research. With this in mind, it is worth briefly describing the key contributions and methods from each area.

Social ontology adapts aspects of philosophy of mind to questions of sociality in an attempt to understand the relationship between the individual and society. While it is obvious from casual observation that many of our daily activities involve joint, shared or collective action, as opposed to an individual acting entirely in isolation, there is much dispute as to how best to analyse these social actions, and this is what social ontology takes as its object. Philosophers working in this area puzzle over questions such as whether or not mental processes are spread across multiple individuals, and how best to conceptually untangle joint action from other forms of action. Where philosophy of mind has until recently had a history of staunch individualism (for example Fodor 1980), social ontology takes sociality and social

interaction as its object of philosophical inquiry. Indeed, philosophy of mind and social ontology make a useful pair, with the former helping to characterize the nature of mentality and the latter trying to better understand the connection between individual mentality and the social realm. Social ontology provides a conceptual starting point for understanding two important components of the distributed-dynamic account of group cognition: firstly, it distinguishes between individual and group level states or processes, and secondly, it identifies what it is that facilitates or constitutes social interaction. By exploring conceptual analyses offered by philosophers in social ontology, it is possible to refine select work from social ontology for our current purposes by adapting it to empirical work from the cognitive and social sciences. This process is bi-directional however, as not only will aspects of social ontology be naturalized, but many of the empirical studies can benefit from the thorough conceptual examinations provided by social ontology. Specifically, in Chapter Two several prominent accounts of shared mental states and joint action are reinterpreted in light of a scientifically plausible account of emergence. These same accounts are then built on in Chapter Three as part of the dynamic-distributed account. The account of group cognition developed in this thesis attempts to answer several key issues in social ontology on its way to a fleshed out theory of group cognition, not in an effort to undermine the enterprise of social ontology but to complement a variety of existing threads within it.

Specifically, I cherry-pick key contributions from Margaret Gilbert's, Philip Pettit's and Michael Bratman's conceptual analyses, weaving together what I think are the most useful aspects for understanding group cognition. Where social ontology is concerned with what it is that constitutes the social world, theorizing about what social things exist and how society relates to the individual and their mental life, this thesis also explores the further issues of whether or not some aspects of the social realm might themselves be cognitive; or more specifically, whether what seem to be primarily social interactions might also be cognitive interactions.

As we will see, Gilbert's, Pettit's and Bratman's different characterisations of joint and collective action have in common the search for a way of understanding the mutual responsiveness to the shared intentions of those individuals participating in a joint or group activity. There seems almost to be a consensus in social ontology that what distinguishes a real group of people acting collectively from an accidental or aggregate group of people is the presence of some form of mutual awareness and responsiveness. This lays the groundwork for the dynamic-distributed account of group cognition, where it is argued that the processes that realise mutual responsiveness to intentions and joint-readiness are dynamic and fluid. To make this case, however, I move beyond social ontology to explore work from the cognitive and psychological sciences that can usefully build on Gilbert's, Pettit's and Bratman's work.

The account of group cognition developed in this thesis relies on social ontology for its foundation but is compatible with cognitive science's characterisation of cognition as involving the transformation of representations, information-processing and computation (Clark 1989; Von Eckardt 1993; Boden 2006). By introducing useful concepts from other areas of research into the philosophy of mind and cognition, it is possible to gain a better understanding in general of cognition as it happens in the wild, and in particular, of group cognition. An obvious area of empirical research that takes groups as it focus is organisational psychology. A common thread of inquiry in this field involves developing explanations of group behaviour in terms of information-processing (Hinsz, Volrath & Tindale 1997). According to this view, groups process information in much the same way as individuals. The focus in this area is on what makes a group or a team of people, often co-operating in formal, organized situations, perform successfully and efficiently. A mix of cognitive and social factors intermix to drive group performance, including factors that influence group motivation, and the group's capacities to search, attend and process select bits of information. While many of these abstract and conceptual models of information-processing in groups are useful for an account of group cognition, many important questions

are left unanswered. What are the mechanisms responsible for the group searching for and attending to different pieces of information? What are the different ways in which information can be transmitted throughout the group and acted on? What sort of information flows throughout the group? And must group cognition be static? Organisational psychology is useful insofar as it enables us to break down the stages involved in information-processing, but we need to look elsewhere for more detailed explanations of the sorts of interactions and processes that drive group performance. In Chapter Three, specific theoretical frameworks are introduced from organisational psychology, but the need to complement them with a more refined characterisation of distributed information-processing is also highlighted.

In contrast to organisational psychology, cognitive ethnography is a method of studying groups of people in the performance of complex tasks in the real world setting in which the tasks are undertaken, emerging from the intersection of cognitive science and anthropology. This sort of research involves microanalyses of social interactions and person-artefact interactions, focusing on the details of how individuals interact with their environment. Typically, this involves videoing groups, collecting diaries or logs and conducting in-depth interviews with group members. For example, Ed Hutchins used this approach to study the cockpit of an aeroplane and the navigation of the navy frigate (Hutchins 1995a, Hutchins 1995b), arguing in the latter case that the important task of remembering the vessel's speed is distributed across multiple crew members and the various measuring devices with which they engage. In contrast to organisational psychology, the emphasis of this kind of research is on the details of how particular processes of information sharing and storage are carried out by the group. Ideally, some sort of alliance between detailed microanalyses on the one hand and abstract models of information-processing on the other should result in an appropriate middle ground for exploring and explaining group cognition. The dynamic-distributed account of group cognition attempts such a meshing conceptually, but also via qualitative data from sports teams in Chapter Four. The data provides insight into the

distribution of knowledge and information across the group, as well as specific moments of group behaviour.

Relatedly, there has been a recent trend in sports psychology to look beyond individual cognition to group and team cognition. Some of this work involves trying to fuse information-processing conceptions of group performance with what is already known about sports performance (Reimer, Park & Hinsz 2006). Studies of athletes' performance at various levels of competition provide insight into the numerous ways in which information can be extracted from the environment and used to guide either the individual's, pair's or whole team's behaviour. Importantly, some areas of sports psychology are accumulating evidence that an athlete's actions are heavily constrained by aspects of the environment such that the skilled perception of select, meaningful aspects of the environment causes the production of specific action response, rapidly and automatically. This is relevant to the current project, as in team sports it is the teammates that make up part of the athlete's environment. Research of this kind is useful because it helps to identify the variety of ways in which information can be transmitted amongst team members. In situations where the environmental conditions are changing and unpredictable, as when a team is faced with a fast paced opposition, alternate ways of sharing information are going to be vital. Sports psychology is also a useful area of empirical research for the simple reason that members of a sports team are so heavily interdependent that if group cognition is to be found anywhere, it is likely to be found here. Examining research on sports teams provides a useful contrast to social ontology's emphasis on static representations such as written communication and pre-planned actions.

Conceptual and empirical research on group cognition or socially distributed cognition as it is often referred to is still in its early stages. This means that the opportunity for inter-disciplinary research spanning a variety of research areas is very much alive. The envisioned place for group cognition research is somewhere between individualist psychology, focusing on what a single individual can do and generalizing from there, and the kind of sociology and

social theory that measures or tracks macro-level trends of populations. Explanations of group cognition are still likely to rely on orthodox cognitive science kinds like memory, decision-making, perception, learning and so on, but our understanding of these will be transformed by what is found out about group behaviour, as it happens in the wild. By bringing together work from organisational psychology, sports psychology, cognitive science and social ontology it is possible to develop a rich account of group cognition.

6. Conceptual Foundations

The study of groups as intelligent entities is only a recent topic for the cognitive sciences and, as such, it is at the fringes of and not part of the mainstream research in this area. Given the newness of this topic, a lot of foundational issues still need to be teased out. In this section, I identify some of the key concepts for the account of group cognition I am developing. The idea is to highlight how key terms have been operationalised on this account and how certain terms are related. The way in which these concepts are defined has been informed by the theoretical commitments developed across this thesis. While they may only make limited sense at the outset, they are further explored in later chapters. Including explanations of these key concepts at this point in the thesis has the advantage of making it clear from the outset the kind of theoretical and empirical claims that are at stake, and how it is they operate in conjunction with the evidence and theorising of subsequent chapters.

The primary explanandum of this thesis is the intelligent behavior of groups of people. As a result, the terms mind, cognition and intelligence are used fairly interchangeably and liberally to denote intelligent behaviour. With respect to these and many other concepts I am a pluralist, endeavouring to create as broad a characterisation of key terms as possible so as to find points of contact not just with supporters but also with critics of the notion of group cognition. Earlier in this chapter, the concept of cognition was discussed in order to explain the characterisation of cognition that applies to groups if, as I argue, it is in fact true that

groups themselves can be cognitive systems. The emphasis here was on the production, transmission and transformation of representations and goal-directed information-processing. We can now turn to an equally important concept for explaining group behavior: action, as a general concept, and as it applies to groups. Again, in the interests of finding common ground for both supporters of the notion of group cognition and critics, I employ a broad view of action. Generally speaking, action is intentional, cognitively-driven bodily and object movement that is produced by a cognitive system. On the basis of empirical evidence from sports and cognitive psychology canvassed in Chapter Four, action ought to be viewed as being intimately connected with and shaped by other cognitive processes including perceptual, attentional and memory processes. On this view, action is shaped by how we see the world, with the kinds of information that is meaningful and useful to us being shaped by how we are acting in the world at a given moment. Cognition and action are thus tightly bound. This characterisation of the coupling between action and perception is defended in Chapter Four on the basis of empirical evidence from sports psychology. This applies and is applied to forms of joint action where two people act together, such as walking together, and group or collective action, where multiple people are acting together such as a dance troupe or soccer team. While many social ontologists, including Pettit, Gilbert and Bratman, distinguish terminologically between joint and group or collective action, on my account of group cognition both can qualify as cases of group cognition if they meet the kinds of conditions set out in Chapter Two.

Having sketched the characterisation of cognition and action that underpins this account of group cognition, we can now move on to outlining the theoretical claims that are specific to this particular theory of group cognition. Put simply, the account of group cognition developed in this thesis highlights the emergent, distributed and dynamic nature of real group cognition. In Chapter Two, a precise formulation of mechanistic emergence is outlined. This particular characterisation of emergence holds that novel properties arise

through the idiosyncratic way in which component parts of a system are organised and interdependent. This organisation shapes what each part, in this case each of the group members, is able to do, which gives rise to emergent properties, in this case the cognitive properties of the group. Across this thesis I explore the variety ways in which groups of people can be cognitively interdependent. Two general forms of cognitive interdependence are identified: interdependent higher-level cognitive states and processes, and lower-level cognitive states and processes. These two different forms are used to distinguish between the different cognitive features that drive successful group collaboration, as informed by a common distinction in cognitive science between levels of cognition. While there is not an accepted, hard and fast definition of the higher-level and lower-level, it is still possible to sketch a general, useful characterisation. In the case of higher-level cognitive states or processes, these are typically mental states like beliefs, desires, intentions, goals and memories. They are potentially available to consciousness, and potentially able to be reported. They may also be of a propositional, truth evaluable nature. Lower-level states or processes on, the other hand, typically operate beneath consciousness and are implicit perceptual and attentional processes. As part of this account of group cognition, it is argued that in many cases lower-level cognitive processes realise higher-level cognitive states. Specific examples of this are described in Chapters Three and Four. As part of showing that group cognition is distinct from aggregated cognition, it is also shown that individual group members can be cognitively interdependent for both higher-level cognitive states or processes and lower-level cognitive states or processes. This will of course be laid out and defended more fully.

Furthermore, the nature of these higher-level and lower-level cognitive states and processes is described as distributed and dynamic cognition. Distributed cognition or socially distributed cognition describes the way in which cognitive states and processes are distributed across multiple individuals. The remembering of an event or planning for a future event may be distributed across the group such that each individual undertakes a sub-process, or

contributes part of a representation, which when combined with those undertaken by the rest of the group forms a new, shared representation or process. The dynamic aspects of group cognition are the way in which group members mutually and continually influence each others' cognitive processing. In many cases this occurs rapidly and automatically, beneath conscious awareness, in response to features of the immediate task environment. The distributed and dynamic aspects of cognition are further explained and defended across the thesis, particular in Chapters Three and Four in relation to sports team performance.

7. Conclusion: Points of Contact and Points of Departure

The dynamic-distributed account of group cognition that is on offer here addresses two kinds of problems. The first involves characterizing the relationship between individual cognition and group cognition, and providing a way of teasing these apart across a variety of possible instances of group cognition, thus showing that group cognition is real and distinct. The second involves identifying the processes that are typical of group cognition, such as how it is that information is shared throughout a group and how members are mutually responsive to each other. Social ontology is the starting point for the answer to each of these problems. In particular, one of the key contributions made by social ontology is its demand for a way of understanding the differences between an accidental group and a real group.

Building on the solution to the first kind of problem, based on work from social ontology, we can enquire as to the nature of mental processes that are shared between individuals. Is planning always involved? How can we understand impromptu and on the fly group cognition? And what are some of the important ways that we communicate with others so as to bring about the shared end? Social ontology is correct in directing our attention to the mutual responsiveness of group members: but how is this achieved? The first set of questions, those pertaining to the project of teasing individual cognition apart from group cognition, will be addressed in the next chapter. Having shown this, it will then possible to further develop

an account of the nature of cognition as realized by groups, which is informed by the cognitive and psychological sciences. In Chapters Three to Five the nature of cognitive processing between individuals is explained, emphasizing the dynamic and distributed aspects of group cognition.

Chapter Two

Groups as Organised Wholes

1. Introduction

Not all philosophers and scientists are comfortable with the idea that groups are cognitive systems. This makes it a strange case for the cognitive sciences because, prior to getting to the nitty gritty details of the nature of group cognition, it is prudent to show that group cognition is something real in the world. Indeed, some might even object that the reality of groups first needs to be proven. In the interest of making progress in explaining group cognition, I am simply going to assume that groups of people are real. Group behaviour is likely to consist of a mix of cognitive and non-cognitive properties, and in explaining such behaviour this thesis focuses on the cognitive properties alone. In this chapter I argue that in some cases groups can be thought of accurately as thinkers or cognitive systems because they are emergent, organised wholes or systems, possessing emergent cognitive properties attributable to the group as a whole. Broadly, the idea is that there are properties that emerge only when individuals interact in certain ways. If this is true, the burden is then on the defender of group cognition, such as myself, to show that these emergent properties can be cognitive. This will be the task of subsequent chapters – identifying different forms of emergent cognitive properties.

For now though, to support the claim that group cognition is real and in some way distinct from individual cognition, I draw on work from the philosophy of science on mechanisms and mechanistic explanations. In so doing, I develop a way of usefully characterising the relationship between individual group members and the group itself, as this is the first kind of problem that an account of group cognition must overcome – the reduction of group cognition to individuals' cognition. As we will see, however, once groups are understood to be organized wholes, claims of reduction are potentially less threatening to the

ontological status of groups as distributed and dynamic cognitive systems. By adopting a mechanistic perspective we gain the added benefit of being able to unify research relevant to group cognition from a variety of different fields including social ontology, cognitive psychology, sports psychology and organisational psychology. Adapting the mechanistic framework to groups of people highlights the importance of cognitive interdependence between group members as the key feature of a group cognitive system.

After outlining the mechanistic characterisation of group cognition, I then borrow from select work in social ontology to explore what it is that makes real group cognition different to accidental, group-like behaviour or the aggregation of multiple individuals' cognition, as in a herd or hive. I suggest that aspects of Gilbert's, Bratman's and Pettit's work in social ontology are compatible with a mechanistic characterisation of group cognition. Their respective contributions provide sound examples of the kind of 'cognitive interdependence' between group members that is central to the mechanistic, emergent view of group cognition I am adopting. Subsequent chapters then build on the combination of the mechanistic framework with social ontology to provide an account of the kinds of factors and processes that underpin group cognition.

The idea that group cognition is emergent has intuitive appeal. We often talk about the group being more than the sum of the parts, and sometimes something special or unexpected can happen when a group of people get together. The way in which a sports team unexpectedly 'gels' when coming together for the first time, or the way in which a band seamlessly carries on from an unexpected mistake by one of its members. As we will see, a group's behaviour may indeed be special and ineffable to spectators, yet this does not come at the expense of scientific tractability. The particular take on emergent group phenomena outlined in this chapter provides a way of rendering group cognition examinable in philosophy and cognitive science.

2. Difficulties for Group Cognition

It is not uncommon to talk of a sports team being disappointed, or a committee planning for the future. Descriptions of this kind involve the attribution of mental states to collective entities, and they feature heavily in our everyday language. We attribute mental states to groups almost as readily as we attribute them to other individuals. Whilst it is not prudent to take everyday language as our single best ontological guide, it points us in the direction of an interesting phenomenon: namely, the behaviour of groups of people. However, there is controversy as to whether or not groups are the entities that possess mental states or whether it is the only the individual members that possess mental states. While ordinary language alone is unlikely to be our best guide to an ontology, the sheer prevalence of mental state attribution is motivation for asking whether or not there is any truth to the way we talk about groups. The possibility of groups possessing mental or cognitive properties has also featured in more empirically informed domains of thought, including biology, psychology and sociology.

The possibility that groups of people or animals posses mental or cognitive properties has a substantial history in the social and biological science. Two of the most prominent forms of group mind theories, as discussed in considerable detail by Robert Wilson (2004, 2005), derive from the social sciences and biological sciences, in the form of the superorganism or swarm intelligence in the biological sciences and collective psychology in the social sciences. The term superorganism has been used by ecologists and biologists when attributing cognitive or mental capacities to groups of organisms, especially social insects like bees or ants (Wilson 2004). On this view, communities of organisms display cognitive-like traits attributable to the whole swarm or hive such as sophisticated communication and perception capacities enabling group foraging. Coinciding with these attributions of mental properties is the claim that natural selection operates on the group as a whole, rather than at the level of the individual organism (Wilson 2004). Underpinning the superorganism tradition

is the observation of an important difference between the individuals and the group: the group as a whole possesses intelligent, mental properties, but none of the individual organisms do (Wilson 2005). Given this distinction, it is not surprising that those theorists working in the superorganism tradition held that the group mind exists at a level which is autonomous from the level of individual organism behaviour. In contrast to this, when we attribute cognitive properties to groups of people, as advocated in this thesis, one undeniable fact is that group members individually possess cognitive properties. While many of the principles outlined as part of the theory of group cognition proposed in this thesis may apply to collectives of animals or insects, the theory itself embraces the relationship between individual cognition and group cognition, explaining how the latter emerges from the former.

The collective psychology tradition developed in the psychological and social sciences, with William McDougall and Emile Durkheim being numbered among the key proponents. Central to this view was the nature of "the crowd," a term used to denote any kind of social gathering or meeting of people, of any number. As Wilson identifies, there are two different strands within this tradition (Wilson 2004, 2005). One strand adopts a negative view of the crowd, according to which the psychology of crowds is very different to the psychology of an individual. Crowd psychology is irrational, deeply emotional and unable to be controlled (Wilson 2004). The other strand embraces the crowd psychology, painting it in a more positive light, attributing major social and cultural achievements to crowds of people (Wilson 2005).

As with the superorganism tradition, proponents of collective psychology advanced nonreductionist accounts of the group mind, whereby properties of the group or collective are separate from properties of individual cognition, and cannot, therefore, be explained by individual, experimental psychology (Wilson 2005). On this view, collective psychology was 'emergent from and thus not reducible to the psychology of the individuals in those collectives and was to be studied as such' (Wilson 2005, p. 268). This view can be

problematic, as it can be difficult to explain the role of individuals in the group or "the crowd" and the relationship between the two. It is particularly difficult to explain the causal powers of the group in relation to the individual. Any appeal to downward causation from group properties to individual properties is potentially metaphysically dubious, as appeals to top-down causation typically are (Craver & Bechtel 2007).

Throughout this and subsequent chapters, I show that, despite what some theorists claim (in the collective psychology and superorganism tradition, but also in the philosophy of science and the philosophy of sport as we see in Chapter Five), reduction and emergence are indeed compatible, if we adopt the right characteristion of emergence. The important upshot of this will be that it is possible to plausibly explain the relationship between group level cognitive properties and the cognitive properties of individual group members. This can be done without appealing to any form of downward causation, and by drawing on the rich traditions of research on individual cognition. It also means we can bridge the gap between research on single individuals, as is typical of orthodox cognitive psychology, and research on larger scale social or group research from sociology and anthropology, which is concerned with groups like nations, cultures and religions. The account of group cognition on offer secures both the reality of group cognition, and also accommodates the individual group members as fully cognitive systems themselves.

While each tradition has its own internal problems, which are dealt with by critics of each tradition, the collective psychology and superorganism traditions draw attention to the difficult issue of which level cognitive properties can accurately be attributed to, and how best to accommodate the mental lives of individual group members. This is an important issue for an account of group cognition, as it gets to the heart of whether or not group cognition is a real phenomenon, worthy of philosophical and scientific enquiry. Critic Robert Rupert characterizes this problem as one of explanatory superfluity (2011). The objection proceeds in the following manner: if an explanation of what appears to be group cognition can be offered

in terms of the individuals' cognition then we will gain no extra explanatory purchase from offering an explanation in terms of the group's cognition. It is therefore unlikely that group cognition is a real phenomenon to be included in our ontology (Rupert 2011). However, as we will see, if we adopt the right version of emergence it is possible to be both a reductionist and an emergentist with respect to cognitive group properties. We can explain cognitive group properties in terms of individuals' cognitive properties without threatening the reality of group cognition.

3. Mechanisms and Groups

Commonly, social scientists and natural scientists frame their explanations in terms of mechanisms. When looking to explain how something happens, scientists endeavour to find the mechanism that is responsible for making that something happen. Economists talk about market mechanisms, biologists talk about mechanisms of inheritance and neuroscientists talk about the neural mechanisms of vision, for example. But what exactly is a mechanism, and how is it that a mechanism can play such an important explanatory role? And, how might this relate to groups of people? There is great interest in these questions in the philosophy of science at the moment and a variety of robust analyses of mechanisms have surfaced. The main aim of this chapter is to develop an account of the relationship between groups and their individual members that draws on some of this work from philosophy of science. Where philosophers have primarily focused on the biological sciences and neuroscience (Machamer, Darden & Craver 2000; Bechtel 2008), I propose that an understanding of mechanism provides a new way of understanding how groups can be thinkers or cognizers in their own right.

Before going on to explain how mechanistic explanation works, it is necessary to start by breaking down the features of a mechanism as it exists in the world – that is, the ontology of a mechanism. We can follow William Bechtel's lead by identifying three key features of a

mechanism: the constituent working parts of a mechanism, their operations, and their organisation (Bechtel 2008). The parts are the different physical components that compose the mechanism, each of which performs different operations. Importantly, the parts need to be organised in a way that enables the mechanism to perform a function. A change to the arrangement of the parts, changes the behaviour of the whole. When the parts are organised or fitted together in the right way, they are able to perform a function that the individual parts, unorganised or organised differently, cannot perform. This makes the mechanism and its capacities emergent, with the whole, or the group at a higher-level to the individual parts, in their unorganised form. It is worth noting that Bechtel, one of the key proponents of mechanism, is open to the possibility of mechanism being applied to groups of people (Bechtel 2009a):[2]

In order to further unpack the way in which mechanisms are emergent, we can adopt William Wimsatt's conditions of aggregativity (Wimsatt 2007). This gives us a way of distinguishing cases where group behaviour is real, distinct and novel, from those cases of an aggregation of individual behaviour. It should be clear by now that this is crucial for showing that group cognition is real and distinct from individual cognition. For Wimsatt, emergence is the failure of aggregativity and is cashed out in terms of the violation of four conditions of aggregativity: inter-substitution, qualitative similarity, invariance under re-aggregation and linearity. The more of these conditions that an entity or system violates, the more emergent it is. If none of the four conditions are violated, then the entity is an aggregate and does not display emergent properties. Of course, there are all of the cases in between. The inter-substitution condition holds that if the arrangement of the parts changes and the behaviour of

[2] Bechtel expresses some reservations about the relationship between the mechanistic framework and its application to accounts of the extended mind thesis, on the basis that the organism using external objects is the locus of cognitive control, and therefore the mechanisms to explain this cognitive control will be internal to the organism. This is different to a group of people or social network, according to Bechtel, because in the coordination of their labour, the group becomes the locus of control: 'Just as individual cells may specialize their operations and coordinate them so as to maintain a multi-celled organism, so individual organisms may specialize their activities and coordinate them to maintain a larger systems such as a social network. In these cases, the social network becomes the locus of control for certain phenomena—those that are carried out by the social network in the service of it' (Bechtel 2009a, p.166).

the whole does not, the entity is an aggregate with respect to this condition. For example, we can change the arrangement of a pile of sand and it will not possess any new properties. The qualitative similarity condition holds that an entity is an aggregate with respect to this condition if the addition or subtraction of a part does not result in a change in the behaviour of the whole. Again, we can imagine this being the case for a pile of sand. The third condition, invariance under re-aggregation, holds that if after the parts are pulled apart and re-assembled there is no change in the behaviour of the whole, then the entity is an aggregate with respect to this condition. And finally, if there is no evidence of co-operation and inhibition between the parts, then the entity is aggregate with respect to the linearity condition. The idea is that the more inter-dependent the parts of a whole are, the more they will violate these conditions, and therefore the more emergent the whole is. Conversely, in cases of aggregation there is no new or distinct phenomenon at the level of the whole.

An important, often overlooked aspect of the mechanistic framework is what it reveals about the relationship between the higher- and lower-level or, more precisely, the relationship between the parts and wholes. While the function that the mechanism performs is emergent in the Wimsattian sense, there is no sense in which the whole causes changes to the parts themselves and their behaviour. This is because the parts compose the whole. If the whole consists entirely of the parts, then it does not make sense to talk about the whole exerting some sort of downward causation over the parts. Instead, it is a matter of constitution (Craver & Bechtel 2007). On a mechanistic view, causation operates within a level, rather than across levels. This means that parts causally effect each other, and organised wholes effect other organised wholes. In the context of groups of people, the corollary is that a group of people can have causal powers on other groups and on the environment, but the relation between the group level and the individual level is one of constitution: a change to the group as a whole is a change to the parts, and vice versa. Furthermore, because of the causal processes operating between the parts or the individuals, their organisation, or how the parts are fitted together,

effects what each of the parts is able to do (Craver & Bechtel 2007; Bechtel 2008)

Thanks to the mechanistic emergence framework, it is possible to make sense of the expression 'more than the sum of the parts' as involving adding up the properties of the individual parts, independent of their interactions and how their interactions change each other. Thus, the whole is more than the sum of the parts when what the whole does is the result of multiple parts interacting and depending on each other for their own functioning. The mechanistic framework has the added bonus of highlighting that the relationship between each of the parts affects how all of the parts behave, and thus the emergent properties. This is especially interesting in groups because often the role that each individual plays in the group depends on the role that the other people in the group play, and this can change over time with dependence growing and agreggation declining or vice versa. Since a physical system can violate anywhere between none and all of the conditions, so too can groups of people. This means that there are likely to be degrees of group cognition, where some groups are more interdependent and better integrated than those that violate only one or two conditions.

Consider for instance the difference between a relay running team and a team of rowers. In the relay team, aside from the baton exchange, each runner performs their task independently of the other members. While each runner is dependent on a particular team member to receive the baton in the right way, which will greatly influence each runner's own performance, the rest of the runner's task will be completed independently of the group. For the relay team, members are less interdependent diachronically, or throughout the course of the whole race, than a team of rowers, with rowers constantly altering their movement in light of how each other rower is performing. For the rowers, they are an interdependent system for the duration of the task as each rower's actions influences their team members' actions; the rowers' actions are coupled together in this context. For the relay team, there are moments of interdependence, and moments of individual performance. These examples speak to the variety of different human groups, and so too does the account of mechanistic emergence as it

permits degrees of cognitive interdependence. Given this diversity, it seems unhelpful to talk about group cognition as operating at its own level, distinct from an individual level. Instead, group cognition is a matter of multiple individuals' cognition being dependent on each other's cognition and action.

To understand how a mechanism produces a particular phenomenon, researchers must, as Bechtel argues, 'look down, around and up' (Bechtel 2009). Looking down involves identifying the parts that comprise the mechanism and the activities they are performing. This requires the mechanism to be decomposed or taken apart, either physically or conceptually, to understand how the various parts fit together, and what activities they perform. As well as looking down, explaining a mechanism requires researchers to look around and up (Bechtel 2009). Looking around involves piecing the parts of a mechanism together and thereby understanding the nature of the mechanism's organisation and how the various parts interact. Finally, researchers must look up, by locating the mechanism in its environment, and accounting for the influence the environment has on the functioning of the mechanism or organised whole. The picture that emerges from this outline of mechanistic explanation is one of multiple levels of explanation, most likely requiring differing methodologies.

A similar unification of mechanism with group cognition has also been proposed by philosopher Georg Theiner (2010a, 2010c; Theiner, Allen & Goldstone 2010b). In particular, Theiner (2010c) develops in great detail the view that there exist degrees of both emergence and group cognition. While this present project is compatible with Theiner's, there is greater emphasis here on unifying disparate fields of research under a mechanistic framework, including social ontology, organisational psychology, sports science and cognitive psychology. This particular project is also focused on explaining the nature of the cognition that a group realises, which takes the form of the dynamic-distributed account of cognition. As mentioned in Chapter One, as well answering the metaphysical problem of the reality of group cognition, this account seeks to explain the factors and processes that facilitate group

collaboration, especially in fast-paced, changing environments. In contrast to Theiner's claims that redundant roles, skills or representations across the group can undermine any claim of group cognition (Theiner, Allen & Goldstone 2010b), in Chapters Three and Four I explore the possibility that some shared or common features are essential for successful group coordination and performance.

By adopting a mechanistic characterisation of groups, we have the resources to distinguish cases of group cognition from aggregated individual cognition, where there is no group level phenomenon. If we adopt a mechanistic view of emergence, we can characterise those situations where group behaviour, especially the associated cognitive properties, is novel and not simply the aggregation of individual behaviour. We can then work on further refining and explaining what exactly the behaviour is, and how the interdependence plays out in specific cases of group cognition.

4. Mechanism and Social Ontology

The mechanistic framework outlined above is neutral with respect to the content of explanations of group cognition. It simply lays out an ontological structure and offers a guide as to how to explain phenomena that have this kind of structure. As a result, it provides a basis for unifying and integrating diverse work on group behaviour so as to enrich our explanations of social groups. Equipped with the relevant conceptual resources from the philosophy of science we can now explore how mechanism assists in making groups higher-level and move forward with some of the issues identified in social ontology, thus rounding out the solution to the first kind of problem: the problem of whether or not group cognition is real. Using the mechanistic framework, we can select and adapt some of the key contributions from social ontology. We can draw on work from social ontology to separate cases of real group behaviour from cases of accidental social behaviour, such as a bunch of strangers moving along an urban street, but without doing so *together*. Various aspects of Margaret

Gilbert's and Michael Bratman's conceptual research on joint action and acting together are useful for our purpose in teasing out the differences between actions genuinely undertaken together and those that merely appear that way. We can then draw on Philip Pettit's work on group action as a starting point for identifying cases where genuine group action is teased apart from aggregated individual action.

Note here the differences in terminology between joint action, collective action and group action. As we will see, Pettit introduces the notion of group action, and outlines the conditions under which it is distinct from joint action. Gilbert and Bratman on the other hand, focus on joint action rather than group action, identifying what makes genuinely doing something with another person or other people, like walking together, different from cases of acting alongside a person or other people, but not *with* them. In the social ontology literature, the terms joint and collective action are typically used to denote acting together as a dyad or as a group. It need not entail claims of a group mind, group agency or group cognition. Gilbert and Bratman are less explicit about treating pairs or groups of people as mental or cognitive entities themselves, whereas Pettit provides specific criteria for cases where groups of people are mental or cognitive entities in their own right. For the account of group cognition being developed in this thesis, both joint, group and collective action qualify as group cognition if the individual cases grouped under those labels feature a kind of cognitive interdependence between those people acting jointly or as a group.

Note also the emphasis placed by social ontology on action rather than cognition. For our purposes, I will give action and cognition the same treatment. The exact relationship between the two will not be addressed given space limitations but, broadly speaking, cognitive processes such as deciding or planning are tightly bound with action, whereby action is not merely the by-product of inner cognitive processing, but an integral part of those cognitive processes. Evidence of the coupling of cognition, perception and action is endorsed in Chapter Four in the context of sports performance. But for now, action and cognition are

treated similarly.

To get a sense of the explanandum of social ontology we can borrow an example developed by John Searle (1990). Searle offers the example of numerous people spread throughout a park who all make their way towards a shelter when it starts to rain. Presumably, this requires very little cooperation between people and is instead a matter of each individual acting alone and more or less independently of the others. At best, each person will be trying to avoid colliding with another person. Searle then asks us to imagine that in the same setting the same individuals act the same way, yet this time as members of a dance troupe performing a piece that involves running to shelter from rain. Unlike the first case, we are probably comfortable referring to the dance troupe's behaviour as a joint or collective action insofar as multiple people are cooperating or working together to achieve the task they are each involved in. We can then ask what the differences are between the two cases? What makes genuinely undertaking a task together different if to the naked eye both cases look the same? Gilbert and Bratman provide answers to questions of this type. Importantly, I suggest that their accounts of acting together provide a conceptual basis for identifying potential kinds of cognitive interdependence between those people involved. To begin, we can start with Margaret Gilbert's model of shared beliefs.

Examining whether or not there is any truth to our folk ascriptions of mental states to groups or pairs of people, Gilbert unpacks what it is for two or more people to form a belief and an intention. In explaining group intentions and group beliefs, Margaret Gilbert argues that individuals form a 'plural subject' (Gilbert 1992). According to Gilbert, forming a plural subject involves two steps. Initially, each individual becomes 'quasi-ready' to participate in a joint action. This is followed by each participant in the joint action becoming jointly ready. Quasi-readiness is distinct from joint-readiness insofar as it pertains to each individual's state of readiness. Joint-readiness on the other hand arises once each member is individually ready and this has been shared, either explicitly through statements of readiness or implicitly,

presumably through gesture, posture, bodily signs and so on. It involves each of the group members being ready to act together as a plural subject. Put succinctly, the plural subject is formed and its members act jointly if, and only if, each individual member acts in light of the sharing of the goal of the action. The plural subject is the agent of the intended action, and is composed of both individuals in a dyad, or all individuals in larger groups. For Gilbert, the formation of a plural subject gives rise to a joint commitment whereby each member accrues a set of (non-moral) duties and obligations, including a duty to see through the shared action and play their respective role in bringing about the shared goal. This is not restricted to joint or shared intentions but extends to joint or shared attitudes and intentional states. When two or more individuals are jointly committed, this is what distinguishes them from accidental groups of people who are not acting jointly. It is the formation of a plural subject on the basis of joint readiness and awareness that is central to Gilbert's account.

In developing her analysis of plural subjecthood and collective or shared belief, Gilbert walks us through an example of a poetry group meeting (1987). The group meets regularly to read and discuss poetry. Gilbert paints the picture of a group that reaches a consensus as to the quality and interpretation of a particular poem. We can imagine the group reading through the poem, discussing its strengths and weaknesses and settling on a particular assessment of the poem. For Gilbert, the interesting thing is that even though a group can reach a decision about a poem, say, that it is well crafted and one stanza is particularly poignant, it can be the case that none of the individual members of the group endorse this view. There might be a consensus across the group as to the quality of the poem, and yet, if each member is individually asked what their own assessment of the poem is, it not too hard to imagine there being a conflict between the group consensus and what an individual might think of the poem. The consensus may not even be held by the majority of individuals in the group. It is just something that arises out of the group discussion. When this happens there is a discontinuity between group preferences and individual preferences. It might be the case

that only the minority of group members hold the group preference as their own individual preference. It could even be possible that no individual holds the preference. Gilbert explains this discontinuity in terms of the different types of roles that individuals can play and the way in which these effect preferences. Where no group has been formed, individuals simply endorse their own, personal preferences and to reach an agreement these can simply be added up. However, when one is a member of a group, each member adopts preferences or intentional states that rely on the role they play as part of a group. Gilbert refers to these preferences as participant preferences (1992). Furthermore for Gilbert, when a person is part of a group, they accrue obligations and duties to act as a member of that group. It is because of this that it is possible for an individual to have personal preferences as well as preferences as a member of a group that are distinct from one another. Cases like these are quite easy to imagine: consider the introduction of a new policy to a work team where the team accepts and endorses the policy yet, individually, the members are not in favor of such a policy. Cases like these suggest that there is something new and novel at a group level.

The discontinuity is useful in highlighting a group property as being different to an individual's property. But what might this mean for cases where individual and group preferences are consistent? Do these fail to be cases of group cognition or group mental states? A sports team, for example, at least a successful, cohesive one, is likely to have continuity between individual and group mental states more often than there is discontinuity; indeed a large part of coaching and team development involves fostering the coherence of individual and team goals and also ensuring that the right kind of knowledge is shared by team members, as discussed in Chapter Three (Eccles & Tenenbaum 2004; Reimer, Park & Hinsz 2006)). So an accurate account of group cognition and group mental states needs to accommodate cases of discontinuity *and* continuity between individual and group states. While this is not to disagree with Gilbert and Pettit, it is to accommodate cases where discontinuity may never arise. On a mechanistic view of group cognition, discontinuity is

irrelevant. I first address this issue here in the context of Gilbert's work, and return to it at the end of this section in relation to Pettit's account of group agency to more fully explain how mechanism captures more cases of group behaviour than those where discontinuity is possible.

If we apply a mechanistic explanation to the poetry group scenario we can see how it is that group intentional states are emergent and ontologically distinct. As unorganised parts, each individual develops their own, personal interpretation and assessment of the poem. However, once the individuals come together as a group, and presumably commit to reaching a decision about the quality of the poem, each individual takes on the role of a member of the poetry group. Through various modes of communication, and the structure of the task, the individuals form a group, or in mechanistic terms, become integrated and organised. Each of the parts (ie each member) is influenced by the other members' contributions, and interactions between them can lead to a group preference that may in fact be quite different from what the majority of members endorse as their own, personal preference or as 'unorganised' parts. It is through the organisation of the parts, in this case the exchange and transference of information, interpretations and ideas, that the group reaches a consensus and thus achieves something new and distinct from the individual parts. By applying the mechanistic framework to this version of the example, we can see that it is an appropriate way of characterising groups as it is consistent with research on joint and collective action. If we use the mechanistic framework as a guide for cases of real group cognition, then the focus is on processes of interaction and interdependence between the working parts, as this is what gives rise to emergent, cognitive group properties.

As another version of Gilbert's poetry example, one informed by my own experience of book clubs, it might be the case that before the poetry group discussion a number of members were undecided about the poem or had not developed fully formed assessments of the poem. Yet, during the course of the discussion in which they, along with the whole group,

endorse the final conclusion, both the group and all of the previously undecided individuals endorse the assessment. If we apply the mechanistic framework here, then the group preference emerges from the processes of communication and idea sharing that took place during the meeting. Even though there is continuity between the group and individual preferences, there are still emergent, group level properties. The group members are somehow mysteriously transformed. Of course, this need not be genuinely mysterious as the empirical cognitive and psychological sciences would be able to identify the processes involved in the group decision-making process, such as those discussed in Chapter Three and Four in the context of sports team performance. Groups are therefore thinkers not only in cases of discontinuity but in cases of interdependence between the parts, where the group's preference, belief or intention is continuous with the individuals'. This is a more relaxed view than that which Gilbert or Pettit suggest. What really matters is the process, not the product. We now have a way of distinguishing genuine, emergent group properties from aggregate properties: that is, by looking for evidence of interdependence and constraint between the parts, or group members. This may or may not result in a discontinuity between individual and group intentional states.

What we can take from Gilbert though, as a means of distinguishing between those people acting together and those that only appear to be acting together, is the presence of quasi-readiness in each individual, that is transformed into joint-readiness. These forms of readiness are represented internally for each person, but we see that there is change for each person, from quasi to joint-readiness, as the prospect of acting together emerges. Of course, there are a variety of ways that this shift can occur, including through the explicit spoken suggestion that people act together, or unspoken bodily suggestions, like gestures and postures, that indicate agreement in sharing the task together, whether it is walking together, moving a table together or joining a street soccer game. Here we can see that something - an intention, a belief, a representation of readiness - changes as a person shifts from acting alone

to preparing to act together. On a mechanistic view, this could be considered evidence of parts (individuals) becoming interdependent. Each person's joint-readiness is dependent on that of the other person. On Gilbert's view, acting together is driven by the formation of a plural subject, where each individual's intentions or readiness is transformed into a joint intention and readiness.

Bratman develops the notion of joint intention as a distinguishing feature of genuinely acting together. The focus on intentions is driven by the view that our actions are guided by the intentions we form towards the world and acting in it (Pacherie 2010). Bratman's account of joint or collective intentionality attempts to capture the interrelatedness of multiple individuals' intentions as part of a joint action, as a pair, or a collective action, or as a group. For Bratman, collective or joint action involves an interpersonal structure of connected intentions that is composed of the intentions of each of the individuals participating in the action (Bratman 1993; Bratman 2009). It takes the following form:

a) intentions on the part of each in favour of our joint activity:
    i. I intend that we J
    ii. You intend that we J
b) I intend that J in accordance with and because of a)i and a)ii, and meshing sub-plans of a)i and a)ii; you intend likewise. (Tollefsen and Dale 2011; Bratman 1993, 2009)

The first of these conditions is rather straightforward. In undertaking a task or activity together, we must both, individually, intend that we undertake the task together. For example, if we are to practice passing a soccer ball together, than I intend that we practice passing together, and you intend that we practice passing together. Here, the kind of interdependence present is in terms of the content of our intentions. For an action to be joint or collective, each

person's intention to undertake the task together is maintained only in light of the others' maintenance of the same intention. This involves everyone being mutually responsive to each person's intention, and changing or updating their intentions of the basis of responsiveness. For each person involved, their own intention in favour of the joint activity persists on the basis of their ongoing knowledge that the other person intends the same. These conditions will not be met in cases that merely *appear* to be a joint or collective action - for example, if we think back to Searle's example of the two groups of people in the park. The first group is a bunch of strangers running to get out of the rain, and the second a bunch, whose actions appear the same as those of the first bunch to an observer, but which are in fact elements of a rehearsed, choreographed performance. In the former case, there would be an absence of individual intentions that we run to shelter together following the pre-planned dance steps. For this group each individual is most likely ignorant of the intentions of the other people in the park. The latter case would meet Bratman's criteria, as each dancer would intend that the group performs the choreographed moves towards shelter. So, from Bratman we can see that one of the ways in which a group of people are interdependent is in terms of each group members' intentions referring to the jointness of the action. Moreover, each person's intentions will be maintained or altered and updated in response to the maintenance or updating of the others' intentions.

Condition B of Bratman's account is particularly useful given its emphasis on intermeshing sub-plans. On Bratman's view, intentions are what guide our planning and preparing for future action. For an action to be joint or collective and, of course, successful, actions must be suitably co-ordinated. If we both agree to travel to the beach together, yet I am left waiting at the bus stop while you drive by in a taxi, then our travelling together has not been very successful. Bratman argues that this would be because of a failure in the meshing of our sub-plans regarding the joint action. So not only does joint or collective action require each person to have an intention that refers to the shared action and the other

participant(s), but that we also have compatibility in the way in which the action will be executed, or in Bratman's terms, that our sub-plans mesh. The full plan for the action may not be articulated fully prior to the action starting, but could unfold as the action unfolds. In cases like these, participants would have to be mutually responsive to each others' actions, as well as intentions. It could also be possible that inter-meshing sub-plans as pursued by each person are different, yet complementary. For example, in soccer, a goalkeeper and defender may each intend to work together to defend the opposition's corner kick. Both players share the same intention, but each will have to act slightly differently. The goalkeeper will be looking to push the ball outside of the danger area, and the defender will be looking to prevent opposition players getting the ball near to the goalkeeper. Despite the differences in each sub-task, both players will know what the other is going to do to achieve their shared intention.

While Bratman's account convincingly separates out genuine joint or collective action from actions that merely appear to be joint or collectively produced, there are some limitations to it for explaining the breadth of cases that I suggest meet the cognitive interdependence yardstick for group cognition. This is due to the emphasis Bratman places on intention and planning. Essentially, intentions are for future actions. If intentions pick out or inform what actions one is likely to perform in the future, then an account of joint or group action in terms of intentions might struggle to explain joint or group actions that are spontaneous, without being planned beforehand (Tollefsen & Dale 2011). In many of our joint or group endeavours, while we can plan for some future actions, or make coarse-grained plans, we must often abandon our plans, or respond on the spot to an unpredictable change in conditions. This is the case in the goalkeeper and defender example. While they might form a shared intention to defend against the corner kick, and while they might be able to form inter-meshing sub-plans that roughly capture what they are both going to try to do, the presence of the opposing team complicates matters. The opposition players introduce variables that cannot be predicted by the jointly acting defenders. To explain how the goalkeeper and

defender manage to save the goal, we have to appeal to other factors and processes beyond intentions and planning. Scaling up to an entire soccer team's collective action, part of their behaviour will best be explained in terms of planned actions, such as overarching game strategies including at what point to bring on particular players from the bench. But there will also be times when what the referee or the opposition players do disrupts the groups' pre-planned strategies. In situations like these the group must rely on other cognitive resources.

Explaining this sort of flexible group action in terms of shared intentions and meshing of sub-plans does not explain how exactly a team is able to change their behavior quickly and efficiently to respond to the changing playing environment. Plans and intentions informing future action cannot be the whole story. As Beth Preston argues, our folk notions of action capture not just planned, future-intended action, but also contingent, unexpected, spontaneous and improvisatory actions (2012). Theories of action that over-emphasise the role played by planning in action production miss an important, and thoroughly prevalent part of human action, namely improvisation, adaptive and on the fly action (Preston 2012). In light of this, it is preferable to endorse a pluralist explanatory framework combining shared intentions and inter-meshing sub-plans with other kinds of mental representations as well as lower-levels of information-processing between group members which can provide richer explanations of joint and group action in its many and varied forms. These possibilities are explored in the next two chapters as part of the dynamic and distributed account of group cognition.

In contrast to Gilbert and Bratman, Pettit develops an account of group agency, where under the right conditions, groups of people have 'minds of their own,' in a sense distinct from the individuals that comprise the group. According to Pettit, for a group of people to constitute a group mind, it must behave rationally, as our scientifically informed characterisations of cognition suggest, through making decisions that fit and are consistent with previous decisions made by the group. Group mental states are therefore like individuals' intentional states and must meet similar requirements of rationality. Like Gilbert,

Pettit notes that interactions between group members can often result in a discontinuity between individual preferences regarding the current decision and the decision that the group reaches, in an effort to maintain consistency with past decisions, and meet the requirements of rationality. This sort of discontinuity suggests that the group can have preferences that are distinct from individual preferences or that groups 'can have a distinct intentional profile from the profiles of their members' (Pettit 2003). Pettit is careful to distinguish between joint action and group agents. In the former case, according to Pettit, dyads or groups of people do not meet the rationality condition for a group mind, but nevertheless act jointly. Group agents, on the other hand, are a special case (Pettit and Schweikard 2006). What then is the condition of group agency?

Central to Pettit's account of group agency is the nature of a group's decision-making capacity and processes. Broadly, groups of people acting jointly are to be considered a group agent if they meet the following three conditions: 1) members establish shared goals, and a means for identifying future goals as they arise; 2) following these goals, members of the group act together to establish a 'body of judgments' for 'rationally guiding action' to achieve these goals; and 3) the group members identify who it is, whether the group as a whole, sub-groups or individuals within the group, that is responsible for acting in pursuit of the group's goals (Pettit & Schweikard 2006). Groups who meet these three conditions will appear to act as a thinker or cognizer. They will form beliefs, pursue goals, protect the group's interests, make decisions, and remember events and facts relevant to the smooth operation of the group - in short, they will appear to be a rational agent. However, as Pettit acknowledges (2003, Pettit & Schweikard 2006; List & Pettit 2011) it is not clear with just these three conditions that the group itself has its own intentional or cognitive profile separate from that of its individual members. Using an analysis of the 'discursive dilemma', Pettit shows that it is possible for the group to make a decision that is truly its own, separate from the preferences of the individual members. Such group level decisions arise through a particular kind of

'aggregation function' (List & Pettit 2011) or decision-making process where the group must reach a decision, drawing on members' preferences but remaining loyal or consistent with the group's previous decisions or preferences. The separation between majority individual preferences and the group's preferences can be seen in the discursive dilemma, as represented by a mocked-up version of this dilemma.

Here, several committee members, (1, 2 and 3) need to make a decision about the suitability of an applicant for membership on the committee. The committee members weighed up their options on the basis of three considerations or premises. Does the candidate have the right qualifications? Do they have enough experience? Did they interview well enough? Below are the answers arrived at by each group member:

Table 1. Group and Individual Preferences in the Discursive Dilemma

	Qualified?	Experience?	Interview?	Appoint Candidate?
1	Yes	No	No	No
2	No	Yes	No	No
3	Yes	Yes	Yes	Yes
Group	Yes	Yes	No	Yes\No

In the table we can see that if the group takes a majority vote, yielding a tally of what each committee member prefers for each premise, then the decision will be to reject the applicant. However, this does not seem to fit with what the group prefers for each consideration. For each premise the group, as a majority, prefers to appoint the applicant. To satisfy the group's preference, and therefore exercise its agency, the committee members should appoint the applicant, and disregard the majority voted preference for rejecting the applicant. To do this, the group will take a premise-driven decision making approach (Pettit 2003; Pettit & Schweikard 2006).

If groups adopt a premise-driven procedure then it is possible for the group to make a decision in its own right, independently of each individual's own preference. A group that can do this to maintain consistency across decisions, and across premises, is operating with its own mind. This shows that it is possible for there to be a difference between what the majority of individuals personally endorse and what the group endorses. While it may not occur in every situation for all groups, for a joint action to be a case of group agency, the group needs to be willing and able to endorse the group's preference at the expense of the majority view, in the interest of maintaining consistency between the group's decisions across time (Pettit, Pettit & S 2006).

I am in part sympathetic to Pettit's efforts at separating out the group's cognition (or intentional profile) from the individuals' cognition (or intentional profile). As we saw at the start of this chapter, the core difficulty for a proponent of group mind/cognition is to show that group cognition is somehow different to individual cognition, such that explanations of group intentionality or group cognition are not superfluous. Pettit's case of premise-driven judgment aggregation achieves this. Cases of this kind are likely to include committees, courts, offices and, on some occasions, maybe less formal groups like families or poetry groups. But what of other kinds of groups that, at least on the surface, appear to behave intelligently? Many groups, who may not obviously use a formal judgement aggregation procedure to make a decision, seem to be capable of acting intelligently. Sports teams and jazz ensembles appear to pursue specific goals, respond selectively to the enivornment, adapt their behavior and process information across the group. Thankfully, on a mechanistic account it is still possible to avoid the problem of superfluous group explanations. This is because, as outlined above, for a group to be a genuinely cognitive group the individual group members need to be cognitive interdependently such that each individual's cognition and behavior is shaped and influenced by the particularities of their interdependencies with group members. On this view, group explanations are not superfluous because although the

cognition of the individual group members needs to be explained, so too does the way in which particular group members shape and influence each other. Adopting a mechanistic characterisation of group cognition broadens the numbers of cases that can be considered as thinkers or minds in their own right beyond the kind of judgment aggregation cases Pettit put forward as candidates for a group mind. In Chapters Four and Four, I draw on empirical evidence from organisational psychology, sports psychology and cognitive psychology to explore the different ways that sports team members can be cognitively interdependent, thus making the teams' behavior emergent and thus novel.

So far we have seen that a mechanistic account of groups provides a response to the problem of the reality of group cognition. In cases where group members are cognitively interdependent, the group is likely to violate Wimsatt's conditions of aggregation and therefore be an emergent system with new and novel properties. From Gilbert and Bratman, we get a sense of the ways in which people acting together can be interdependent. In Gilbert's case, this is achieved through both or all those people acting together in the group shifting from quasi-readiness to joint-readiness to act together. We also saw that it was possible, through the interactions of poetry group members, for a group-endorsed preference to emerge that was not endorsed by the majority of group members. In a similar vein, Bratman identifies further internal, mental properties present for each individual case of genuine joint action. These are intentions to act together that are persistent only so long as they are shared by each of those members that is acting together. Interrelated intentions and intermeshing sub-plans are further means by which two or more people can be interdependent. For both Gilbert and Bratman however, there is no discussion of the kinds of immediate, moment-to-moment processes that sustain successful joint or group action, especially in cases such as sports where the joint action can be unpredictable. In the following chapters, I explore different forms of cognitive interdependence, informed by empirical science, that are present not just in cases of pre-planned group cognition, but also in more fast-paced, on the fly forms of joint

action typical of sports. Finally, we saw from Pettit that it is possible to find a clean break between individual cognition and intentionality, and the group's cognition and intentionality. While these cases satisfy Pettit's quest for a discontinuity between the group and individuals, a mechanistic account allows for a more inclusive view of group cognition, including cases where judgment aggregation procedures are not the dominant form of decision-making.

Mechanistic emergence allows us to pick out cases of group cognition where there is no discontinuity between individual cognitive states, and the group's cognitive state. While Pettit's account holds that groups only need to be *willing* to select group preferences over individual preferences, rather than doing so in actuality, mechanism can also speak to cases where such a discontinuity may never occur. This is because the way in which the continuity is achieved may involve some kind of cognitive interdependence between group members. Being part of a group may transform the degree to which an individual endorses a preference, or it may result in an individual endorsing the preference in their capacity as an individual person (an 'unorganised part' if you will) or as a participant, *à la* Gilbert. To identify this kind of interdependence and transformation, the group cognition theorist needs to track the processes of group decision-making, and each individual's contribution to the group's behaviour as a whole, which is akin to looking down, around and up, in mechanism's parlance. Such decision-making processes may be fast and unspoken, and may be realized by lower-level cognitive processes like the kind described in Chapters Three and Four. Mechanism therefore broadens the category of group cognition to cases of overlapping cognitive states between the individuals and the group, where there is cognitive interdependence between the parts.

5. Overcoming the Reductionist Objection

The idea of groups as organised wholes is appealing because it fits neatly with work in social ontology and, as we will see in subsequent chapters, empirical work in the cognitive sciences.

Up to this point, though, the skeptic might view the mechanistic characterisation as an explanatory luxury that might not really add anything to pre-existing analyses of groups and joint action. We can now turn to the real benefit of mechanism: the way in which it allows one to deal with claims of reductionism that crop up in both social ontology and as objections to the notion of socially distributed cognition. By characterising groups as organized wholes we can see that reduction does not threaten the claim that groups are thinkers in their own right. The objection I want to counter is that if we can reduce group cognition to individual cognition, then groups are not real and distinct (Rupert 2011). Of course, there is an intermediate case: if we can reduce group cognition to individual cognition, groups may still be real and distinct, but only with regard to their non-cognitive properties. A mechanistic account of emergence can be applied to show the reality of both cognitive and non-cognitive group properties, but the focus here is on what this metaphysical framework can do for the reality of cognitive group properties. I want to challenge the claim that reducing, or explaining, behavior in terms of individual cognition need not lead to the conclusion that group cognition is not real and distinct. In earlier sections of this chapter I gestured at the possibility that groups are real and distinct, but it is important to address directly the problem of reduction. We can start by questioning what reducing groups to individuals means. As I see it, it can mean one of two things. We either reduce group behaviour to explanations of the cognitive mechanisms of a single individual, generalizing those mechanisms to all individuals in the group and claiming that group behavior is some kind of brute aggregation of these individual cognitive mechanisms: call this Reduction One. Or, we reduce group behaviour to explanations of each of the distinct individuals' behaviour and their interactions and relations: call this Reduction Two. On a mechanistic characterisation of group cognition, Reduction One does not apply and Reduction Two can be accommodated.

Reduction One would hypothetically involve explaining what a group does, for example a jury reaching a decision, in terms of identical cognitive mechanisms in each

individual, and some brute aggregation procedure, such as a vote. In fact, it might be possible to describe Pettit's judgment aggregation in this way. Rather than the group jointly making a decision, each member calculates their own decision, and then the majority vote wins. If the premise-driven approach is taken, then a proponent of Reduction One might argue that all that is happening here is the aggregation of each group member's preference to reach a group endorsed decision. Both forms of reduction are informed by orthodox psychology and cognitive science most notably in terms of what Fodor refers to as 'methodological solipsism.' According to Fodor, psychology and cognitive science ought to endeavour to identify the cognitive processes of a single individual without reference to anything outside the individual. We can then make generalisations based on these findings to the rest of the group, as each group member uses the same internal cognitive mechanisms. But as we have seen, central to a mechanistic account of group is that many groups to which we can attribute group cognition consist of parts, or members, that are doing diverse things, with diverse cognitive processing, yet are interconnected and interdependent. For example, groups may consist of a diversity of expertise and skills or a diversity of formal roles, as is the case for a soccer team, a family remembering together and an office of lawyers and clerks. Understanding the behaviour of a single member will not explain the behaviour of the other members, or of all the members and how it is they can jointly and successfully complete such demanding tasks. Reduction One would only be successful in cases where each of the individuals is doing similar things, using similar cognitive mechanisms and influence and interdependence between group members is minimal. These are cases of aggregation, in mechanistic Wimsattian terms, or of accidental groups, in social ontology terms. In cases of aggregation it is expected that an explanation of a single individual should generalise to the group, because each of the members are sufficiently similar. For more richly interdependent groups, the nature of the interdependence and the particularities of the members' interrelation shapes each group member's cognition. Specific examples of how this happens are described

in the following chapters.

To fully appreciate this we can look at an example developed by long-time distributed cognition critic, Robert Rupert. In objecting to groups as distributed cognitive systems, Rupert argues that the distribution of foraging animals can be explained in terms of a single individual animal's behavior and some brute or simple form of aggregation (Rupert 2011). In this example, food is distributed such that 80% is in one location and 20% is in another. Interestingly, animals foraging for this food are distributed correspondingly, with 80% in one area and 20% in the other. Using this example, Rupert argues that there is no need to posit a group mind, or a group cognitive system, in order to explain this distribution because the data can be explained in terms of single animals searching for food. For this example, I would agree with Rupert that no group-based explanation is needed. However, according to the mechanistic framework, this example seems to be more a case of aggregation, rather than emergent group behaviour. Presumably, the animals are all doing similar things and not acting interdependently when choosing where to take food from (this is similar to Searle's group of people running from the rain, as opposed to the interdependent dancers). Rather than this being a case of organised and co-operating parts, this is an example of aggregation. It is because of the lack of differentiation of the parts and a lack of integration across the animals that the groups in this example do not have novel causal powers and therefore there is nothing at the level of the group to explain. While explanations of single individuals can explain the foraging data, it seems unlikely that Reduction One can explain groups that are emergent, organised wholes and thus does not threaten the reality of group cognition. So if Reduction One is not a threat to groups as cognitive systems, when understood as organised wholes, let's see how Reduction Two fares.

Reducing to explanations of individuals and their interactions seems like a very plausible position to hold. After all, operating within a naturalistic world-view means that groups are only comprised of members and their interactions, and maybe certain artefacts, but

there is no other mystical substance that constitutes a group and that needs explaining. If we want to explain how a committee makes a decision, it might be possible to do so by explaining the knowledge of, and decision-making procedures used by, each of the committee members, as well as the ways in which information is shared across members. Working within a mechanistic characterisation of groups, this sort of reduction would not undermine the argument that groups are thinkers in their own right. Instead, we would have a case of reduction without levelling, or explanatory reduction without ontological levelling (Bunge 1977): real cognitive groups consist of two ontological levels, the whole and the parts, but the group is explainable in terms of the parts. This is plausible because the whole is comprised by the parts and to explain its behaviour we need to 'look down, around and up,' explaining the parts, their arrangement and the context of the whole mechanism (Bechtel 2009b). For example, Gilbert's poetry group could be explained in terms of each member's contribution and the ways in which the members exchanged ideas and accepted or rejected various bits of information. Each group member's diverse contribution, skills, role, and cognitions could be explained, and so too could all the exchanges both verbal and bodily or gestural, between group members. This form of reduction is compatible with a mechanistic characterisation of groups, because it explains the essential elements of a mechanism: the parts, and how they are arranged or the interactions between them. Once this kind of explanation is developed, we would get a picture of the whole group's behaviour. However, one response to this view available to Rupert and like minded objectors is that the interactions themselves are social not cognitive. It is then the burden of the rest of this thesis to show how the kinds of interdependence between group members are cognitive in flavour.

6. Conclusion

By adopting a mechanistic interpretation of groups it is possible to distinguish between cases of emergent and novel group cognition and cases of aggregated individual cognition. Not only

does mechanism work to unify disparate areas of research concerning groups and sociality, it also overcomes claims of reduction and superfluity of group cognition explanations. A robust way of distinguishing between group cognition and aggregate individual cognition shows that group cognition is a real phenomenon, existing in the world to be explained by philosophers and scientists. It has the added benefit of emphasising the need for researchers to explain the processes of interaction between individuals within the group and the ways in which these facilitate and maintain successful group performance, and shape what the group members are able to do. Having outlined the mechanistic characterisation of group cognition, we can now turn our attention to a specific domain of group cognition: sports team cognition.

Chapter Three

Team Cognition and Enduring, Higher-Level Interdependence

1. Introduction

Some of the most impressive feats of human collaboration can be seen on the sports field, especially from high-performing teams of elite athletes. Insight into how successful collaboration is made possible on the sporting field can inform our understanding of successful collaboration in other domains, such as music, performance (Stevens, Malloch, McKechnie, & Steven 2003), dance (Kirsh 2010), surgery, office-based work teams and scientific teams (Giere 2002). In this chapter, moving on from the general account of group cognition outlined in Chapters One and Two, sports team performance is treated as an archetypal case of group cognition. In sports like basketball, soccer, gridiron and water polo, teams must coordinate the actions and cognitions of multiple individuals. As with many group-based tasks, some teams are more successful at this than others, and this can even be the case at elite levels where all teams are composed of elite individual performers. Successful performance requires that different individuals, each with their own skills, intentions, patterns of action and cognition, be brought together in the right way at the right time. Given the temporal, perceptual, emotional and work pressures typical of elite level sports, this type of successful coordination can be difficult to attain. As a result, sports teams can vary in their level of expertise just as individuals can. In this chapter, a framework for explaining successful sports team performance, as an exemplar of effective group cognition, is developed.

Theoretical and empirical investigations of sports team performance are expeditions into human cognition and sociality insofar as multiple people are acting together on a cognitively demanding task. Moreoever, sports team performance offers a useful contrast to other group endeavours that are driven by more deliberative, discursive practices such as jury-

room deliberations and committee-decision-making. As foreshadowed in Chapter One, describing and explaining sports team performance creates fruitful points of contact between philosophy of mind, philosophy of sport, social ontology and the cognitive sciences, to enhance our understanding of not only sports performance, but of human collaboration more generally. This chapter focuses on the kinds of factors that are important for successful collaboration on the sporting field, particularly those that arise and are honed as the team spends more time playing and training together. Such factors are one possible result of the team's shared history, and are an enduring property of the team. More context-specific, immediate factors are added to this framework in the following chapter, but for now the concern is with the more diachronic aspects of team performance. In particular, this chapter identifies the kind of higher-level cognitive states and processes that accumulate over time through experience and enable a team of individuals to successfully coordinate their actions. These higher-level states, as I argue, are one way in which team members can be cognitively interdependent, thus making sports team cognition a real phenomenon. They may arise in teams with a shared history, but can also be present in teams coming together for the first time. These possibilities are explored throughout the chapter.

In Section Two, the account of mechanistic, emergent group cognition from Chapter Two is applied to sports teams to support the claim that sports teams are cognitive systems in their own right. In Section Two, I outline a framework for explaining the successful performance of sports teams, contrasting a team comprised of individual experts and a team that can aptly be described as an expert team to highlight the variability across teams, and the possibility that there is something about how team members interact that is causally efficacious in the team's performance (Eccles and Tenenbaum 2004). The framework is developed from existing theoretical and empirical research in cognitive and organisational psychology, as well as informal interviews, observations of athletes, and semi-structured interviews and surveys with sports teams (see Appendix 1 for full procedural details of our

two studies, one with a touch football team, the other with an OzTag team)[3]. It captures both the synchronic and diachronic cognitive life of sports teams, and emphasises the multitude of ways in which group members, or working parts, can be cognitively interdependent.

2. Beyond the Entities of Social Ontology

On the basis of the integration of mechanism with social ontology, the space has been cleared for group cognition to begin to find a foothold, as real, scientifically tractable phenomenon in the world. Recall that there is a way in which group cognition (or group behaviour) is distinct from individual cognitive properties, that is, by group members being interconnected and interdependent such that emergent properties of the group arise. We saw that good first candidates for this sort of interdependence are those entities identified by Gilbert and Bratman: joint-readiness, we-intentions, inter-meshing sub-plans and mutual responsiveness to each others' intentions. This chapter and the next builds on the previous chapter by further exploring the notion of cognitive interdependence, as it applies to sports teams. The present chapter explores various ways that co-actors can become cognitively interdependent, other than in terms of the interconnectedness of intentions. While interconnected intentions and mutual awareness are likely to be present in most sports teams, further factors and processes are likely to be efficacious in the groups' performance. Merely intending that we act together to defend our goal does not necessarily lead to a successful defence: other factors like skill, speed of the decision-making and communication can play a role. These are the kinds of factors and processes that are described in this and the following chapter in order to round out an account of sports team cognition. Bratman and Gilbert's accounts are useful for picking out some of the elementary cognitive states involved in sports team cognition, but we can further supplement their work to gain more explanatory leverage. Moving on from social ontology, some of the questions that now need to be answered are: what cognitive processes

[3] Full results from these studies are to be published elsewhere. Only those results that are directly relevant to the distributed-dynamic account of group cognition are presented here.

are involved in group performance? What are the lower-level processes that drive or realise we-intentions and mutual responsiveness? What other cognitive features accompany these? What else helps with the performance of the task? How are we-intentions updated and communicated or shared in the heat of the moment, on the fly? In this and the next chapter, I identify different forms of cognitive interdependence, those that operate as higher cognitive processes, as outlined in this chapter, and those that are at a lower-level of cognition, operating quickly and automatically. These two chapters are therefore intended to expand social ontology to sports teams as well as introducing important empirical evidence to social ontology

3. Sports Teams as Emergent and Distributed Cognitive Systems

Team performance can vary as much as individual performance even at an elite level, in ways that are surprising and unexpected. For instance, teams can be composed of individual experts, yet not succeed when they perform together. In the sporting domain we see examples of teams who look promising 'on paper' (i.e. a "team of experts") but whose performance does not meet expectations of the team officials or the fans. For instance, England's 2010 Soccer World Cup team was comprised of individual stars drawn from high performing English Premier League teams, yet as a whole the English team performed poorly, failing to make the quarter finals. Occasionally, we see a similar phenomenon in Australian Rugby League, where star players from club teams are selected to play representative Rugby League in the State of Origin, yet do not perform with their usual, expected brilliance in the representative side. Of course it is also possible for brilliant individuals to 'gel' and get things right from the start. This mix of cases raises questions as to what vital components could be missing in situations where a team of brilliant individuals does not perform well. When is the sum of the expert parts not enough, and why? What kind of extra, team-wide factors are missing?

We also see examples of teams who succeed unexpectedly. For example, at the 2004 Olympic games in Athens, in a shock defeat of the United States, Great Britain won the men's 4 x 100m relay in athletics. Strikingly, none of the athletes in the British relay team had made the final of the individual 100m, yet in the relay they crossed the line ahead of the favourites from the United States (with athletes who placed 1st, 3rd and 4th in the individual 100m final). This team could be considered an "expert team," since it is the team as a whole that is especially skilled, rather than any of the individual members. Similarly, in 2004 Greece created an upset when they won the European Soccer championship, having never won an international soccer tournament. The calibre of athletes in other teams was far superior, featuring players from the dominant club teams in the National leagues in Europe. Many players in the Greek team were not even in the starting lineups of their local club teams in the Greek national league. The Greek team's success was attributed to their "team spirit," an adage familiar to anyone who has ever been part of a sporting team or been a team supporter. But what exactly might team spirit involve? How can we make scientific sense of this popular language? Unlike the teams expected to perform well on the basis of their high achieving individual members, Greece's 2004 football team and the British relay team displayed expertise not obviously attributable to each of the individual team members.

Together, these examples raise interesting questions, the answers to which can further our understanding of how multiple people are able to collaborate successfully. Namely, why do some teams "gel" or have an apparent sixth sense? What does an expert team have that a team of experts sometimes lacks? What factors might contribute to an expert team's seemingly magical success? In this chapter, the focus is on trying to pin down the key cognitive factors characteristic of an expert sports teams, especially those factors or cognitive properties that are higher-level and enduring, potentially emerging and developing through time as the team plays and trains together. In the following chapter these more enduring factors are supplemented with more immediate, context-specific factors that facilitate group

performance.

Adding to the literature and research on group performance and cognition, I focus now on expert sports teams as examples of distributed cognitive systems, where players are cognitively interdependent and interconnected. Not only is it theoretically useful to characterise sports teams as distributed cognitive systems, it is also an informative methodological approach. Typical studies of distributed cognition (see Hutchins 1995a; Kirsh, Muntanyola, Jao, Lew & Sugihara 2009) involve microanalyses of a group's performance in "the wild", outside of the lab, as we saw in Chapter One. The group is audio-visually recorded and interviewed, usually over an extended period, providing insight into group processes. Using this method we can look for evidence of how information is shared throughout the group, and track representations between group members. For sports teams, this might include cataloguing the ways players communicate (e.g. code words) and how movement, gesture and body positioning become more meaningful for a team over time. Distributed cognition is therefore useful in understanding sports team performance both theoretically and empirically.

Earlier, I introduced the descriptive distinction between a team of experts (e.g. the 2010 English World Cup team) and an expert team (e.g. the British Olympic relay team). These cases point to the unpredictability of team success. The striking thing about both the British relay team and Greek soccer team is that the individual team members are not the best in the field, and yet the team members who made up both the opposition teams consisted of athletes who had excelled in the field. In these cases, something must have gone right in the way the team worked together since individually each player is not as good as each individual in the opposing, defeated team. As a way of explaining this kind of superior performance, whether it is that of a team consisting of non-elite[4] experts such as Britain's relay team, or a

[4] Here elite and non-elite are used to separate athletes who have achieved different degrees of success. Note that is a relative notion. Elite athletes in this context are those that are at the top of their field. In the case of the relay team, making the Olympic finals in the relevant individual event would qualify as elite compared to those who

team of elite experts where each player is at the top of their field, we want to understand the nature of the team member's interdependence. One form of cognitive interdependence that could explain successful team performance is the presence of enduring, higher-level cognitive interdependence, such as knowledge and intentions that are dependent on the way in which players interact with each other. In the following sections I discuss examples of enduring, higher-level cognitive interdependence.

In folk parlance, these emergent properties are what make the whole more than the sum of the parts. To support the aggregate-emergence distinction, we can draw on accounts of mechanistic emergence employed in Chapter Two (Wimsatt 2007; Bechtel 2008). Critically, the way in which the parts are arranged determines the overall behaviour of the system. A change in the arrangement of the parts changes the behaviour of the whole system because the parts are interdependent and mutually influence each other's functioning (Machamer, Darden & Craver 2000; Wimsatt 2007; Bechtel 2008). This interconnectedness and mutual influencing is what distinguishes an emergent system from a merely aggregate system. The behaviour of an aggregate system is not the result of interdependent component parts. Importantly, interconnectedness between component parts is a matter of degree, with some systems being more emergent than others due to the component parts being more interdependent (Wimsatt 2007).

For all sports teams, there is an obvious sense in which players are interdependent and mutually influence each other, displaying emergent properties. Team sports such as hockey, football, and basketball can clearly not be played by a single individual and require multiple individuals who have specific, differentiated roles and need somehow to communicate their intentions with each other. In this sense, the players will always be interdependent, but only in a limited sense. This is in keeping with Wimsatt's characterisation of emergence, whereby there are degrees of emergence. The more conditions a system violates, or the more

---

did not. And in the soccer example playing in the starting line up of a European club team would be elite, compared to making training squads as most of the Greek players did.

interdependent the parts are, the more emergent the systems' behavior is (Wimsatt 2007). For our purposes, this is important as it means that the more interdependent the parts are, the less likely it is that the team's behavior can be explained in terms of an aggregation of individual behavior. Instead, the interactions between team members need to be explained. The exact ways in which an expert team is more interdependent is discussed in the following section.

Broadly, the proposal is that richly emergent teams display properties that are not found in minimally emergent teams (i.e. the whole is more than the sum of the parts). As a result, explanations should proceed in terms of the interactions between team members, and therefore the team as a whole. It now remains to be shown exactly what these properties of richly emergent teams are.

In the previous section I introduced the distinction between an expert team and a team of experts, borrowed from organisational psychologist David Eccles (Eccles & Tenenbaum 2004; Eccles & Johnson 2008). The reason for invoking this distinction, and using Greece's national football team and Britain's relay team to exemplify it, is that it highlights the possibility that factors other than individual excellence can drive successful team performance. We can attribute the success to the team as a whole. As we saw in Chapter Two, this does not mean that we are attributing the behavior to a 'supermind', or to the overriding of individuals' behavior, but to the particular ways in which the individuals interact, and behave interdependently. We can speculate that in a mere team of experts, each member has expert skill relevant to the game, but something is not quite right about how the team enacts those skills together. In an expert team, we can speculate that while each member may not possess elite expert skill like those in the team of experts (although some expert teams may be composed of individual experts) they are somehow able to effectively collaborate. There is something about how team members interact that leads to their success. In what remains of this Chapter, I explore the different kinds of higher-level cognitive states and capacities likely found in an expert team that enable them to collaborate successfully. In keeping with the

mechanistic characterisation of group and team cognition as emergent, arising from interdependence between members, I emphasise the sense in which these higher-level, enduring cognitive states are a form of interdependence between team members. This kind of higher-level interdependence may arise in a team through their history of playing and training together, as team members' own skills and knowledge are transformed by other team members.

Obviously though, successful collaboration is not a necessary outcome of a shared history, as many amateur teams with a long history of playing together may struggle to perform well, even against newly formed teams. And of course, newly formed teams without a shared history can collaborate successfully, and impressively. In these cases, players perhaps rely on skill and knowledge that is generic in the sense that it is derived from players' experience with the sport in general, and in other teams in general, rather than the current team of which they are a part. In the next chapter, I explore the ways in which newly formed teams are cognitively interdependent on the basis of low-level perceptual processes. We can now explore the kind of shared factors and knowledge that make members of expert teams interdependent, enabling them to smoothly and swiftly coordinate their actions under severe time and perceptual constraints. Here I am seeking to identify tendencies, and the kinds of things that can happen, all other things being equal, such that if a shared history does bring successful collaboration, the causes and factors involved will be of the kind described below.

4. Enduring and Higher-Level Cognitive Interdependence

Mutuality and "sharedness" are essential parts of team expertise, potentially leading to the experience of "gelling," "running like clockwork" or "being on the same page." Here the term 'sharedness' denotes the cognitive states and processes held in common by all or most team members. Sharedness takes many different forms, including shared moods, emotions, intentions, goals, knowledge and skill, all of which can be important for team expertise. For

example, as we have seen, from social ontology we learn that shared intentions and a common understanding of how to achieve a particular goal are important for being able to act together (Bratman 1992; Gilbert 1992). Shared mood and emotions are also going to play a role in determining the quality of a team's performance (Totterdell 2000). In my own studies with a representative Oztag team (see also Appendix 1), the coach made the following observations about his team's performance. Commenting on a semi-final match they were expected to win but lost, he noted:

> In attack we just sort of lost our direction, lost our momentum just sort of didn't really want to be out there towards the end of the game and had no go forward whatsoever…Towards the end of the second half though I think a lot of girls, their heads went down…and I think they sort of thought they weren't gonna win the game.

This quote emphasises the way that shared moods and emotions seem to affect a team's performance. The passage suggests that the players mutually influenced each others' attitudes and moods, leading to patterns of play, especially in defense, that were hard to disrupt and change in the heat of play. This is an example of the possibility that not all emergent team cognition leads to a positive or successful outcome. In this example the team presumably meets the requirement of cognitive interdependence, as the team's overall mood or attitude could be attributed to the kind of verbal exchanges between players, the body language of players, and other forms of affect contagion, where players mutually shape each others' moods and attitudes. Given that this account of team cognition is underpinned by mechanistic emergence, it gives us a way of capturing all kinds of team cognition with varying degrees of success. In a sense, team cognition parallels individual cognition, in that teams (like people) have varying degrees of cognitive skill.

While this quote hints at the possibility of emergent but detrimental cognitive and mental properties, it raises a conceptual and methodological problem for explaining sports performance, and team cognition. The above quote is part of the coach's assessment of how the team played, and what went wrong. While this can be informative, as the coach has spent considerable time training with the team, watching their performances, and learning about and developing the team's style of play, a coach's perspective may not capture the reality of a team's performance. It can be easy for coaches, players and spectators to fall into patterns of description and faux explanation that invoke sports lore and folk notions of skill, rather than capturing the mechanics of performance. This means that tapping what is real about a team's performance, for example, why it was that at a certain point in the match discussed above the team's performance changed, can be difficult. In order to best tap what is going on during a team's performance, a mixed methodology is going to be the most useful. Qualitative data can be checked against video footage, or against lab-based controlled experiments such as those described in the following chapter. And in turn we can make sense of lab-based data, by situating it within the team's own narrative, as represented in interviews or logs from the team. The upshot is that when we are trying to explain team cognition, and group cognition, capturing as diverse a range of cases as possible, a hybrid methodology will be the most reliable. This is especially important to identifying the kinds of shared factors and states that drive team performance.

While sharedness in many forms is important, not everything is going to be shared across a team. Within a team each player may have a different skill set, different roles to play and ever different styles of playing. Despite such differences, some things need to be held in common by all or many team members. These shared or common cognitive states and factors need not threaten the claim that groups are emergent cognitive systems, rather than aggregated individuals. Although there may be homogeneity for some skills, affect, knowledge and intentions across the team, this homogeneity may arise through interactions

between team members, which initiate and sustain these throughout the group. This is what can make the group interdependent, to varying degrees, despite each group member instantiating the same plan, intention or skill. If we think back to the discussion of mechanistic emergence in the previous chapter, we can recall that the system is emergent when each part depends on the others for its own functioning. In the context of a sports team, each team member depends on the others for their own skill, intention or knowledge despite these states being similar or held in common across the group. The remainder of this chapter focuses on shared knowledge and shared skills, as examples of the important factors and states shared across a team that underpin team expertise. In Chapter Four other, more automatic and immediate forms of sharedness and mutuality are discussed.

Shared skill and knowledge in an expert team is multi-faceted. In keeping with mainstream cognitive science and philosophy of mind, cognition and action involve some form of mental representation that guides the system's behaviour. Similarly, shared knowledge and skill consists of mental representations shared across team members. The exact nature of the cognition involved in skilled action is itself an area of ongoing and unresolved debate, and is predominantly examined in terms of individual cognition, not skilled team or group action. The two extreme positions are the classicist account of skilled action, and anti-representationalist, phenomenological accounts of skilled action or 'absorbed coping.' On the classicist account, skilled action is a form of propositional knowledge that over time is accessed automatically and unconsciously, but still retains a propositional form, able to be articulated by the expert (Ericsson & Simon 1980). The anti-representational view on the other hand, as exemplified by Hubert Dreyfus, holds that expert, skilled action unfolds automatically, and unconsciously, in the absence of internalised propositional representations, like if-then rules, and is consequently unable to be articulated (Dreyfus 2002a, 2002b). So the debate amounts to a disagreement over the extent to which the mind is involved in skilled action (Montero 2010): does rule-based cognition mediate between perception and action or is

it more a matter of an individual's bodily attunement to relevant action? And is the nature of representational content propositional or otherwise? Surprisingly, neither camp acknowledges the sharedness of many skilled actions, such a team sport, overlooking the question of whether or not expert action operates similarly in an individual setting as it does when acting as part of a group or pair.

My hunch is that neither extreme in this debate is entirely right, or captures the complete nature of intelligent, skilled action. Although this has predominantly been an issue for the explanation of individual skill and an area of unresolved, ongoing debate, the distributed-dynamic account of team cognition need not be harmed by the diversity of thought in this area of research. Instead, we can welcome this debate, embracing a pluralist conception of the kind of representation involved in skillful action. It seems unlikely that all skillful action can be fully and accurately articulated, given the speed at which many actions are produced, and the added difficulties posed by folk descriptions of skilled actions that over - emphasize certain aspects of performance. But by the same token, it is also unlikely that experts have nothing true to say about their skillful actions (Sutton 2007). This creates a conceptual, empirical and methodological problem, but a juicy problem nonetheless, centred on trying to separate the true talk from false meta-talk. Relatedly, we must ask: how does an expert's self-report relate to what the athlete can actually do? While some cognitive psychologists are skeptical of the veracity of self-report on the basis of our tendency for confabulation (Nisbett & Wilson 1977) others maintain that under the right conditions self report can accurately capture aspects of a person's mental life (Bortolotti & Cox 2009) and can be used as an effective methodological tool (McPherson & Thomas 1989; Williams & Ericsson 2005). These questions are notoriously difficult, and, given such debate, it is not in our immediate interests to come down strongly in support of either side. Instead, it is possible to construct a hybrid framework of cognition in skilled performance as it applies to teams. This means that we are hostage to whatever researchers work out about individual cognition.

Nevertheless, adopting a pluralist view can be useful, especially when trying to address the problem of skilled action representation as it applies to teams, a topic which is only just surfacing.

While this account of shared skill and team cognition remains pluralistic, for conceptual convenience we can cluster different representational kinds together, to explain some of the cognitive features of shared action. Broadly, we can use three categories of knowledge or representation to explain team cognition. These are fully declarative knowledge, part-declarative procedural knowledge, and inarticulatable skill. Fully declarative shared knowledge is 'knowledge that,' and can be reported by team members. For some skills, patterns or moments of play, individual team members are able to fully explain what they did, how they did it and why they did what they did. Fully declarative shared knowledge can also consist of semantic knowledge about rules of the game or tendencies in other team members. Partially reportable procedural knowledge is knowledge that can only be articulated to a third party in part, but can be fully demonstrated through action, this is 'knowledge how.' For this kind of shared knowledge, the extent to which this knowledge can be articulated, expressed or demonstrated will vary across the team. An individual's explanation of their action may be very coarse grained, and as minimally declarative as cue words, shared imagery or idiosyncratic metaphors that scaffold team action, inducing particular joint actions and decisions. It can also involve shared understandings of embodied information such as gestures, posture and gaze. And finally, there are skills that are unarticulated. Players can demonstrate and repeat these skills, but they are shut off from conscious awareness. At best, players are aware that they can perform certain actions, but have no insight into how they are executed.

These forms of shared knowledge and skill can arise via the history of playing and training together, of acting and doing.[5] It is enduring but updatable, changeable and dynamic

[5] This is consistent with what is known about individual expertise, namely that amount (Ericsson 2006) and type (Smeeton, Williams, Hodges & Ward 2005) of practice correlate with level of expertise. Perhaps for expert

insofar as shared knowledge influences the team's action and action influences shared knowledge. However, these forms of knowledge and skill may also be found in newly formed teams that are able to successfully collaborate. This means that there can be differently-balanced solutions to the same task demands, that is, the same sport. In fleeting teams, a combination of generic shared knowledge and skill, that is, knowledge of the sport and others playing that sport in a variety of teams, may enable players to collaborate successfully. This means that teams without a shared history *can* play really well together. In newly formed teams there will be evidence of cognitive interdependence, but it may take the form of lower-level perceptual interdependence that updates higher cognitive states, like knowledge, plans and intentions, as explored in the next Chapter. The upshot of this hybrid, tri-partite clustering of representations is that when investigating team expertise experimentally or ethnographically, a variety of methods are required to tap the kind of knowledge that is available for self-report and those that are not, which requires a mix of qualitative and quantitative methods.

5. Existing accounts of shared knowledge

The view that shared states and factors, especially shared knowledge, are crucial for an expert team and the coordination of multiple individuals' action is compatible with a shared mental model account of team performance (from organisational psychology) and transactive memory theory (from cognitive psychology). A shared mental model consists of shared knowledge or information about a system and how that system works (Cannon-Bowers, Tannenbaum, Salas & Volpe 1993; Klimoski & Mohammed 1994; Hinsz, Vollrath & Tindale 1997; Reimer, Park & Hinsz 2006). In the sporting domain, an individual expert will have a more refined, detailed and complex knowledge of how the system (i.e., their particular sport) works (Ericsson & Lehmann 1996). On the above tripartite clusters of representation in skill,

---

teams similar correlations might exist. The amount and type of practice (i.e. playing and training together as a team) can for some teams correlate with the team's level of team expertise.

this kind of knowledge would either be fully declarative, or knowledge that is in part articulable but fully demonstrable.

Applying shared mental models to sports teams, Reimer, Park and Hinsz (2006) described how group information-processing theories can apply to sports teams. Similarly, Eccles and Johnson (2009) explored whether team coordination could be explained in terms of shared knowledge. The account developed herein is in a similar vein to this work, emphasising the role played by sharedness in team performance. Complementing this work, I emphasise the kind of *interdependent* knowledge that members of an expert team are likely to have. This might include awareness of the skills of other team mates, and awareness that one's own skills may change alongside particular team mates. I also note that an individual's skills can be transformed idiosyncratically by playing alongside particular team members, where neither team member is able to report this transformation, making it an example of the third type of representation discussed above: unarticulatable, but real skill. Specific examples from original research with a touch football and an Oztag team are provided to highlight the existence of this kind of interdependent knowledge.[6]

In emphasising the sharedness of knowledge, shared mental model accounts of team performance place little emphasis on *how* knowledge is actually shared and accessed by the team as a whole. On this point, transactive memory theory can be informative. Like a shared mental model, a transactive system involves knowledge held in common or shared across the group. In a transactive system, each person has their own expertise or skill set, which is especially likely in sports, but also has knowledge of other people's expertise (second order knowledge) (Wegner 1986). Although transactive memory theory has not been applied to sports teams, many aspects of this theory are highly relevant. For instance, given the heterogeneity of expertise in sports teams where different positions require different skill sets (for example the middle in OzTag is a playmaker whereas the winger is a runner), knowledge

[6] See appendix for a description of this study and the sports involved.

of each other's expertise would allow players to rapidly access the type of skill required at a given moment. The notion of combining heterogeneous skills, and tracking who in the team knows what (or who has particular skills), highlights the interactive nature of team performance. Another key aspect of transactive memory theory is that transactive systems develop and mature over time. Groups of people are better able to track and utilise each other's expertise, the longer they engage in task-relevant situations. The history shared by the group can enhance the transactive system and therefore the performance or expertise of the group as a whole. Knowledge of others' skills becomes more refined, and expectations about teammates' actions can be more reliably generated. It is not clear how quickly transactive systems can form, and so it may be possible that such an interdependency can arise in a newly formed team in the earlier stages of playing together. The processes by which this can occur are described in Chapter Four, highlighting the ways in which lower-level processes can update and change higher cognitive states like shared knowledge and shared expectations.

6. Shared knowledge and Shared Skill

In the spirit of transactive memory theory and shared mental model theory it is possible to highlight the kind of shared knowledge and skill that an ideal expert team would likely have. Table 1 is an outline of the kinds of shared knowledge and skill that would be found in an ideal team of experts and an ideal expert team. In real-world teams, the kind of knowledge and the extent to which it is shared across team members will vary. We can now turn our attention to the kinds of shared knowledge and skills that we would expect to find in an expert team, compared to a team of experts (see Table 2).

Table 2. Shared Knowledge and Skill in a Team of Experts and an Expert Team

Type of Knowledge	A Team of Experts	An Expert Team
Knowledge of the game	a) Formal rules b) Rules of Thumb - For general game play - NOT specific to other players	a) Formal rules b) Rules of Thumb - For general game play - Specific to other players
Knowledge of own skills	a) Good knowledge of own skills (including skills that cannot be explained)	a) Good knowledge of own skills (including skills that cannot be explained) b) Knowledge of how skills are transformed by others
Knowledge of others' skills	a) Limited knowledge of others' skills	a) Good knowledge of others' skills - Explicit knowledge - Implicit knowledge
Knowledge of what others know about you	a) Limited knowledge of what others know about you	a) Good knowledge of what others know about you
Skills – able to be performed, and observed, but not fully articulated	a) Skills that remain more or less constant across different games, teams and tasks	a) Skills that remain more or less constant across different games, teams and tasks, and work in combination with b) Skills that arise through or are transformed by others

Knowledge of the Game

*Formal Rules***.** As part of an individual athlete's knowledge of the game, they will

have detailed semantic knowledge of the formal rules of the game. This kind of knowledge will be similar for both the team of experts and the expert team. It would include such knowledge as offside rules, and the number of touches or tags allowed before the ball is handed to the other team. It is possible that this kind of knowledge is not causally efficacious in an expert team's shared skilled action, but it may play a role in constraining the action responses available to players at any given moment. For example, an attacker may curb their run so as to avoid being offside.

Rules of Thumb. Rules of thumb are also part of a player's knowledge of the game.[7] These include if-then rules that can govern a player's choice of action. A rule of thumb in Touch football might be: "if we get the ball in our own half, then we have to drive it (run it forward) for the first 3 or 4 touches before doing anything risky or before running from dummy half". This rule of thumb knowledge will be the same for both expert teams and teams of experts. However, in an expert team there are also likely to be rules of thumb that take into consideration the characteristics of particular team members. Rules of thumb will vary dramatically across individual team members, different teams, coaches, styles and cultures. For example, in an OzTag team, a general rule of thumb is that the attacking team will kick on the 4th or 5th tag to gain more territory. Typically a defensive player will drop back to cover these kicks. However, sometimes the defensive team fails to drop a player back to cover kicks early in the tag count. If the attacking team notice this and can easily give the ball to their best kicker, their rule of thumb will be to kick the ball early (i.e., on the 2nd or 3rd tag). If this rule of thumb is shared, the team can act on it quickly and seamlessly with little need for detailed verbal instructions. This kind of knowledge is likely to be declarative in the sense that players are able to articulate and describe the rules and how this influences their play. However, during play team members may not be consciously aware of the rules, but can instead act on them automatically, and unconsciously. Rules of thumb knowledge about

[7] The distinction between formal rules and rules of thumb or heuristics has also been made by Montero & Evans (2011) in the context of expert chess playing.

particular team members will only be found in an expert team and suggests interdependence between the players insofar as one player's skill affects another player's actions. This knowledge is likely learnt only from playing together as a team, or at least from spending time training, studying footage of or watching particular players. Importantly, knowledge of rules of the game or rules of thumb, for both kinds of teams, may be in the form of unconscious representations, which guide behaviour without being attended to. The exact representational nature of this knowledge is a difficult issue. While players will be able to verbalise these rules, and explain how they operate during play, the inner-goings on as a player is acting are much less clear. So it remains an open conceptual and empirical question.

Knowledge of Own Skills

This is the kind of knowledge that a player has of their own skills. This knowledge may be declarative (e.g., verbally expressed) or procedural (e.g., enacted) or a combination of both. A member of a team of experts will have this sort of knowledge of their own skills and expertise, as well as their weaknesses. For example, they might be able to list their strengths and weaknesses, talk about their playing style, or enact the way they side step around a defender. A player who is skilled at sidestepping for example, might explain how they draw a defender to a certain point before quickly transferring their body weight in order to suddenly change direction. They might also physically demonstrate the point at which they decide to change direction. This kind of knowledge would also include knowing that you are a ball player, a runner, a kicker, etc. For instance, a good ball player would know that they could draw the opposition into a particular pattern by looking in one direction while passing accurately in another direction. Knowledge of one's own skills will be the same for members of a team of experts as for members of an expert team. The degree of detail may vary according to individual expertise, such that the more expert a player is, the more detailed is their knowledge of their own skills.

An important element of an expert team that would not be found in a team of experts

is knowledge of how one's own skills can be transformed by other players. This might be as simple as saying: 'I change the way I weight my pass when I play next to player B,' or 'I throw a long cut-out pass if I can see player Z out wide'. In more detail it could be expressed by OzTag players in the following ways: 'If I am playing alongside Player X on the last tag, I am more inclined to pass the ball to them instead of kicking because I know they are fast and are capable of side-stepping an opponent' or, 'I can use my speed to best effect when I play with Players A and B because I know they can draw the opposition and create a space just wide enough for me to get through without being tagged'. This kind of knowledge can also exist without being fully expressible, or indeed, without being expressed at all. Instead, it could be tapped through empirical observation and experimental manipulation whereby players are tested performing certain actions first with strangers and then with fellow team mates, and the success of the action is scored.

In contrast, in a team of experts everyone will be highly skilled but the players will have limited knowledge of each other's skills. For example, in a team of experts, a player may be more inclined to take the safe option by throwing a dummy and running themselves, rather than throwing a risky pass to a player who may not be expecting it.

Own Skills they cannot Explain

Players from both a team of experts and an expert team may potentially have skills and patterns of action and movement they cannot explain. A player may know that she has these skills but cannot articulate or break down how they work. They can be observed, and are repeatable, but are performed so quickly and automatically that the player does not have conscious access to any kind of declarative or intentional state. This might be akin to the folk-coined phenomena of being in 'the zone.' Alternatively, there may even be implicit skills that she is able to perform but does not know she has. These skills exist, and can be observed. Skills of this kind might include magic moments such as sidestepping multiple players. Members of both a team of experts and an expert team will have these kinds of skills.

However, members of an expert team, but not a team of experts, should also have knowledge that their skills change when playing alongside particular other players, or as part of the expert team, but cannot explain the ways in which these changes occur. Their capacity for expression may be limited, and involve merely identifying that something different happens. Some possible examples might include – 'we have this uncanny ability to beat a defender,' or 'things feel more organised when we defend together,' and more obvious sentiments such as 'I'm able to do things I otherwise could not do, I play better with player X.' The types of skill that cannot be explained by an expert an interesting empirical question, and may even vary from team to team, and between individual experts.

Knowledge of Others' Skills

In a team of experts, knowledge of other team members' skills will be minimal, when compared to the kind of knowledge acquired from acting with specific other players. Knowledge of other team members' skills might be obtained by having observed particular teammates playing, rather than having played with them. For the expert team, though, knowledge of other team members' skills will be akin to second order knowledge in a transactive memory system. This knowledge will be specific and detailed and will include knowing how another player's skills can be utilised to change one's own performance (this is obviously connected to the aforementioned knowledge of a player's own game). It can facilitate the mutual and deliberate adjustment of one's own behaviour to a teammate's behaviour. For example, in OzTag, if I know Player A likes to run in one direction and then step back off her left foot and throw a short pass to her right, I will angle my run so as to expect this particular pass from her. Similarly, in Touch football, if I know that Player Z is quick off the mark and can run well from acting half, then I will play the ball if I am next to this person (rather than pass the ball to them), to allow them to run from acting half. With this knowledge it is then possible to coordinate actions and use the skills of a particular player to attain a particular outcome. This illustrates the interdependence in a team where one person's

skills affect what another is able to do, and vice versa.

Knowledge of What Others Know About You

This kind of knowledge will be present in an expert team but not in a team of experts. In an expert team it builds on the knowledge of other teammates, by adding metaknowledge about who knows what about who. This kind of knowledge can help coordinate actions in a dynamic, fast-flowing game where there is little time to verbally express one's intentions and plans. To illustrate, player X knows that player Y likes to throw long balls to her because she can use her speed to get outside her defender. So player X has knowledge of what Player Y knows. In a reciprocal fashion, Player Y knows that Player X will be expecting her to throw a long ball. So Player Y also has knowledge of what Player X knows. Moreover, Players X and Y both know that this long ball is a good option because it has worked sometimes in the past, and the current situation is similar enough to previous situations. Based on this common knowledge there is little need for verbal communication. This type of knowledge helps establish expectations of what is likely to happen in particular circumstances. In some instances, bodily and gestural cues may be sufficient to communicate an intention to perform an action, such as pointing, or moving one's body in a particular direction or glancing and turning one's head to indicate the direction play will move.

This constellation of shared knowledge and skill can be subjected to empirical scrutiny. We can take each kind of knowledge or skill and see how shared it is. We can ask players to describe set plays or footage from a game and explain what they and their teammates were thinking. We can compare verbal memory reports to actual footage of the team's performance, and look for moments that surprise players or that they cannot explain. And then to get a sense of what is shared in the team, answers could be compared. Furthermore, as a way of beginning to work out whether reportability is indicative of a player's level of expertise, comparisons can be made between the amount of detailed description provided by players with different levels of expertise (McPherson & Thomas

1989; McPherson & McMahon 2008). While the exact nature of the cognition involved in skilled action is still greatly contested for individuals and, by extension, teams, it is possible to speculate about the various forms of knowledge and skill shared by an expert team.

In summary, members of an expert team are likely to have knowledge of how one's own skills are transformed by others, detailed knowledge of other players' skills, and knowledge of what others know about them, as well as skills they cannot explain that are transformed by playing alongside specific other players. This extra knowledge, which is typical of an expert team, is a sign of interdependence. It connects one player to the other players in an expert team, making the team's performance emergent. In a real world setting, there may therefore be a correlation between the amount of detailed shared knowledge and skill a team has and the degree of interdependence between players, such that the more knowledge and skill players have that is specific to their particular team, the more interdependent the players are. And, based on the metaphysical view developed in Chapter Two, interdependence between players or working parts may be seen to give rise to emergent properties, in this case, cognitive properties, attributable to the group as a whole.

As we saw earlier in this chapter, the successes of expert teams such as Britain's relay team or Greece's national football team are instances of something more than individual excellence leading a team to success. The way in which this is achieved is through the sharing of relevant forms of knowledge and skill. While this kind of sharedness can arise in teams that have a history of playing and training together, a shared history is not necessary for emergence or for team expertise. Indeed, teams with a shared history may be hugely unsuccessful. They may be interdependent, but this could be an unhelpful interdependence, such as shared negative attitudes or mood arising in teams that have played through a whole season losing most games, for instance. Even a newly formed team can develop shared knowledge and shared skills, possibly operating with generic knowledge of the game and knowledge of playing with other teams, that become refined and more idiosyncratic as the

team plays together. The way in which shared knowledge and other higher cognitive states are updated and changed on the fly is addressed in the next chapter.

Based on the knowledge constellation outlined above it is possible to speculate as to how the coordination of multiple players' actions can be achieved under severe time and perceptual pressure. Of course, the relationship between such knowledge and coordination needs to be more robustly investigated, but we can get a general sense of how the relationship might operate. A detailed knowledge of the skills of other players assists verbal and non-verbal coordination by creating expectations of how a teammate is likely to behave in a given situation, without the need for a detailed discussion of how best to proceed. Moreover, knowledge of what others know about you can also reduce the need for laborious, detailed verbal communication and may perhaps facilitate fast paced, non-verbal coordination. For example, during the course of my ethnographic research an OzTag player said "sometimes I see where (Player X) is looking and then I notice her look at me so I know I should run at the gap where she looked." Here it seems that because both players share knowledge about what the other knows and can do, their posture and gaze have meaning that enables fast, non-verbal coordination. More specific examples of non-verbal, situated coordination are discussed in the subsequent chapter.

7. Finding shared knowledge

To find out what kind of knowledge team members had about the game and the team, the extent to which this was shared, and the kind of things that players were able to articulate, in collaboration with Dr Rochelle Cox I developed a semi-structured interview to run with a team of amateur Touch Football players. Dr Cox was both a player in the team, and an investigator making her an informed participant. We interviewed 7 amateur, non-professional touch football players (3 male and 4 female) from a team of 11 who play in a weekly social competition. Our procedure was based on a course-of-action analysis (Theureau 2003), where

participants separately viewed four segments of game footage and were asked to discuss their on-field behavior and experiences at the time. There were two attacking clips and two defending clips, and players were asked questions about the on-field play, including being asked to identify whether something was set-play or not and, if it was, to explain what it was called, whether or not it was planned, and if so, how, and what they were thinking as play unfolded. Full procedural details of this study are included in the Appendix, including all the questions we asked the participants.

Broadly, we were interested in the extent to which players had a similar understanding or shared knowledge of each of the key moments. For example**,** do players describe set plays in a similar way? Do they call set-plays by the same name? How are set-plays initiated or planned? Is there agreement about how they are initiated? And are there any differences between what the more experienced players know and what the novices know, as well those who have played together longer?

Attacking 1. For this clip, all players recognized that a set-play was performed and all players say that it was planned in advance. All players claimed that the set-play was initiated via a code word and that the male players who were involved in the set-play initiated it. Five experienced players (including Cox) said the code word was "City". A less experienced player, Linh, said the code word was "S-Auto" and another less experienced player, Amy, did not know the code word, although Amy thought that a code word would have been used. Three of the experienced players also identified an error in this play, while the other players did not mention the error.

*Attacking 2***.** For this clip, all players recognized that a set-play was performed and all players said that it was planned in advance. Six players said that the move was initiated using a code-word. And the least experienced player, in terms of playing touch football and playing with this particular team, thought it was initiated on the fly by "just knowing" what to do. Four players identified the correct code word, and of the two least experienced players, one

did not initially know what it is called, before eventually identifying the code name, where the other player did not know what it was called.

Defending 1. For this play six players spontaneously mentioned the presence of the National representative player in the opposing team. Four players recognised that the defence has spread out prior to the try: "defence gets spread out and holes form"; "they made our defence slide from one end to the other". Most players identified that their team's mistake was not 'dropping short' properly or at the right time. Six players suggested a lack of communication was a problem.

Figure 1. A still from the Defending 1 Clip. The participating team is the team defending the opposition's attacking play.

Defending 2. For this play, 5 players identified that Wayne had made a mistake. However, each player explained the mistake differently. Less experienced players offered more general and obvious explanations, for example, Amy stated that "Wayne got beaten," and Linh suggested that "Wayne didn't make the touch in time." The less experienced players knew that he had made an error yet could not explain the nature of the error. More experienced players offered an unprompted explanation as to why he didn't make the touch. Wayne said, "I've positioned myself incorrectly," and Matty explained that "because Wayne has come in they have naturally got an overlap." Five players mentioned that the defence broke down due to a lack of communication.

While this is only a small study, working with one team, and focusing on a single

game, there are still some striking features of this particular team. Obviously, we need to be careful when making generalizations from this team to sports teams in general, but the pattern of answers we got from the team members are promising, especially given that the performance of sports teams is a developing area of inquiry. One of the initially striking features of the data we collected is the variety of things that team members seem to need to have knowledge of. For example, codes names for plays, who plays what position and when, who plays what kind of role in each play, what each player should do in each position and who should initiate each play.

This is in keeping with the kinds of shared knowledge outlined above. From our touch football work it is evident that players need shared knowledge not just of the rules of the game or of general touch football practices, but knowledge specific to the team itself: the code names of each set-play, whose role it is to call the play, who does which part of the set play, and who has to communicate to who. As we can see from the less experienced players' mistakes in identifying the set-plays, this kind of knowledge is acquired over time.

The differences between the knowledge of the more experienced and that of the less experienced players is interesting because it appears to be a matter of degree in several cases. If we look at the team's identification of errors, especially for Defending 2, there are differences in what is identified that correlates with a player's experience. All players identify that a mistake was made, but the more experienced players offer a more detailed explanation of why the error occurred. So while knowledge is shared or common across players at a general level, at a much more detailed, specific level it is only shared by those players that are more experienced both in terms of playing touch football and playing in this team.

A further suggestion in the data is that a team can have more refined and detailed shared knowledge about one part of their game than about others. For this team there is a sense that players are 'on the same page' when they are attacking, insofar as the majority of players can at least run the set-play regardless of whether or not they can report the code-

word (although most know the code words), and there is mostly a consensus that the set-plays were planned in advance, initiated using a code-word and even a consensus concerning who makes the call and who hears or needs to hear to call. In defence, however, it's a different story for this team. This is not real consensus as to what mistakes are being made and why in defence, with several players in the team describing very different styles of defence that the team is trying to run, or some players even stating outright that they are confused about the defensive style the team is working with:

> *Linh* - "this dropping short business … when you're on the field and you're playing it's actually very difficult."
>
> *Rochelle* - "our team doesn't defend drives properly … they don't rotate properly."
>
> *Amy* - "I get really quite confused when it's like hit defense or not."
>
> *Marissa* - "on offense when you're running a set play it's easier because everyone is on the same page and everyone knows what they're doing. Defense is much more … um… you're improvising a little bit."

It might be the case that the team struggles in defence simply because they are an amateur team playing in a social competition. One way forward would be to compare this team with other teams in the competition to see if this is symptomatic of amateur teams or specific to this team. It could also be the case that there are important differences between the kinds of task a team confronts during a match. Conceivably, in attack, when the team has control of the ball they may also be more in control of the play, including what happens next, where it happens and when. Perhaps in defence, at least in touch football, it is the opposition that seems to have the control. Defending may be a more reactive, improvisatory task than attacking. This would have implications for an account of shared knowledge in a team, as it

would have to be sufficiently flexible and updatable to allow the team to respond to an opposition who controls when, what and how play will unfold. So there's a sense in which shared knowledge, whether reportable or not, which is acquired through repeated exposure to the task albeit, perhaps, with slight variations over time, can foster effective coordination in the team. However, there is also a sense in which players and the team as a whole are at the mercy of what is unfolding right in front of them. In the next chapter, the enduring sharedness and mutual knowledge described so far is supplemented with more context-specific aspects of sports team performance.

8. Conclusion

One reason why expert teams can perform with such fast, fluid coordination is that players have a shared knowledge of the game and the team, which enables efficient communication and reliable generation of expectations as to what teammates are likely to do in what situation. This knowledge can accrue, and can become more detailed the more a team undertakes task-relevant activities together, although shared history does not necessarily entail shared knowledge and skill. The players in an expert team are cognitively interdependent because what they are able to do and what they are likely to do is affected by what they know of and do with other teammates. However, as we can see from the kinds of time and perceptual pressures created by a charging opposition, an expert team must also be attuned to specific, salient parts of their immediate playing environment. They must be able to collectively adapt and change their routine to new, sometimes unpredicted information. These kinds of context-specific factors are explored in the following chapter. Expert team performance, as the archetypal case of group cognition, is driven by the combination of that which the team has experienced on previous occasions, manifesting in enduring, updatable shared knowledge, and that which is unfolding immediately in front of the team.

Chapter Four

Dynamicising Interdependence: Coordination "On the Fly"

1. Introduction

Many team sports, particularly at an elite level, are extremely fast-paced. Actions must be performed swiftly and efficiently. Players have severe restrictions on the time they have to make a decision – for instance, 'should I pass around the defender or take them on?' – and decision-making must often be done with obscured, minimal perceptual information. These kinds of time and perceptual pressures are not unique to elite athletes: even amateur players are under pressure to decide and act quickly. One of the key differences between elite, expert athletes and novices is how well they perform under such time and perceptual pressures, where visual information may be limited and where there is minimal time to pick a course of action. This could also explain why many amateur teams seem frantic and unsettled, with players appearing panicked and rushed compared to the more measured, flowing performances we commonly see and expect of elite teams. Of course, the degree and kind of temporal and perceptual pressures will vary across sports. For example, there are clear differences between lawn bowls and basketball in this regard, and there are differences even across different levels and playing styles within the one sport. Despite such variability, it is still possible to identify some general features of team cognition in a changing, not entirely predictable environment. In this chapter, the focus shifts to more fast-paced team behavior and the ways in which it can be explained. Where the previous chapter dealt mostly with teams who have a shared history of playing and training together, the kinds of processes identified in the present chapter are not exclusively found in teams with a shared history and are present in fleetingly formed teams, or established teams facing novel, unexpected and fast paced task environments. By focusing on the processes that allow for effective coordination of behavior in the heat of the moment, it is possible to understand not just team cognition in

enduring teams, but also team cognition in more fleeting teams, or teams in the initial stages of playing or training together.

The key theoretical devices that have been put to work in this thesis so far are still relevant for newly formed groups and for understanding fast, fluid team performance. Newly formed teams, as distinct from accidental groups of people, will form a shared intention to undertake the task together and will be mutually aware of acting together. Further, in cases where the task is completed successfully, intermeshing sub-plans are likely to have operated. Moreover, many teams that have a shared history must also confront and successfully navigate unexpected, novel situations during play whereby context-specific factors will greatly influence the players' actions, with novel situations being absorbed into the team's enduring shared knowledge. The view being developed here is that, as with individual cognition, successful navigation of the world, in this case sports performance, is driven by the way that past experiences have shaped the cognitive system, as well as more immediate features of the current environment. Which factor is more efficacious, experience or immediate context-specific aspects of the environment, will vary depending on the kind of sport, the kind of playing style a team has, how long the team has been playing together, and most likely many other factors beyond the scope of this thesis. As outlined in previous chapters, cognitive interdependence is the defining feature of group cognition. Such interdependence, as we will see, can take a variety of forms and can be present in on the fly sports, making it possible for both newly formed teams and teams with a history to be plausible candidates for cognitive groups.

The aim for this chapter is to identify other forms of cognitive interdependence arising in the context of novel and unexpected play, including cases where teams of people play together for the first time. The chapter is cumulative insofar as the account of on the fly team coordination grows out of what is known about individual cognition in sports, which is then extended to teams of athletes. The first move is to invoke lab-based research on skill

acquisition to show that individual sports performance is deeply situated and embodied where athletes rely on minimal information from the particularities of their immediate playing environment to act. This notion is then applied to the kinds of information provided by team members, in the processes of coordinating a team's actions. Finally, drawing on research from cognitive psychology on alignment, I identify forms of cognitive interdependence that arise in groups or teams without a shared history. Building on chapter three, this chapter creates a picture of dynamic, situated and embodied cognitive interdependence between team members.

2. Dynamicising Team Cognition

When taking sports teams as an object of study, it is essential to acknowledge the great diversity of teams and tasks that exist. Within the one sport different teams may have different styles of play, different strengths and weaknesses and different histories of playing and training together. And then of course, teams will be different between sports. A 400 metre running relay is a closed sport, with generally no direct contact being made between competitors and with only a select few moments of interaction between members of the same team, contrasting with more open sports like basketball and hockey, where the event unfolds through constant interaction between two teams, and within each team. The distinction between open and closed sports is not obviously hard and fast, but more of a continuum. Thoroughly closed sports are stable and relatively predictable, where participants know exactly what they need to do and when they need to do it. For example, swimming or running a relay. In thoroughly open sports, play is continuous and there are a multitude of variables that determine how a game unfolds, such as constant interaction with an opposition, as we see in hockey or basketball. Many sports have a mix of open and closed elements, and exist somewhere on the continuum between fully closed and fully open. American football, for instance, is likely to consist of more pre-planned, set or closed plays than sports like soccer, where set plays are not the dominant form of play. An account of team cognition needs to be

sensitive to this kind of diversity.

In chapter Three, the emphasis was on the role that shared knowledge plays in successful team performance, and as one of the ways that team members are cognitively interdependent. As with shared intentions, mutual awareness and inter-meshing sub-plans, shared knowledge alone is not enough to explain team cognition across all these diverse cases. Further empirical questions remain. How might we explain the performance of teams who have only recently formed, whose shared history might be limited? Similarly, how can we explain situations where established teams confront a new, changing and unpredictable playing environment, as is typical of open sports, or an encounter with a foreign opposition, or a change to the teams' line-up? To answer these questions, we need to supplement social ontology, organisational psychology and qualitative research with lab-based research with expert athletes. This does not mean that the entities identified in social ontology, like shared intentions and mutual responsiveness, and shared knowledge, are not efficacious in facilitating group performance, but that they are likely to be realized or driven by further, lower-level factors involved in novel, unpredictable, multi-variable situations such as when the team makes a mistake, or when the opposition chooses an unexpected course of action. We need a way of explaining fast, immediate, flexible team cognition that can accommodate team members' use of internal representations of past experiences stored in shared knowledge and other higher-level cognitive states, but that captures how it is that these states can change and update. Many successful teams display some kind of shared flexibility and adaptability, a capacity to respond to a rapidly changing, not fully predictable environment, suggesting that there is something special about the way they respond to particular features of the current task rather than relying exclusively on previous experiences to guide the team's response. There is therefore a need to explain the more dynamic, situation-specific aspects of team and group cognition. This and the previous chapter culminate in a framework for explaining team and group cognition, in which the team's behaviour is caused by higher-level mental states, such

as we-intentions and shared knowledge, that are realized by lower-level perceptual processes.

So what does it mean for team cognition to be described as dynamic? And what does it mean to claim that sports performance is deeply situated and embodied? At first pass, dynamic cognition in the context of team or group cognition refers to the way in which group members mutually and continuously shape and influence each others' cognitive processing as they adapt to the evolving demands of the environment. This cognition is deeply situated and embodied in the sense that cognitive processing is constrained by each individual's bodily interaction with specific parts of their immediate environment. The feel of the environment, of the actions a person performs, and the idiosyncrasies of their own bodily movements and those of their fellow team members shape that person's perception, action and decision-making, thus making cognition situated and embodied.

To get a sense of what these concepts are intended to capture, why they are important and their usefulness, it is worth reflecting on a moment of sporting team brilliance that exemplifies the speed at which team members must adapt to each other's behaviour and to the changing environment. The moment was part of the 2005 Grand Final between two Australian club teams in the National Rugby League.[8] The play began with West Tigers' Five-Eighth Benji Marshall running the ball. He stepped around three defending players, moving from the centre of the field on a diagonal run towards the left sideline, all the while heading in the direction of the try line. As Marshall moved toward the sideline Pat Richards, his supporting winger, drifted inwards on a diagonal run, just shifting behind Marshall. The remaining unbeaten defender, preoccupied with Marshall, who was still carrying the ball, closed in on Marshall. As this happened Marshall glanced to his left, towards the sideline, as the defender moved in from the right. Despite glancing to the left, Marshall then flick passed the ball backwards but to his right, placing it between himself and the defender. The ball was then comfortably, and yet surprisingly, collected by Richards, now with open space to run the

[8] Footage of this spectacular play can be found on *Youtube* at http://www.youtube.com/watch?v=2nGVHfacYtU, accessed February 4th 2012.

ball for a try. The whole play lasted for 12 seconds, for most of which time Marshall was moving from the centre of the field to the sideline.

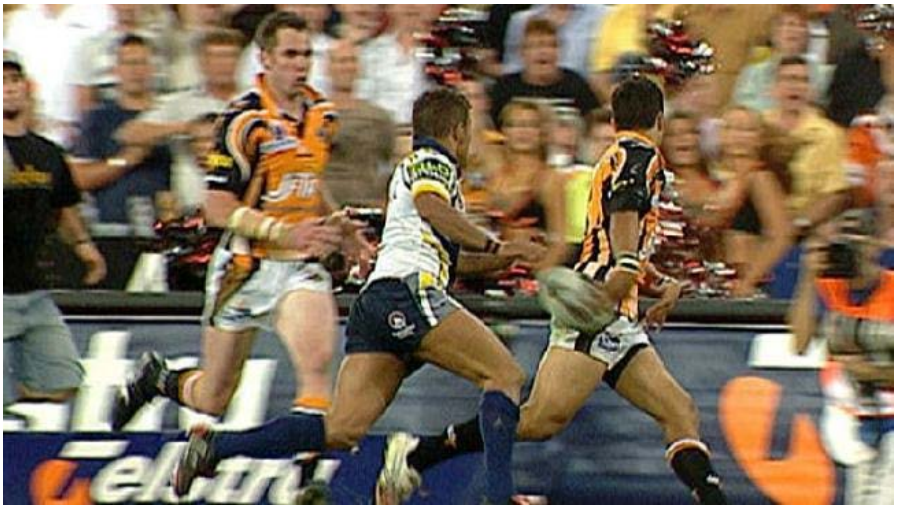

Figure 2. A still of the moment after Marshall releases the ball. Photo from *The Telegraph*, available at http://www.dailytelegraph.com.au/sport/nrl/marshall-the-man-rivals-fear/story-e6frexnr-1225914998283, accessed 2nd February 2012.

Marshall and Richards' performance is astounding for several reasons. The play happened swiftly, appearing to unravel and evolve as each defender was beaten. Moreover, Marshall is known for his individual skills of improvisation and unpredictability no doubt making him a difficult player to support. And yet, the play was a joint action, unfolding effectively because both players complemented each other's actions. How was this achieved? It is unlikely that this would have been a pre-rehearsed set play, and even more intriguing is that verbal communication between Marshall and Richards would have been limited so as to fool the defender. And yet despite the speed at which this play unravelled and the lack of explicit instruction, Richards was able to time his movements appropriately so as to support an unpredictable Marshall. These are the kinds of sporting team moments that this distributed-dynamic account of team cognition is attempting to explain. What processes and factors enable players to act together so smoothly and effectively, on the fly?

Using mechanism as the yardstick for distinguishing between cases of emergent, real group cognition has led us to look for ways in which two or more people can be cognitively

interdependent. In this chapter the forms of interdependence are less anchored in the groups' shared history, and more in their immediate task environment. While shared knowledge and history can develop through time, there are also ways in which people rapidly and automatically constrain and shape each other's cognition, on the fly and in the heat of the moment. These are the dynamic aspects of team and group cognition, whereby two or more people mutually influence each other's cognitive processing, including perception, action, and even language processing, beneath conscious awareness. This kind of interdependence is deeply situated and embodied in that the immediate behaviour of each interacting person is what shapes the others' cognition. Capturing the kind of cognitive interdependence that exists between teams and groups, necessarily involves investigating these dynamic, embodied and situated aspects of performance.

Worth highlighting here is that the notion of dynamic cognition is itself an area of study in philosophy and cognitive science, focusing on describing the underlying system dynamics of individual cognition (van Gelder 1998; Beer 2000). In the context of the theory of group cognition being developed here, the notion of dynamic cognition is employed to capture the way in which cognition unfolds on the fly, and to describe how higher-level cognitive, representational states can be updated, and shaped by newly acquired experiences. In the evidence discussed below, it seems more plausible that team cognition and team coordination in sports is, at least in-part, driven by a player's and the teams' immediate interaction with the environment. These dynamic aspects of group cognition involve action and perception being deeply coupled, unfolding together, which facilitates fast fluid action and coordination between team members. Importantly for the view developed here, the emphasis on the dynamic, context-specific aspects of team cognition does not come at the expense of a representational characterisation of cognition, either for individual cognition or team and group cognition. Instead, I use the term dynamic cognition more loosely than some dynamical systems theorists might: this particular account can be distinguished from the anti-

representationalist aspects of Anthony Chemero's dynamicist research in philosophy of cognitive science (2009) and Keith Davids' (Davids, Button & Simon 2008) dynamicist research in sports science.

Here, 'dynamic' describes the nature of multiple individuals mutually influencing each others' cognition, in the moment, on the fly. While it may be possible to describe and model team behaviour, with some form of dynamic modelling, as a self-organising system, explanatory leverage can be achieved by also explaining the kind of declarative and procedural knowledge individuals and teams rely on in the completion of a cognitive task. Thus, the view developed here is pluralist, accommodating aspects of a dynamical account of cognition, endorsing select commitments such as that of Araujo and Davids (2006) who state that 'decision-making behaviour is best considered at the level of the performer–environment relationship and viewed as emerging from the interactions of individuals with environmental constraints over time towards specific functional goals,' while also maintaining that mental representation has a role to play in cognition, for individuals and teams, across a number of domains and not just sports. With these theoretical commitments laid out, we can now turn our attention to the kind of empirical evidence that supports the characterisation of team cognition as dynamic and thoroughly situated and embodied.

3. The Ultimate Case of Situated, Embodied Cognition: Elite Individual Sports Performance

One of the most active areas of research in sports science and sports psychology is skill acquisition and expertise research. Skilled, expert performance has been intensely studied by sports and cognitive psychologists in both lab-based and field settings. Working in this area, researchers seek to find the key cognitive or information-processing differences between experts and non-exports when performing relevant sport-specific skills, and subsequently to ascertain how expert information-processing can be developed or trained. Despite the breadth

and sophistication of such research, there is very little research specifically on team performance. Research on skill acquisition and expertise in sports teams is only beginning to emerge as an area of research. This is a shortcoming in sports science, but something that is slowly shifting (see Eccles & Johnson 2009 and Reimer, Park & Hinsz). However, research with individual experts provides insight into some of the processes that are likely to be in operation when multiple individuals act together in a sporting context. From individual sports research there are two key themes that usefully inform an account of team cognition. Firstly, specific features or information in playing the environment afford certain actions for expert athletes. Expert athletes are able to detect a wider variety of useful information quickly, and in the context of changing conditions, to make predictions that afford certain actions. Non-experts, on the other hand, detect less information and are less able to predict or anticipate things like the direction of a tennis serve or the line and length of a cricket delivery. From this we learn that experts are better attuned to the environment in which they are acting, which leads to the second theme. The performance of a skill, physically preparing and moving one's body shapes how the relevant features of the environment are perceived, which in turn shapes how the action is performed. The upshot is that action and perception are tightly coupled. These two themes are exemplified by the research discussed below on individual expert sports performance. In the following section, these themes are expanded to apply to sports teams, identifying analogues of expert information-processing and the tight coupling of action and perception, as this operates in teams.

Much of the individualist research canvassed below reveals the ways in which salient, meaningful aspects of an athlete's environment are picked up rapidly and automatically by the expert athlete, and used in the production of an action response. The athlete is sensitive to particular features of the environment that lend themselves to particular actions. Drawing on ecological psychology and Gibsonian terminology (Gibson 1986), many sports scientists characterize these salient aspects of an athlete's environment as affordances or action

possibilities (Aruajo & Davids 2006). Fajen et al describe the role of affordances as '…opportunities for action provided by the environment' (Fajen, Riley and Turvey 2008, p.86). Fajen et al further note that:

> 'Affordances describe behaviors that are possible at a given moment under a given set of conditions…and allow for the prospective and moment-to-moment control of activity that is characteristic of fluent, fast-paced behavior on the playing field,' (Fajen, Riley and Turvey 2008, p. 79).

This kind of research provides evidence that individual sports cognition is driven, at least in part, by what the environment affords an athlete, suggesting that action production depends on the particularities of an athlete's immediate perceptual and bodily experience of the world. Dynamicist research in sports performance also emphasizes the way in which action and perception are coupled, whereby perception directly causes action. This is in opposition to what Susan Hurley (1998) terms 'the classical sandwich', known more generally as the cognitive sandwich, view of cognitive processing, which is most typically associated with classical accounts of cognition. On this view, cognition is 'sandwiched' between perception and action, with perception providing input for central processing and with an action response being the output of this processing. This view describes action exclusively in terms of the internal processing of mental symbols or representations, and not, as Wilson and Clark (2008) note, in terms of the individual or the environment in which the individual operates. Current sports science research rightly challenges this view by highlighting the way in which perception and action mutually shape each other. As we will see in the data from the experiments described below, this is a plausible view to take as it speaks directly to the way that expert athletes perform in the world: quickly and flexibly.

Based on this dynamicist research and its emphasis on the embodied and situated aspects of cognition, it is possible to speculate as to how this might apply to teams as a whole. I suggest that team members themselves provide affordances and action possibilities that are

idiosyncratic to that particular team member, so that not only are each member's own actions and cognitions coupled, but so too are the team members' coupled together as play unfolds. Importantly, the aspects of the environment that act as affordances, which, as we will see below, will differ depending on the sport and the task, do so in part because of the athlete's personal history, their experiences of similar sports and tasks – so that some parts of the environment become more meaningful or more useful as an athlete's experience grows and diversifies. This accumulation of experiences that renders the world more meaningful and ultimately more useful is exemplified by the differences between expert and novice participants in the experiments below, where the experience of the athletes is presumably significantly more than the novices.

Operating within the area of skill acquisition research in sports science, for several decades Bruce Abernethy and a number of colleagues have been investigating the unique information-processing capacities of expert athletes. This work is underpinned by the repeated observation that information is heavily constrained in sport because of the speed at which an individual must decide and act, and how fleeting information sources can be. Across this research, Abernethy and colleagues have established that experts are better able process the same information that non-experts process. That is, experts process this information faster and more accurately than non-experts, and use it to generate predictions. Furthermore, Abernethy et al have shown that experts pick up different kinds of information, and often more information overall, than do non-experts (Abernethy 2003). This work strongly suggests that an expert's perception is attuned in specific ways to action that they have experience of viewing.

A common way of exploring the mechanics of expert information-processing and the role played by sports-based affordances or the way in which specific aspects of the environment facilitate particular actions is through occlusion paradigms, as used by Abernethy and colleagues. These paradigms involve presenting subjects with footage from

their area of expertise, usually a single sportsperson performing an action typical of the sport, such as a kick in soccer or a stroke in tennis, but editing that footage so that various stages of the action being performed are omitted. Participants are then asked to indicate where the kick or the stroke is likely to land. These studies have been conducted by a number of different researchers, across different sports, and different skills. The motivation for using these paradigms stems from the speed at which athletes need to decide and act. Given such speeds, complex internal computations that calculate the right course of action are likely to prevent a swift response, so sports scientists have postulated that expert athletes are skilled in picking up minimal information that remains relatively constant across hugely varied, changing conditions (Abernethy 1993). The occlusion paradigms are used across many different sports as a way of identifying what minimum information is most effective. It is this minimum information that creates the affordance which guides the athlete's action. As a further step in investigating an affordance-based explanation of skill, some researchers combine tasks that identify the important information with tasks that test the relationship between the information and the action produced by the athlete, thus testing the extent to which action and perception are coupled.

The following experiments examine how this information is used and speak directly to the tight coupling between action and perception. Farrow and Abernethy (2003) conducted two experiments using a temporal occlusion paradigm to investigate the capacity of tennis players to anticipate an opponent's service and the relationship this anticipation has to the participant's action response. The first experiment compared experts to novices and the second compared experts to players with intermediate skill. For both experiments, two different response conditions were used. For the coupled condition, participants were asked to respond as they would in a real game, facing the same service. In the un-coupled condition, participants were not required to act but to verbally provide a prediction of the direction of the service. In experiment 1, participants showed significantly superior accuracy in the coupled

response condition when the ball flight was able to be seen. Experiment 2 supported this finding (Farrow & Abernethy 2003). Particularly interesting for our current purposes is that in experiment 1, across all occlusions, experts showed more robust superior accuracy in their predictions in the coupled condition, than in the un-coupled condition. These experiments provide evidence for an embodied, situated account of sports cognition, where performing the action shapes and guides an athlete's perceptual processes. For the player returning the serve, actually performing the task changes the kind of information picked up, and the way that this information is used. This study shows not only that a certain kind of information is useful to the athlete, but that it becomes useful or meaningful only through performing the action such that action is informing perception, and vice versa. So not only is an expert athlete's cognition dynamic, it is also embodied and shaped by the experience of the task.

In a similarly focused study, Mann and colleagues (Mann, Abernethy & Farrow 2010) explored the relationship between perception and action in a more fine-grained manner. Working with skilled and novice cricket batters, Mann and colleagues used a temporal occlusion paradigm. Batsmen were asked to predict the direction of the balls bowled at them across four conditions tapping different degrees of perception-action coupling. The first condition required batsmen to verbally report the direction of the ball, the second condition required the batsmen to produce an action response using only lower-body movement, the third condition allowed for full-body movement but without a bat and the fourth condition mimicked the usual game response of full body movement with the bat. The striking finding in this study was that the skilled batsmen's anticipation, that is, the accuracy of the batsmen's assessment of the direction of the ball, improved the more their bodily movement was coupled with the assessment. This was the case when no ball flight information was provided, and also when early ball flight information was not occluded. Even minimal bodily movement improved a skilled batsman's accuracy, but the most accuracy was achieved through full-body movement with a bat (Mann, Abernethy & Farrow 2010). As with the above Farrow and

Abernethy study, this study suggests that there is something gained by physically performing the task. The way in which the body and, interestingly, the bat, moves shapes a batter's perception of the direction of the ball. This is further evidence of the embodied, situation-specific nature of individual sports cognition. The challenge now is to develop these kinds of findings for an account of sports team cognition, emphasising the dynamic and embodied aspects of performance.

Using a form of the occlusion paradigm, in a series of experiments Abernethy and colleagues (Abernethy, Zawi & Jackson 2008) sought to capture the ways in which an expert athlete is able to anticipate an opponents' actions, on the basis of that athlete's 'attunement' to movement patterns observed in the opposition. Twelve expert, elite and twelve non-expert badminton players were tested, and participants were asked to predict the depth of an opponent's stroke. In experiment 1 all participants were shown a mix of film footage of an opposing player's serving action and a point-light display version of the same action. The film and point light display clips are viewed as if the participant is on-court facing the opponent. 32 strokes were shown, in both point light and film form, and each stroke was shown under five different temporal occlusions. The occlusions were 167ms before racquet to shuttle contact, 83ms before contact, point of contact, 83ms after contact and for all of the flight of the shuttle. After viewing each clip, participants predicted the depth of the shot, whether it would be short or long, and the direction of the stroke. The researchers found that experts were better able to predict the force of the stroke than the non-expert players. Experts out-performed non-experts in predicting forces across all five occlusion phases. Abernethy et al (2008) noted in particular that the experts were able to extract more information than the non-experts at each phase, especially for the first phase 167ms before contact. Interestingly, this superior performance of experts was present for both film footage and point display across all trials, with experts performing slightly better on point-light display than film footage.

In the second experiment from this study, Abernethy and colleagues sought to

ascertain what specific information players use, and if there are differences between experts and non-experts. In particular, they were concerned with whether experts used the same information as non-experts and are just better at using it, or whether experts picked up important information that the non-experts were not able to pick up or, as Abernethy et al describe it, that the non-experts were not 'attuned' to. Participants were shown point-light display footage of strokes similar to those in experiment 1, and each stroke was shown at the same five occlusion phases used in experiment 1. There were four conditions employed: under the first condition the footage displayed the motion of the shuttle and the racquet; the second condition displayed the motion of the shuttle and the arm holding the racquet; the third condition displayed the motion of the shuttle upper body without the arm holding the racquet; and condition four displayed the motion of the shuttle and the motion of the lower-body. Abernethy and colleagues found that experts were able to accurately judge stroke depth from the racquet alone, whereas non-experts could not. Interestingly, when shown the arm movement and shuttle both experts and non-experts failed to pick up information from across the first two occlusions. The experts did however out-perform the non-experts for the upper-body and lower-body conditions across all occlusion phases. This suggests that there is information that experts are attuned to which non-experts are not, and that this information is picked up quickly and early in the opponent's action. The only evidence of information pick-up for non-experts was from the upper body in the first phase of occlusion. Across these experiments, the difference between the experts' and the novices' prediction of stroke force and placement highlights those bodily features of the opponent that a badminton player relies on to anticipant the outcome of an opponent's stroke. From these results we can see that experts are attuned to and able to pick up a wider variety of information than non-experts shown the same footage, for the same duration. As this was the case across all occlusion phases, it suggests that experts pick up this information prior to seeing the opponent's action in its entirety. From this we can see that an expert experienced in responding to actions

similar to those used as stimuli in this study is attuned to different aspects of the same environment as non-experts. Presumably then, it is this fast, automatic and honed perception that produces the expert's action response.

Using a similar occlusion-style paradigm Muller and colleagues (Abernethy, Zawi & Jackson 2008) undertook four experiments with cricket batsmen of different skill levels, investigating their ability to pick up advance information on the basis of which they could anticipate the kind of balls bowled and their length by both swing and spin bowlers. As with the Abernethy at al study (Abernethy, Zawi & Jackson 2008), this study sought to identify the key information that allows batsman to anticipate the course of the ball, as part of formulating an action response. Rather than watching footage of the bowler, the batsmen wore glasses that randomly occluded various stages of the bowlers' action and the flight of the ball and had to perform a strike action for each ball. Each ball was occluded either prior to the ball being released from the bowlers' hand, prior to the ball bouncing, or not occluded at all, leaving all information available to the batsman. The contact between bat and ball was assessed to measure the quality of the batsmen's interception of the ball and the foot movements of the batsmen were recorded and evaluated to gauge the batsmen's assessment of ball length. Six of the batsmen involved were considered highly skilled, while the other six were considered low skilled. The study found that expert or highly skilled batsmen were better able to judge short ball length on the basis of information from before the ball is released from the bowler's hand when compared to the low skilled batsmen. They also showed a superior tendency to use information about the flight of the ball, both before and after ball bounce, to perform more bat and ball contacts (Muller, Abernethy, Reece, Rose, Eid, McBean, Hart & Abreu 2009). Muller et al's study picks out the information that expert batsmen use to generate particular actions, namely, pre-ball release information and pre and post bounce information. The particular actions performed by the batsmen, the nature of the contact between bat and ball, are dictated by their capacity to detect this information. Thus, the production of an expert's

interception of the ball emerges from their skilled anticipation of where the ball is likely to land. This suggests that the action-decision making process is a situated experience, whereby a batsman's actions are, at least in part, determined by their immediate perception of the bowler's actions. As well as indicating some kind of coupling between perception and action, this also suggests that the batsman's own actions are shaped by particularities of the bowlers' actions. This situation-specific environmental influence on a batsman's actions is an example of the kind of dynamic deciding and acting that unfolds as sports people perform the tasks typical of their sport. Batsmen swiftly alter their body and actions in response to the particular features of the bowler's actions. Note though that the extent to which this is dynamic is limited by the fact that batting is a relatively closed skill. There is limited mutual influence between batter and bowler, as each bat and ball contact is a closed event. In other sports, we could expect to find the actions of multiple players influencing the actions of multiple other players, such as in soccer or hockey.

Also working with cricketers, McRobert and colleagues (McRobert, Ward, Eccles & Williams 2011) employed the occlusion paradigm to investigate the role played by context-specific information. They used a simulated cricket-batting task where 10 skilled and 10 less skilled cricket batters responded to video footage of opponents bowling a cricket ball under high and low contextual information conditions. The high condition involved the batsmen being exposed to a bowler multiple times, providing the batter with more information with which to anticipate the length and course of the ball and to produce an appropriate batting action. In the low condition, batters batsmen were only exposed to the bowler for a single delivery that was the same as the last delivery shown in the high condition. Unlike the above two studies, the kind of information manipulated here is not from within a single delivery of the ball from the bowler, but across multiple deliveries. Batsmen were required to play a stroke for each ball just as they would in a real match, and after completing the stroke were asked to mark on a provided page where they expected the ball to land. Participants were also

asked to describe their thinking out loud as the bowler prepares prior to the ball being released, and then again immediately after the batter has responded to the delivery. The eye movements of the batsmen were also recorded during the bowler's actions. The study found that the skilled batsmen were more accurate in predicting where each ball would land. They also displayed more efficient gaze and search patterns, reducing the amount of time they fixated on the bowler the more they faced the same bowler. This indicates that a wider variety of information is being picked across the deliveries. Interestingly, the more skilled batsmen provided more detailed verbal reports of their thinking processes both prior to the ball being released and after they had responded to the anticipated ball. Overall, all players improved their accuracy and shifted their thinking in the high condition as they encountered more actions from a single bowler (McRobert, Ward, Eccles & Williams 2011). This study is particularly interesting for our purposes because it shows the way that a batsman's cognition attunes to the actions of a bowler the more they interact together – in particular, the way in which a batsman's search patterns are transformed by the kind of information provided by the bowler. As with the Muller experiment, here we see evidence of a batter's actions being shaped by the context-specific information they acquire from those they are interacting with. This appears to be a dynamic process, shifting and changing through time as batsmen distill the key minimal information provided by the recurring bowler, as indicated by a decrease in visual fixation time. Rather than relying solely on an internal store of appropriate responses, the batter shifts and adjusts not only their actions but also their perception in response to the particularities facing them.

To recap, from the above studies there are two central themes that emerge and inform the distributed-dynamic account of team cognition. The first is that specific features or information in the environment afford certain actions, as is the case in the occlusion based experiments where select information informs an athlete's action response. Experts with a rich history of task-related action are attuned to their immediate environment in such a way

that they can extract and use a wider range of information than non-experts to produce an appropriate action response. The second theme is that the performance of a skill, physically preparing to move and moving one's body, shapes how the relevant features of the environment are perceived. This was the case for the first two studies, where prediction accuracy was superior when participants were performing actions rather than passively perceiving an opponent. Relevant for the distributed-dynamic account is that an individual's history plays a role in how they perceive their immediate environment, automatically detecting what is meaningful and useful. Yet, this information pick-up is also shaped in part by the individual performing or preparing to perform the action as they pick up the information. This suggests that experts have an internal model, but that it operates rapidly, and usually automatically, and that it is also sensitive to specific parts of the immediate environment.

So far these themes and findings have only been applied to individual cognition, and not to team cognition. The challenge, then, is to find evidence of 1) specific features of team members that are picked up as information, ultimately providing affordances for action, and 2) evidence of situation-specific or embodied factors that change or transform each individuals' perception, or, more generally, cognition – just as physically acting shapes cognition, so too might socially acting. Finding evidence of these kinds of phenomena will help to build an account of team cognition and cognitive interdependence on the fly. Unfortunately, the possibility of team-based affordances and the embodied aspects of team behaviour have not been directly tested. However, we can add weight to this proposal through select sports science work in combination with studies of 'alignment' from cognitive psychology.

4. Social Affordances and Alignment

Building on what we learn from research on individual sports performance, two claims can

now be made here: firstly, that teammates are affordances for action based on their provision of the all important minimal information, and, secondly, that teammates are cognitively coupled or interdependent in a situation-specific manner. To support the first claim, or at least its possibility, we can draw on Steel's research in sports science. Steele et al (Steel, Adams & Canning 2006) provide evidence to suggest that athletes' perception is attuned to picking out familiar athletes based on minimal information. Working with 15 touch football players, Steel presented the players with 400msec long video clips of people running, whose familiarity to the participants varied from high to low. The footage was created using point-light displays, which meant that most of the distinctive information about the runner was omitted, leaving only the mechanics of their gait as represented by the movement of the point-lights. The runners were a mix of participants' teammates, players from opposition teams and non touch football players. Participants were asked to identify anyone who was familiar, and to rate their certainty. Strikingly, participants were significantly above chance at recognising familiar runners, including teammates and opponent players. This was despite both the short duration of the clips, and the reduction of distinct information. Steele and colleagues replicated this study with water polo players (Steel, Adams & Canning 2007). Given the way that water can obscure a player's view of other team members in this sport, Steele tested player's detection of fellow members' swimming gait on the basis of minimal visual information, using point-light display. Participants viewed footage comprised of a series of clips that were each less than 1 second in length and which depicted a random mix of team members and non team members swimming freestyle. Participants were asked to identify familiar swimmers and then rate their certainty. As with the first experiment, participants were significantly above chance in identifying those swimmers who were from their own team.

Steele's studies do not address the coupling of action and perception as some of the other sports science studies do, and they do not make a connection between action responses and the identification of a teammate under a reduced information condition. Nor does Steel

invoke the notion of affordances. What the studies do show, however, is that athletes are able to extract key information about team members, and distinguish them from non-teammates, under severe time and perceptual constraints. One way this capacity is likely to arise and develop is through athletes sharing a history together, attuning them to pick up key information from a co-actor on the fly. But given that participants could also distinguish opponents from strangers, the shared history can be very minimal, suggesting that these capacities would also be present in a newly formed team.

Presumably, the more one sees a teammate swim or run, the better able one is to detect that teammate quickly, and with little distinct information such as skin or clothing colour. It is from this minimal information that an affordance for particular actions can be gleaned. Recognising a familiar player from a non familiar player in the heat of the moment, be it through peripheral vision on the touch football field, or through the white wash of water in the pool, will afford a range of action possibilities that detection of an opponent would not. So just like isolated segments of the bowler's action in the above studies, minimal information about gait can create opportunities for action that are specific to the information being picked up on the fly. Note, though, that other players need not only afford certain action when one is familiar with that player. In teams that do have experience of playing together, players would detect idiosyncratic affordances unique to specific players. But for teams with no history, such as those newly formed, players might rely on more generic information about types of action, players or tendencies gleaned from their own history of playing in sports teams. Of course, this will shift and change through time the more familiar players become.

We have seen so far in this chapter that elite sports performance is driven, at least in part, by the fast, automatic pick-up of minimal, meaningful pieces of information. Steel's work shows us that that kind of information can be social, or can be about familiar actors. It is therefore likely that teammates provide affordances for actions that are situation specific, detected on the fly under severe time and perceptual pressures, and presumably shift as the

environment changes, creating a situated, embodied and dynamic performance. While these factors may facilitate successful team performance, they do not really meet the criteria for real group cognition as they do not provide evidence of cognitive interdependence. They do however show that expert sports performance is dynamic and thoroughly situated and embodied, so we should expect interdependence between team members to be of a similar ilk. Maintaining a focus on the context-specific, automatic, immediate aspect of sports cognition, we can draw on studies of alignment from cognitive psychology as evidence of dynamic, automatic and situation specific cognitive interdependence. Unfortunately, the following examples are not taken from sports but from orthodox lab-based studies. Despite this, it is possible to see how similar interdependencies might exist in sports teams, and this will be addressed in the next section.

Some of the most compelling evidence of the way in which people undertaking a shared task influence and shape each other's cognition comes from an area of research in cognitive psychology that has been brought together by philosopher Deborah Tollefsen under the term 'alignment studies'. In these studies, the emphasis is on the automatic, unconscious ways in which interacting with others constrains each individual's cognition. Alignment refers to the ways in which two or more people co-ordinate their cognitive and behavioural states, but which need not involve fully reportable mental states like intentions or beliefs (Tollefsen & Dale 2011, p.12). The idea is that when people act together in pursuit of a common or shared goal, they do so by affecting each others' behaviour across multiple levels of cognitive processing, from low levels of perceptual processing, to the syntactic structure of each others' speech. Evidence from alignment studies shows the various ways that individuals constrain and shape each others' behaviour, when engaged in a joint activity. These processes drive the co-actors' cognitive processing, enabling them to decide, remember and act together.

For example, experimental psychologists Richardson and Dale (2005) investigated the

relationship between speakers' and listeners' eye gaze patterns to identify the extent to which the eye gaze patterns are 'coupled' or constrained by the communication between the pairs. Four participants' speech was recorded either describing a picture of the cast from *Friends* or watching a clip of *The Simpsons*. Those participants looking at the *Friends* images were asked to "talk about the relationship between the characters" and provide their opinion of them and talk about their favourite episode. Those shown *The Simpsons* clip were asked to describe what happened in the clip and what they thought about it. The speakers were recorded and their eye movements were tracked. The remaining thirty-six participants listened to the recordings while watching either *The Simpsons* clip or looking at the *Friends* cast photo. After analysing the patterns of eye movement, researchers found that the listener's eye movements matched the speaker's, with a 2 second delay. Richardson and Dale further found that the more similar the listener's eye movements were to the speaker's, the better the listener's comprehension of the speaker's speech, when questioned about what had been said by the speaker. The significance of this is that the verbal cues provided by one person direct the gaze of the other, and constrain what information they attend to. As the speaker was not present when subjects listened to the description, it is likely that the verbalisations alone influenced the listener's attention, not gesture or posture. The study reveals one of the ways in which an individual's behaviour is shaped and constrained by another, which then shapes the other person's behaviour in turn. The verbalisations are not just carrying content about the picture but are also playing a role in constraining what the listener visually attends to, as the listeners attended to the picture or clip in much the same way as the speaker had originally. In a real world context, the expectation would be that the constraint operates on both participants in a dialogue.

In a more sports-relevant study, Richardson and colleagues (Richardson, Marsh, Isenhower, Goodman & Schmidt 2007) examined the presence of alignment in rocking chair movements. This is more sports-relevant because it deals with co-actors' perception and

action mutually influencing each others'. The researchers investigated whether interpersonal coordination would occur between two people when sitting side-by-side in rocking chairs, either intentionally or unintentionally. In the first experiment, investigating intentional coordination, 24 participants in 12 random pairs were instructed to coordinate their movement in either an inphase manner, or antiphase manner. There were two conditions whereby pairs were instructed either to fix their gaze directly ahead, seeing their partner only peripherally, or on the arm rest of their partner's chair. Participants were told to coordinate their rocking either with gaze fixed ahead or on the other's arm rest, depending on which condition they were part of. Surprisingly, Richardson and colleagues found that there was no difference in the stability of coordination between both conditions. This suggests that when coordination is intentional, the information needed to achieve it can be picked up either focally or peripherally. In order to create as close an analogue to real world social encounters as possible, Richard et al performed the same experiment again, but without instructing the pairs to coordinate their actions.

In experiment 2, 16 new participants unknown to each other were assigned randomly to pairs. So as to maintain the coordination's status as unintentional, participants were told that they were testing the ergonomics of the chair and must be tested together to save time. As with the first study there was a focal gaze condition and a peripheral gaze condition, disguised to the participants as a test of postures in the chairs. Participants wore earmuffs to block out auditory cues. Participants in each condition were asked to practice rocking, and participants were instructed to ignore their co-participant. It was found that unintentional rocking did occur for those participants who were visually coupled, with participants focusing on each other's arm rest. Coordination was achieved when the movements of the chairs correlated better than chance. This study highlights two important things. The first is that two people can coordinate their actions automatically and without conscious intent. The second is that this is done on the basis of detecting visual information about the others' movements. This is further

evidence of lower-level processes by which two or more people can shape each other's actions and cognitions, automatically and swiftly.

In a similar vein, select studies in sports psychology have drawn attention to the way in which parts of the environment constrain an athlete's cognitive processing, in much the same was as non-athletes constrain cognition in the alignment cases. Cordovil and colleagues (Cordovil, Araujo, Davids, Gouveia, Barreiros, Fernandos & Sidonio 2009) explored the decision-making processes of basketball players, arguing that decisions emerge from the player interacting with the external environment. Players were put in attacker–defender pairs, and given specific verbal instructions as to the style of play to adopt, including conservative, neutral or risk taking. When attacking, the instruction that a player was given co-varied with the time it took to cross the midline of the court. If given the conservative instruction, players were slower to cross the mid-line. Interestingly, the cues also co-varied with the speed at which the attacking player broke symmetry with the defender. While the mode of communication is verbal, this study suggests that players are not being guided exclusively by an internal, static representation of how to act, but are being shaped at least in part by environmental particulars, on the fly. As with the alignment studies, this is evidence of the way in which a jointly performed action is driven in part by moment-to-moment shaping of each others' cognitive processes and is not entirely the result of pre-formed intentions or mental models. Cordovil et al (2009) also found that the direction and timing of an attacker breaking symmetry was constrained by the height of the opponent relative to the height of the attacker. Both of these studies suggest that any internally stored rules for decision-making or any higher-level cognitive states must work alongside or be realized by more situation specific processes, such as situated information pick up. Perhaps such rules or higher-level states are causally efficacious, operating incredibly quickly, but are also sensitive to present environmental idiosyncrasies. So, the decision is shaped by external factors such as size of opponent or specific verbal cues. The player is responsive to the variation in the environment

and adapts his or her behaviour accordingly, on the fly. This is consistent with the identification of dynamic, embodied and situated factors that we saw in the sports science examples in Section Three and the alignment studies.

5. Alignment in Sports Teams

Finally, it is possible to synthesise the findings from sports science discussed in Section Three with the concept of alignment, with a view to sketching how automatic processes of cognitive interdependence fit with athletes' detection of minimal information, and the coupling of action and perception in fast-paced sports. The following sports psychology study provides a useful example to work with. Here we can see evidence that performance involves a complex interplay of the pair's and each individual's history of experiences with moment-to-moment processes of situated and embodied information pick-up. Knowledge and expectations, of the kind discussed in Chapter Three, accumulated through task experience and experience of interacting together inform moment-to-moment action: but the relationship need not be one-directional. Conceivably, as the team confronts unpredictable situations, such as returning the opponents' serve, new information can enter the team's shared knowledge. So not only does the team's history of experience play a role in the team's successful negotiation of the task environment, but new information also feeds into this stored experience altering the team's model of the task. We can see an example of this interaction between memory and situated, perceptual processes in the following study.

In research conducted with doubles tennis teams, Blickensderfer and colleagues (Blickensderfer, Reynolds, Salas & Cannon-Bowers 2010) tested the link between shared expectations of team members and implicit co-ordination between team members. Implicit co-ordination is the co-ordination of team members' actions without the use of explicit verbal instruction. The researchers hypothesised that greater shared knowledge and shared expectation is associated with greater implicit co-ordination, where shared knowledge is a

combination of task experience and expertise (in this case, time spent being coached) and team familiarity is the history of playing in the particular doubles team being tested. This is similar to the higher-level cognitive states discussed in the previous chapters, especially the interdependent, shared knowledge that can arise across a team's shared history.

The researchers surveyed 71 high performing teams to ascertain the degree of their task and team familiarity. The teams' matches were then filmed. Raters of the footage identified the degree of 'relative position' between team members. Relative position is a form of implicit co-ordination where teammates 'adjust and adapt their positioning with respect to each other's positioning during team performance and maintain positioning half a court apart' (Blickensderfer, Reynolds, Salas & Cannon-Bowers 2010, p.492). It is performed under severe time constraint and therefore cannot rely on overt communication between team mates. This is not unlike many mundane cases of joint action where the action may not be easily verbally communicated due either to time pressures or to the difficulty in describing the task linguistically. While Blickensderfer and colleagues do not explain the relationship between perceiving the movements of one's teammate and adjusting one's own positioning in response, we know from individual sports science that experts are able to use minimal slices of information, picked up on the fly while performing the action, to shape how the action is executed. And we know from the alignment studies that people are able to quickly and automatically shift their behaviour in response to the behaviour of those they are acting with, a form of low level cognitive interdependence. So, in the case of implicit coordination in doubles tennis, we can imagine that the same features will be at play. Doubles tennis players will become attuned to picking up particular pieces of information from their partner, which affords certain actions. It might be the turn of their partner's head or the angle at which their feet are placed that is picked. This information provided by the partner will then influence the production of an action response from the other partner and vice versa, in a mutual shaping of each other's performance. As with the alignment case described above, partners become

attuned to each other's cognition and action.

Importantly for the distributed-dynamic account of team cognition, when the survey and video ratings were correlated, Blickensderfer et al (2010) found that shared expectations were a predictor of effective relative positioning. So not only does the amount of tennis coaching received by players influence the team's effectiveness (individual skill is indeed important), but so too does the history of collaboration with a particular teammate, which presumably enhances the team's capacity for relative positioning or, in social ontology parlance, mutual responsiveness between teammates. It therefore seems likely that expectations of what the other group member is likely to do, and sensitivity to changes in their behaviour during play of the kind suggested by individual sports science and alignment studies, mediates an individual's own performance and leads to the successful co-ordination of joint action.

Note also that the changes in individual behaviour were not driven solely by explicit verbal communication between teammates, but also by automatically shifting body positioning relative to shifts in a teammate's positioning. The shifting of body positioning, we can speculate, may be driven by the detection of minimal information of the kind explored in the above experiments. The movements of a teammate's hips or their feet may provide information about their next move, which as perceived by their teammates, guides the teammate's own action response. Conceivably, both team members are mutually shaping each other's actions and perceptions. The idea of a relative position really nicely brings together the kinds of processes discussed so far. The adjusting of each others' positioning relative to each other, is likely achieved through the pickup of key information – for instance, the angle of a foot, the movement of a player's hips – and is subjectable to further empirical work to ascertain the kind of essential information that different teams use to act together. In this case, these pieces of key information are rendered meaningful on the basis of having experienced similar enough situations to the present one to form expectations (whatever representational

form these might take) about what one's teammate is likely to do. This information can become meaningful through a team's shared history, but could also emerge in newly formed teams on the basis of each member's experience with other team. This means that team cognition is driven not only by experience of acting together with particular team members, but also by experiences of undertaking the task with other people.

There are two important things we can take away from this example for a theory of group cognition. The first is that there is an interesting relationship between representation of past experiences similar to the present task and the task at hand that guide action but are not shut off from new experience. Instead, these shared models or knowledge are updated with new perceptual experiences. The second thing we learn is that how one person acts is dependent upon how another person acts, such that teammates are mutually dependent on each other's actions and cognitions for their own actions and cognitions. This mutual influencing is driven through the mix of perceptual information and higher cognitive information such as expectations and intentions that unfold quickly, at times beneath conscious awareness.

6. Conclusion

Across this chapter we have seen evidence of the way in which meaningful aspects of an athlete's environment shape the way they execute the relevant action. This suggests that action and perception are deeply coupled, with an athlete's behavior being shaped by particularities of their environment, as their performance unfolds "on the fly." This individual-based evidence, we can apply to team performance. The ability to rapidly pick up information from a teammate, and have this automatically shape one's own action, helps to drive successful coordination between team members. The combination of past experiences, similar to the task at hand, and particular features of the present environment can facilitate fast and fluid coordination of teammates' action under pressure and in unpredictable

situations. The attunement or alignment of each team members' perceptions and actions to one another is a form of situated, embodied cognitive interdependence. This kind of interdependence may be found in newly formed or fleeting sports teams, therefore qualifying them as instances of group cognition.

Chapters Three and Four have provided two different ways of getting at different aspects of the same phenomenon, namely, team cognition as an exemplar of fast, improvisatory collaborative action. Combined, the chapter highlighted the cognitive factors and processes that enable people to coordinate their actions on the fly. The kind of higher-level cognitive states described in Chapter Four, like shared knowledge and skills, can be realized by the lower-level perceptual and alignment processes described in this chapter. For teams with a shared history and those teams that are newly or fleetingly formed, different combinations of higher-level states and lower-level processes are likely to be present. A team with a shared history may draw on shared, common knowledge and skill, whereas a newly thrown together may be more reliant on situated and embodied factors. In both cases, there are different forms of cognitive interdependence suggesting that teams possess emergent cognitive properties.

Chapter Five

A Philosophy of Cognition for Small Groups and Teams

1. Introduction

Human groups are hugely diverse, and vary along a number of different dimensions, as do the tasks they perform. Some groups consist of a relatively small number of biologically related people, while other groups consist of hundreds of people, such as companies or political parties. Some tasks that groups undertake may have relatively few perceptual and time constraints, such as a book club's monthly meeting, while other tasks undertaken may be severely constrained by time, such as a team of paramedics administering first aid at the scene of an accident.

Given the prevalence of group-performed activities, human groups make potentially important and interesting objects of study for philosophy and the human sciences. In this spirit, developing a scientifically and philosophically plausible account of group cognition has been the overarching purpose of this thesis. In this, the final chapter, I revisit the central components of the account of group cognition developed across the previous chapters. In the first section, the explanation of sports team cognition is recapped, emphasizing the variability of teams and the synchronic and diachronic aspects of the cognitive interdependence, which given the metaphysical framework of mechanistic emergence, qualifies sports team cognition as real group cognition. The social ontology of sports teams is also explicitly addressed. Next, having revisited the account of group cognition, as informed by research on sports teams, the various strands of research used in this account are explicitly brought together creating a picture of mutual benefit between seemingly disparate areas of research. Finally, some general principles will be extracted from the mechanistic characterisation of group cognition and the case study of sports team cognition, and applied to other kinds of groups, generalizing this account of group cognition.

2. The Framework for Team Cognition

As this thesis is concerned with explaining group cognition, the focus has been on those sports that demand that multiple people coordinate their actions, namely team sports. Team sports have been treated as an exemplar of group cognition, as in many cases they exemplify the way in which two or more people can be cognitively interdependent, both synchronically and diachronically and through various combinations of enduring, higher-level cognitive states and processes and lower-level, situated processes. The task now is to provide a coherent picture of how these kinds of cognitive interdependence might fit together, thereby laying out the framework developed across the previous chapters for characterizing and understanding sports team performance.

One of the striking features of the world of sport is the range of sports that exist. Within the broader category of sports, the subset of team sports is itself hugely varied. Given such variation, a framework for investigating and explaining team cognition needs to be able to accommodate it. This variation operates over a number of levels or dimensions. For instance, there are numerous sports a team can partake in that will demand different skills of the players, from cricket to volleyball. Within a single sport there can be different kinds of teams or different playing styles, from those teams guided by pre-rehearsed routines to those that improvise more readily. Finally, within a single team members must complete different but complementary tasks, from the playmaker who sets up an attacking opportunity in touch football, to the winger who must outrun the final defender. Other variation includes doubles tennis, where the structure of the game means that opponents face each other and are never spatially intermingled, unlike soccer or water polo where players from both teams are intermixed. Other differences lie in the style of play a sport demands, as in touch football and American football which both rely on running set plays when contrasted with netball or basketball where set plays are not the dominant form of play. The nature of the sport is likely

to determine the kind of training a team undertakes, with some running through the same set plays time and again, and others providing as many opportunities as possible for the team to improvise in game-based situations. The implication that this diversity of teams and team sport has for this account of team cognition is that the extent to which higher-level and lower-level cognitive states or processes are respectively causally efficacious in the team's performance will vary depending on the sport or the team. For most sports teams, both higher- and lower-level states and processes will be involved in the team's performance with different degrees of dominance. Before discussing examples of this variation, it is worth recapping the account of team cognition on offer.

In short, the view defended across this thesis is that group or team cognition is real in cases where group members are cognitively interdependent. This is derived from the mechanistic characterisation of emergence described in Chapter Two. This entails both that these are matters of degree and that there is very often at least a small amount of group cognition in play. As we have seen, there are both synchronic and diachronic forms of interdependence, consisting of higher-level and lower-level cognitive states and processes. Drawing on social ontology and mechanism, we found a way to distinguish between cases of genuine group action or behavior and merely aggregated individual behavior. In genuine cases, group members display some form of interdependence that gives rise to the group's emergent behavior. From Gilbert and Bratman we learn that this takes the form of some kind of awareness of acting together.

However, it is not immediately clear that the problem of distinguishing between genuine group behaviour and aggregated individual behavior is present in team sports research. This is a further issue that needs to be addressed for an inclusive account of group cognition. If team sport is inherently collaborative, as I have been suggesting, then does every team qualify as a cognitive system by default? Despite this inherent collaborative nature, some sports though are designed entirely around the aggregation of individual performances.

Rally car competitions and some team-based equestrian events are considered team sports in that multiple people are responsible for the end result. But with these sports the group's performance or final results consists in the adding together of the individual results. Team members would complete the task separately, be it their lap around the course or dressage exhibition, and the time or points scored by each member is then added together to determine the team's result. In these cases, the way in which the team's performance comes about is via a brute aggregation procedure where individual contributions are summed up. We can contrast these cases with cases from other team sports that are similarly structured, such as the Ryder golf tournament or the Davis Cup in tennis, where it is common for individuals to participate separately but to experience some kind of performance and phenomenological, felt transformation through playing as part of a national team.[9] Here, unlike the rally car competition, there is some kind of interdependence between the parts that affect what each team member does, presumably by way of some form of internalization of the cognitive interdependence.[10] Of course, similar experiences of transformation may occur in some rally car or equestrian teams as well. The point here, though, is that because of the aggregative structure of the sport it is possible for there to be only a minimal form of interdependence.

If we use mechanism as our yardstick, as I have been advocating, then we can see that in cases like the rally team and the equestrian team, each person's or 'part's' behavior is not obviously directly dependent on the behavior of the other parts. There might be some minimal interdependence, for example, each team member may form the we-intention that they participate in the event together, and yet each team member may also form intentions about their own actions as a way of helping the overall team performance. Where a team's

[9] For example, in a post-match interview during the Davis Cup Argentinean tennis player David Nalbandian explained his on-court performance in the following way, "That's Davis Cup…I really like it, I play better when I'm playing for my country and today we both played a great match from the beginning and it was unbelievable for us." The quote conveys Nalbandian's experience of being transformed as a player, when playing as part of the Argentinean Davis Cup team. This quote is taken from 'Argentina has a Point to Prove', post match report by Emily Forder-White, available at http://www.daviscup.com/en/news/articles/argentina-have-a-point-to-prove.aspx, 20th February 2012.

[10] This internalisation of interdependence and transformation of players' affect is the sort phenomena picked out by the Social Manifestation Thesis, according to which certain mental properties, like particular emotional states, arise only in the presence of other people (Halbwachs 1992; Wilson 2004, 2005)

performance is the result of adding together individual performance, we would expect few if any of Wimsatt's conditions of aggregation to be violated. This would mean that the cognitive achievements of the group could be attributed to each individual's performance in isolation of the other team members' performances. However, as Wimsatt notes (2007), it is difficult to find a real world system that is not in some minimal way emergent, perhaps by violating a single cognition. On the account of team cognition proposed herein, the rally team and the equestrian team would possess emergent properties in the form of we-intentions, and inter-meshing sub-plans, to compete together as a team, but other aspects of their performance would be less interdependent. The implication of this is that some teams will be more cognitively interdependent than others across the task or sport that the team undertakes.

The kinds of cognitive independence that comprise this framework for explaining team cognition are higher-level cognitive states realized by lower-level cognitive processes. We can now revisit each form of cognitive interdependence and the relationship between the two as it applies to sports teams. From Gilbert and Bratman's conceptual research we have identified candidate higher cognitive states that drive group or joint action. For Gilbert, joint-readiness to act together is a requirement, and in being jointly-ready a pair, group or team forms a plural subject. This is fairly straightforward for a sports team. Since playing as part of a team is necessarily collaborative, it is fairly uncontroversial to suggest that members of a team will be jointly-ready to act together. What exactly this joint-readiness refers to is not so obvious. What are the players jointly-ready to do? My suggestion is that, as with most of the higher cognitive states identified by Gilbert and Bratman, there would be nested states of joint-readiness. Overarching joint-readiness to play a soccer match together subsumes sub-joint-readiness to defend a free kick together. The idea is that there are some states of being jointly-ready that refer to the whole of the task, during which further states of joint-readiness arise.

The same holds for Bratman's we-intentions. Each member of the team may possess

the intention that they train or play a match together. Nested within this intention, or stemming from this intention, would be a further series of intentions that guide players' actions while they undertake the task that was the subject of the original overarching we-intention. Of course, as we know, sports are very fast paced, so the generation and updating of these intentions occurs quickly, such that it might be difficult for players or even researchers to individuate intentions. This is where lower-level cognitive processes are important for explaining fast-paced joint or group action. What is useful about Bratman's account is his acknowledgement that intentions shift and develop as the joint or group action shifts, making his account of joint and group action suitable for a synthesis with the dynamic, changing and automatic aspects of sports performance. On Bratman's account, joint or group actions are sustained through each participant being mutually responsive to each others' intentions as they are maintained or updated throughout the joint activity. A pair of forwards in a soccer team need to be mutually sensitive to what the other is intending to do, they need to anticipate each others' actions so as to assist and complement each other by running to the right gap, or holding up the ball for just long enough for the other to move into the right space. Mutual responsiveness is likely achieved, or realized, by the kinds of lower-level processes described in Chapter Four. The automatic and rapid detection of the angle of a team member's hips or the direction their head is facing provide information about their intentions, while simultaneously producing an action in the observing team member as their intention shifts with this new information. Here, the intention and the action merge, and are distributed across both players: they intend and act together. The intention-action of each player is dependent on the intention-action of each other player involved in the team's play.

We can make the same case for the inter-meshing sub-plans picked out by Bratman. Recall that for a shared action to be successful, participants must be aware of the means by which the action or task is to be achieved. It is conceivable that a sports team would have some kind of overarching plan or strategy, within which there are nested sub-plans. In folk

terms this might be the "game-plan." The coach, manager, training staff or the team itself may develop some kind of general plan as to how certain attacking plays should unfold. In a soccer team for example, it might be that if a defender receives the ball out wide, they pass inward to a midfielder who plays the ball forward, but out wide again to the wingers. This plan might be intended to stretch the opposition's midfield, creating space. For this play to work, the midfielders need to know what the defenders are likely to do and the wingers need to know what the midfielders are likely to do. This exemplifies the importance of mutual knowledge of the sub-plans. But with sport, sub-plans rarely unfold exactly as planned. At some point the ball might be intercepted by the defence, and the team needs to adjust, update their sub-plans and intentions – in short, to improvise. One thing that still holds from Bratman's emphasis on sub-plans is the mutuality requirement. Even when the team improvises, all players need to adapt and need to know or sense what to do. As with the updating of we-intentions, sub-plans are likely to be updated or shared through low level processes of alignment, and the situated perception of essential information that guides action production, such as bodily movement or the angle of a team member's run. Joint-readiness, we-intentions and intermeshing sub-plans are higher-level cognitive states in that they can potentially be individuated, are potentially open to cognitive awareness, and are potentially able to be reported. They can be enduring, for example, we-intentions that can span and refer to an entire match, or a season, but that can also be updated, changed or abandoned over the course of the joint action.

In Chapter Three, added to the mix of higher-level cognitive states and processes efficacious in team cognition were shared knowledge and shared skills. These are similar to we-intentions, sub-plans and joint-readiness insofar as they are enduring, yet updatable, changing in response to new tasks and experiences. We saw that shared knowledge and skills can take many different forms, from semantic knowledge of the sport, to procedural, inarticulable skills tapped only through observation. As with we-intentions and inter-meshing sub-plans, the kind of knowledge and skills described in Chapter Three can be idiosyncratic to

particular team members, whereby certain procedural knowledge and skill manifests only in the context of playing alongside certain team members, or playing as part of a particular team. This kind of knowledge and skill can accrue and change over time, updated and refined on each new, but related experience.

On this account, shared intentions, shared knowledge and shared skill are formed and mutually acted upon 'on the fly,' making higher cognitive states dynamic, and sensitive to the immediate demands of the environment. This occurs through players aligning with each other, mutually shaping each other's perception and action and using meaningful chunks of fleeting information provided by team mates, the detection of which induces particular actions in the receiver of the information. When two tennis players (a doubles team) need to respond to an opponent's stroke, they need to do so quickly. In many cases the ball will not be obviously directed to a particular member of the team, but headed for a space between the two players. Somehow, the pair needs to decide who will return the shot. While one player is returning the shot, the other is not inactive. There is no time for either player to stop anticipating where the ball will go, where their partner is and where they should be themself. The non-receiving player must pick up information from their teammate about where the shot is likely to land, and where they therefore need to be to return the opponent's following return stroke. In such a short amount of time it is not possible to explicitly communicate intentions by way of detailed verbal instructions. In fact, there is unlikely to be time for simple cue words. Instead, each player's actions are produced through the detection of fleeting information from their partner's action. Information is meaningful because it is similar enough to past experiences with this player or with other players more generally, or it becomes meaningful because it stands out as unexpected. Together the team negotiates the space between the expected and the unexpected, adapting to each other and producing an appropriate and often successful action response.

At the start of this section, the variation across different sports and different teams was

highlighted. Now that the relationship between higher cognitive states like intentions, plans and knowledge, and lower-level processes of alignment and automatic information pick up has been fully laid out, it is possible to gesture at some of the ways in which this relationship might vary. One possibility is that some sports will require teams to be more dependent on pre-planned strategies than others. This is the case in American football where teams create playbooks, or collations of instructions for particular set-plays, the conditions under which each should be employed and, in many cases, alternate options for completing the play that are dependent on what the opposition does. Here, verbal cues and hand gestures will guide players' actions, along with internalized shared knowledge of the plays. We-intentions and the corresponding sub-plans, are likely to be more enduring, and stable. This does not mean that more embodied and situation-specific factors have no role to play. The information provided by team members' posture, body positioning and so on are likely to work as cues that signal a particular passage of play to be implemented. Players would then quickly "consult" their inner playbook and enact the set play as rehearsed. In cases like this it is clear that many intentions and sub-plans are formed prior to the game, or prior to particular set-plays being set in motion. The plans themselves may not change too much, but lower-level processes will act to kick off rehearsed routines.

For a soccer team, in contrast, while there might be game strategies or plans, such as relying on a certain formation of players to increase the number of attacking opportunities at vital points of the game, there tends to be less space for thoroughly planned set-plays as the sport is more open, with more evolving variables at any one time. It is unlikely that a team will rely on a playbook. Here, intentions and plans are likely to be updated and refined continuously on the bases of alignment and low level perceptual and attention processes. Of course, this variability also extends to ability levels and playing styles. Some teams may be more reliant on pre-planned patterns of play based on explicit discussion of players' skills and knowledge, while other teams may be encouraged to improvise and develop their flexibility.

The nature of the relationship between higher-level cognitive states and lower-level processes is likely to be an evolving empirical question. The proposed account is useful insofar as it can apply to a range of different sports and teams, and as we will see in Section Four, to different kinds of group. It does however give rise to a range of interesting theoretical and empirical questions. How can athletes and their teams improve their capacity to pick the right kind of information? Can lower-level processes be open to top-down influence, such as verbal instructions to change or disrupt gaze and attention patterns? How can a team's shared knowledge and shared skill be enhanced in a way that increases not only semantic knowledge but also the team's ability in performing certain skills? These questions suggest that the framework, emphasizing both the distributed and the dynamic aspects of team cognition, is fruitful with regard to further research, while also providing a theoretical foundation for understanding team cognition. We have seen how, in the case of sports teams, higher- and lower-level cognitive states and processes come together to underpin team cognition. We can now look at the bigger picture, and the kind of philosophy that emerges from synthesizing theoretical research with sports science.

3. Theoretical and Methodological Implications: A Philosophy of Small Group and Team Cognition

At the outset of this thesis, through a discussion of three possible candidates for group cognition, a soccer team, a law firm and a family, several questions central to an account of group cognition were identified, which were then to be exemplified in Chapters Three and Four by the case study of team sport. The first of these questions was the metaphysical problem of the reality of group cognition: how can we plausibly characterize group cognition as distinct from aggregated individual cognition? To answer this question, the metaphysical framework of mechanistic emergence was adopted in Chapter Two. The remaining two questions raised in Chapter One were inter-related. What are the key factors or processes that

drive successful collaboration, making genuine collaboration distinct from accidental, seemingly collaborative action? Relatedly, how does a group of people coordinate their actions on the fly, under changing and unpredictable conditions? These questions were answered in Chapters Three and Four, in the context of sports team cognition. In addressing these three questions, the distributed-dynamic account of group cognition emphasizes the different kinds of cognitive interdependence between group members, and the relationship between higher-level and lower-level cognitive states and processes, that enable the group to respond flexibly to a changing environment. The sports team case study developed across Chapters Three and Four addresses each of these questions, as an exemplar of the more general phenomenon of group cognition. We saw in the previous section of this chapter how the higher-level cognitive states and lower-level perceptual processes fit together, in the context of sports team performance. We can now take a step back and look at how the sports team case study not only helps to address the key questions raised in Chapter One, but also creates novel theoretical links between existing theoretical frameworks and areas of research. It can be seen as an example of a mutually beneficial exchange between sports science and more theoretical areas like philosophy, organisational psychology and cognitive psychology, creating a coherent philosophy of team and small group cognition.

In Chapter Two the metaphysical framework of mechanistic emergence was proposed as a conceptual apparatus for securing the reality of group cognition. As well as playing this role, that is, providing a way of distinguishing between aggregate individual cognition and genuine group cognition, it can also shed light on live philosophical issues yet to be adequately explored by philosophy of sport, sports science and philosophy of mind. By adapting mechanistic emergence to sports teams, we can make technical sense of folk claims of teams being more than the sum of the parts. For example, there is growing anticipation across many of the European football leagues that many less financially well-endowed teams are on the verge of out-performing the financially well-off powerhouses of European football.

Commentators have attributed this to some kind of group or team level property such as the 'collectivist spirit' or playing as more than the sum of the parts. Sports writer Rob Hughes has observed the collectivist spirit of particular teams in the English and German Football leagues. In the English football league little-known Norwich City has progressed through three divisions in three seasons, and in the German Bundesliga, Borussia Mönchengladbach recently beat long time favourites Bayern Munich. Similarly, in the European Champions League, Cypress' team Apoel was the first Cypriot team to make it to the play-off stages of the League, and the match commentator attributed the team's performance to the coach's ability to get the team playing as more than the sum of the parts.[11]

As pieces of popular commentary, the above statements involve no theorizing as to what a collectivist spirit entails, or what it means to be more than the sum of the parts. But based on the kind of metaphysics espoused in this thesis, it is possible to piece together a philosophically and scientifically plausible explanation of what this might involve. As we have seen, we can attribute the team's performance to the particular ways in which players are interdependent, how it is they work together, constraining and shaping each other's performance, whereby group cognitive properties emerge. By integrating mechanistic emergence with analysis of team performance it is possible to explain real world cases of sports teams acting as more than the sum of the parts. Mechanistic emergence provides us with an explanatory heuristic for investigating such performances, to make sure we capture the phenomenon in its entirety. As we have seen, this involves 'looking down, around and up,' explaining the individual players, their interactions, and how these fit together in the production of the group's behaviour. Through explanations of this kind, guided by a mechanistic characterisation of emergence, we can begin to make sense of how and why a team without a single superstar may be able to overcome an opposition with multiple

[11] See R. Hughes 'The Collectivist Spirit Rises in Europe,' *The New York Times*, January 22nd, 2012, available at http://www.nytimes.com/2012/01/23/sports/soccer/23iht-soccer23.html?_r=2&pagewanted=1, accessed February 15th 2012. See also Olympique Lyonnais v Apoel, first aired February 15th 2012, available at http://theworldgame.sbs.com.au/video/2196410305/Lyon-v-APOEL, accessed February 16th 2012.

superstars.

Furthermore, mechanistic emergence also helps us to negotiate a traditional stand-off between reductionists and emergentists that has made its way into the emerging literature on sports team performance (see Steel 2004 for an explication of how this stand-off has played out in the philosophy of science). In the philosophy of science, it has been suggested that there is a tension between reduction and emergence in that truly emergent properties, by their very nature, cannot be reduced in or by explanations of lower-level properties. This is similar to the explanatory superfluity objection raised by Rupert (2011), and dealt with in Chapter Two, where the reality of group level cognitive states is threatened by the possibility of explaining such properties in terms of explanations of individuals' cognitive states. In the context of sports team performance, philosopher Jean Fancis Gréhaigne, gestures at a similar stand-off:

> 'For Sartre, the group, from a dialectical perspective, will always be defined by action, by its common praxis, and not as a super-organism or as a collective consciousness. These last two models for understanding the group are two ways of… conferring an apparent autonomy upon it. It is important to remember that this autonomy is no more than an appearance, since the reality of the group is not, in fact, independent of the action of its members.' (Gréhaigne 2011, p.44.)

Here, Gréhaigne seems to be making a distinction between the team or group as an autonomous entity on the one hand, and a bunch of individuals' actions on the other. Readers are given the impression that a kind of reductionist understanding of a sports team, where the team is defined by individuals' action, is distinct from, perhaps even in opposition to, an emergentist understanding of sports teams as possessing group level, novel or autonomous

properties. However, as we have seen, with the right characterisation of emergence, namely, mechanistic emergence, it is possible for there to be both emergent properties at a group level, and reductionist explanations of those properties. This is the idea of 'reduction without leveling.' For a genuinely, mechanistically emergent group, there will be novel properties attributable to the group as a whole, but which only emerge through the ways in which group or team members are interdependent. To explain these novel properties, the individuals' cognition and their interactions need to be explained. This kind of explanation, or in mechanistic parlance, 'looking down and around,' is a form of reductionist explanation, but does not threaten the ontological status of group or team properties as novel and emergent. This makes the kind of metaphysical position developed in Chapter Two suitable for explaining team performance, and also necessary for clearing up philosophical problems lurking in the literature on sports teams.

This kind of benefit is reciprocal in that, as well as improving our understanding of sports teams, we can also expand the range of mechanistic emergence, adding groups of people to our existing ontology, as previously informed by mechanistic emergence. This creates a picture of nested mechanisms, from cellular and molecular mechanisms within each individual, to an individual's neural mechanisms and emergent brain regions, to groups of people as mechanisms or emergent wholes. On this view, groups of people, including sports teams, fit neatly with a biology and psychology that carves the world into emergent, organised wholes.

Not only does the sports team case study expand and reinforce the metaphysical view outlined in Chapter Two, it also provides important lessons for social ontology, organisational psychology, and philosophy of mind and cognitive science more generally. These areas can be brought together to form a coherent framework for understanding group cognition. By exploring the nature of sports team coordination, we gain valuable insight for social ontology into how collaborative actions can change and adapt over time, and how this can happen

quickly, and without explicit communication. The cognitive entities highlighted by social ontology can be usefully applied to sports team cognition, therefore broadening the kind of collaborative actions that social ontology can account for. We see from the integration of social ontology with organisational psychology and sports science that the *a priori* methodology of social ontology is compatible with empirical research, hopefully encouraging more of an interactive dialogue between social ontology and science. The case study shows that the entities posited in social ontology can find some kind of empirical credibility, but also acts as an expansion of social ontology by highlighting what other aspects of collaboration need to be addressed in order to fully understand how two or more people act and think together.

Future integrative work between social ontology, sports science and cognitive science could centre on identifying what kinds of other mental entities exist in social interactions. We saw that a sports team's decision-making can be distributed across or shared by multiple people. Further analyses could focus on what it is for a goal to be shared by a team or group of people, or how decision-making processes operate in domains other than sport. In-depth study of particular cases, such as sports teams, can help bolster the credibility of conceptual analyses. Taking Gilbert's hypothetical poetry group as an example, this could be further developed through ethnographic research of actual poetry groups, interviewing members before and after the group meets, either separately or together. Just as sports science benefits from social ontology's identification of higher cognitive states, social ontology can benefit from the kind of methodology employed by sports scientists.

We can also find similar forms of mutual benefit between sports science and organisational psychology. Presently, there is no coherent theory of sports team cognition in organisational psychology, so this case study is in part a first step towards a credible theory, compatible with and drawing on organisational psychology's work with other groups, such as work teams. By introducing sports science to organisational psychology, we are able to get a

sense of the kind of processes or mechanisms that drive group performance. Organisational psychology typically identifies stages in information-processing across a group, representing these stages with flow-charts of each cognitive state. While this kind of abstraction captures something real about a group or a sports team, the picture is much more complete when sports science evidence of the kind described in Chapter Three is added to these more abstractly described states, such as shared knowledge. Building on the synthetic projects of this thesis, one possible way forward would be to look for analogues to the kind of lower-level processes identified in sports research in other domains, such as high pressure work environments – for instance, a surgical team in an operating theatre. There seems to be little reason to doubt that lower-level processes of alignment and affordance responsiveness would not be present in non-sports domains.

For the domain of sport itself, and for sports science research, organisational psychology is a useful area to draw on as it may also provide opportunities for analyses of sports performance as a form of work and paid employment for the professional athletes involved. When sport is a form of work we can expect this to have some kind of effect on a team's performance. The natures of such effects could be ascertained through ongoing conversation between organisational psychology and sports science. Sports provide an untapped area of research for exploring fast-paced work environments. Moreover, organisational psychology, through identifying the abstract cognitive state realized by a group of people, provides a means of making generalizations across various kinds of groups and teams, acting together on a variety of different tasks from sports to office work and musical performance. When used alongside social ontology, one of the added dimensions that organisational psychology offers is to highlight the diachronic aspects of group behaviour and the fact that a group member's and the group's history shape the group's synchronic cognitive states, including intentions and knowledge. By introducing into this mix the mechanistic emergence framework, we can also make sense of how group cognitive states can be

accurately attributed to the group as a whole, rather than merely assuming that such attributions are accurate, as some organisational psychologists tend to do.

By bringing each of these areas of research together it is possible to create a far-reaching, coherent and integrative framework or philosophy for understanding team and group performance, as characterized in terms of group cognition. The once disparate areas can be brought together in ways that extend and improve aspects of each research tradition. Perhaps more importantly, the framework creates real and clear ways forward to further improve our understanding of group and team cognition.

4. From Teams to Groups

Several important things have been achieved in this thesis: firstly, a way of characterizing group cognition as real and distinct from aggregated individual cognition has been developed using mechanistic emergence. Secondly, the sense in which the emergent properties are cognitive has been explained. And thirdly, the factors and processes involved in successful, on the fly collaboration have been identified. So far, the kinds of cases of group cognition considered have predominantly been from team sport. Yet this is intended as an account or theory of group cognition in general. Therefore, it is necessary to draw some general principles from the preceding explication of this particular account of group cognition.

Crucially, group cognition is nothing but individuals and their interactions. Importantly though, each person crucially depends on the others in the group for their own cognitive processing in completing the task together. The upshot of this is that phenomena of group cognition should generate a plurality of explanatory projects. In explaining the performance of a single group, whether it is a sports team, a team of paramedics, a family remembering together or a poetry group, a variety of different methods will be required to explore the various sub-explananda of group cognition. One component of this will be individual-level explanations including theory of mind explanations, investigating each group member's own skill set and abilities and how these change across different contexts and

conditions. And then of course group-wide explanation, which will involve tracking who knows what, who has what skills or abilities and who does which task or sub-task and how these are related. This will map the distribution of knowledge, representations and skills throughout the group, and how these come together. And finally, it may also be useful to investigate the cognitive profile of sub-groups within the larger group, such as dyads or triads, to ascertain whether or not group members perform differently alongside different group members. There can also be different explanatory goals or motivations for investigating group cognition. It may be to improve a single team or group's performance, or to understand what happens to particular individuals when they work with different groups, or it may simply be to create richer understanding of cognition as it plays out in the wild.

This account of group cognition is far-reaching in the sense that it can apply to a variety of different groups. This is mostly because it captures many of the key dimensions along which a group can vary, especially those that affect the nature of the group's cognitive processing. It captures both the synchronic and diachronic aspects of group cognition. Some groups, like families, will have long and rich histories of acting together, whereas other groups, such as newly formed social netball teams, may have a limited history or no history at all. The kinds of shared knowledge and skills may therefore vary depending on the group's history. The account also captures the individual and group dimensions of group performance. This means that it is applicable to groups that violate at least one of Wimsatt's conditions. Even aspects of rally car team or equestrian team performance can be accommodated by this account. Because it is able to capture each of these dimensions, it can apply to many groups across many different domains. Of course there will be vast differences in tasks and how exactly the group performs them, but we can still make generalizations across all cases. The key generalization is that there will be some combination of higher cognitive states with lower-level processes, but to different degrees. For example, a planning committee may have a fully planned agenda that they have to stick to in reaching a final decision. For each agenda

item, though, the group might engage in unstructured discussion. Here we can imagine hand gesture, speech prosody and posture providing information between committee members that guides the decision that the committee ultimately makes.

Two further general points can be drawn from this account of group cognition. The first is that there are degrees of group cognition, wherein some groups are more cognitive than others. The second is that not all emergent properties are beneficial properties that necessarily lead to the group's success. The notion that there are degrees of group cognition refers to the possibility that, at any one time, for a group, cognitive processing can either be attributed entirely to an aggregate of individual cognition, or to a mix of aggregate and sub-group cognition, or entirely to the group as a whole entity. In cases where a group is thoroughly cognitively interdependent, the cognitive processing can be attributed to the group as a whole. This would be the case for a poetry group, a soccer team or a team of surgeons. For less interdependent groups, only select episodes might be accurately described as group cognition, where there is a form of cognitive interdependence present. For instance, the moment in which relay runners exchange the baton, in contrast to when a single individual is running their own leg of the race, is a mix of group cognition and individual cognition. An office team is another example. In some offices many tasks can be undertaken by individuals in isolation from one another, and yet, in meeting situations, the team may operate as a fully fledged distributed cognitive system, sharing not only linguistic representations but also bodily and gestural information that is detected automatically and unconsciously.

Implicit in the discussion of emergence and groups so far is the notion that not all emergent properties are advantageous for the group. This is because emergence is a metaphysical concept and not an account of how the cognition operates, or its effects. Therefore, it is possible that some properties may be emergent insofar as they emerge from the interactions and interdependencies between group members, and yet they do not necessarily promote the group's success. An example of this would be the phenomenon of

contagion as it applies to mood, affect or memory. Contagion involves the kind of dependence between two or more people that would qualify it as emergent, yet it can disrupt or negatively affect cognitive performance. As well as in terms of affect, contagion has also been associated with memory, particularly the way in which a person's autobiographical memory can change to seemingly accommodate misinformation from another person (Loftus 1979; Barnier & McCauley 1992; Roediger, Meade & Bergman 2001). This is a real possibility for groups of people remembering together. Sharing a memory can change what one recalls and reports on future tellings of the same memory. Similarly, it is possible that mood and affect can operate in a similar way. An example of this in folk speak is that of a team 'putting their heads down,' just as the OzTag coach observed in our study from Chapter Three. In these cases spectators claim to observe some kind of negative or defeated attitude spread through the team, which can affect their performance. This suggests that one important part of future research on group cognition and group performance should be to investigate how to promote beneficial interdependencies between group members, and how it is possible for groups to overcome negative patterns of mutual influencing that become grooved and entrenched overtime.

To make the move explicitly from sports teams to other groups we can re-introduce the two remaining examples from Chapter One: the legal office and the family remembering together. Based on the particular metaphysical approach defended in this thesis, for a legal office and a family to count as genuine cognitive groups, they would need to display some form of cognitive interdependence that gives rise to novel properties attributable to the office or the family as a whole. For the family remembering together, we can speculate that both higher-level and lower-level cognitive interdependence might be at play. There is a very small amount of research on some of these processes, but what exists highlights the existence of family-wide, shared long-term memories (Shore 2009) and the development of a child's self-narrative as shaped by familial interactions (Bohanek, Marin, Fivush & Duke 2006). In the

case of a family sharing a memory, we can imagine that while each family may have their own different mental representation or memory of, say, their last joint holiday, when the family discusses their memories it is conceivable that through integrating these previously existing representations a novel memory can emerge. Perhaps it is a memory that is richer in detail than each individual's own prior memory, or maybe the content is altogether new. In this case the new memory is the higher-level, emergent property that consists of the combination of other memories. We can also imagine that lower-level processes could also help with the emergence of a group memory. Certain gestures a family member uses when talking, cue words, manipulation of physical objects or patterns of eye movement may play some causal role in how the memory is shared, and emerges from the complex combination of individuals' memories. In cases like these, a family meets the criteria for a distributed, dynamic cognitive system.

We can make a similar case for a legal office. Higher-level cognitive states, such as a we-intention or a memory, may be manifest in the record keeping system implemented by the office. Email systems, white boards, and an internal phone network provide the means by which individuals can express their own intentions to undertake a task jointly with another member of staff, or to do it alone. Across time, a legal office is likely to vary in the extent to which cognitive properties can be attributed to the group as a whole. Some cognitive tasks, like deciding and delegating may be done by a senior team member alone. For other tasks, such as working out the case history for a complex matter, may involve a combination of higher-level and lower-level cognitive interdependence such as distributed decision-making or we-intentions to act together, realized by lower-level forms of interdependence such as affect contagion, and gestural and postural information sharing. Just as sports teams can be subject to theoretical and empirical investigations of their distributed and dynamic cognitive profile, so too can families and legal offices.

We can also imagine that a legal office would fit nicely with Pettit's criteria for group

agency, as outlined in Chapter Two. Recall that on Pettit's account, for a group to be an intelligent entity in its own right, it must be willing to pursue consistency of preferences, decisions and actions across time, and be willing to do so even when it means rejecting majority vote individual preferences. Here, we can imagine staff members of a legal office making decisions on the basis of a group vote about the whether or not to take on a particular new client, or deciding which applicant to appoint to the firm. In contrast to the sports team examples from Chapters One, Three and Four, the legal office is more likely to use some kind of formal aggregation procedure whereby preferences are expressed declaratively and manually tallied or aggregated to reach the group's decision. For the case of the potential new client, while each staff member may not individually approve of representing the client, the office may hypothetically have a history of representing clients of this kind, in which case the group would be consistent across time were they to accept the client. On Pettit's view, the legal staff would be acting as a group agent. However, on the kind of account of group cognition developed across this thesis, it may be possible for the group to count as cognitive even when the group endorses the majority's vote. Provided that the way in which the decision is made involves some kind of cognitive interdependence between group members, the group is cognitive. This means that the framework developed across this thesis can be applied and extended to groups in other domains.

5. Conclusion

This thesis culminates in a framework for explaining group behaviour that is able to capture human groups in all their variety. Real group cognition, on this mechanistic emergence view, involves group members being cognitively interdependent. This kind of interdependence can be enduring, existing at the level of higher cognitive states like we-intentions, shared knowledge and joint-readiness. It also exists at a lower-level in the form of automatic, unconscious perceptual processing that realizes higher cognitive states, maintaining and

updating them. In emphasizing the distributed and dynamic aspects of group cognition, this account accommodates and indeed depends on a range of research methods to further understand group behaviour. The framework is made plausible on the basis of existing conceptual and empirical research, but also clears the way for further research, having overcome the metaphysical problem of the reality of group cognition, and by identifying the kinds of processes that typically drive group behaviour. The upshot is that to more fully understand how groups are able to behave intelligently, we need to capture the richness of group member interaction with a variety of methodologies, from conceptual analysis to ethnographical field work and lab-based studies of collaboration, bringing diverse research together in exciting and novel ways.

Appendix

The Sports of Touch Football and OzTag

Touch football and OzTag (modified versions of Rugby League) are played on a grass field 70m x 50m with 6-8 players per team. Teams play in a line across the field facing each other. The object of the game is to score trys and to prevent the opposition team from scoring. The opposition attempts to touch or tag the person with the ball. Once a touch/tag has been made, the player with the ball must "play the ball". This involves rolling the ball backwards between the legs (for Touch) or rolling the ball backwards with one foot (for OzTag). Another member of the attacking team will pick up the ball and can either run with the ball or pass the ball. The opposition cannot move forward to make another touch/tag until the attacking team has played the ball. A maximum of 6 touches/tags are allowed before possession is awarded to the other team.

For both the below studies ethics approval was issued by the Macquarie University Ethics Committee (Reference: REF: HE23FEB2007-R05039).

Study One: Coordination in a Touch Football Team

The participants were of mean age 30.71 years (*SD* = 5.15), had been playing touch football for a mean 7.29 years (*SD* = 3.04) and playing as a team for a mean of 4 years (*SD*=1.63). There were three experienced players who had played touch football for more than 10 years and two inexperienced players who had played for less than four years. We selected this team because one of the authors, Rochelle Cox, was a team member. We emphasized to players that there was no obligation to participate especially given Cox's involvement with both the team and study. Written informed consent was attained from all participants and we financially reimbursed them for their time.

We filmed a weekly touch football game using a handheld video camera from the sideline, with the operator moving up and down the sideline with play. From this footage we

created four short clips that depicted key moments in the game. The key moments were selected by Cox based on her experience as a member of the team, her experience playing touch football more generally, and her role as an investigator in this study. There were two attacking clips and two defending clips. The first attacking clip (attacking 1) involved a set-play that the team enacted from tap-off (i.e., after the opposition had scored a try) and resulted in a try. The second attacking clip (attacking 2) involved a set-play that the team enacted when they had possession in their own half. The final pass was a forward pass and thus it did not result in a try. The first defending clip (defending 1) involved the team defending close to their line where the opposition scored a try. And finally, the second defending clip (defending 2) involved the team defending mid-field where the opposition scored a try on the wing.

Between 1 and 7 days after the game, we interviewed the participants individually, asking them to view each clip and asked a series of set questions developed devised by us on the basis of existing conceptual research in organisational psychology (Eccles & Johnson 2009). We informed participants that they could watch each clip as often as required, rewinding and pausing at any time. After watching each clip, we asked participants the same set of questions and audio recorded their responses. For both attacking clips, we asked the following questions:

1) Can you talk me through what is happening here?
2) Was this move planned in advance?
3) How do you think this play was initiated?
4) What influenced what you did in this play?
5) What were you thinking as this play unfolded?
6) What were you feeling while it was happening?
7) What do you think each of your teammates involved in the play were thinking? Starting with (player name).

For both defending clips, we asked the following questions:

1) Can you talk me through what's happening here?
2) What went wrong?
3) Why do you think it went wrong?
4) How did your team fix the problem to prevent it happening again?
5) What were you thinking during this play?
6) What were you feeling during this play?
7) What do you think each of your teammates involved in the play were thinking? Starting with (player name).

We also prompted participants to provide more information and elaborate on particular comments.

Next, we asked participants to rate on 7-point likert scales, how much they remembered about the match, how well the team played, and how well they played as an individual. Specifically, we asked:

1) On a scale of 1 to 7 how much do you remember about the match as a whole, with 1 meaning you don't remember much at all and 7 being you remember a lot?
2) How well did the team play? 1 meaning the team played very poorly and 7 meaning the team played very well.
3) How well did you play? 1 meaning you played very poorly and 7 meaning you played very well.

To conclude, we asked participants how long they had been playing touch football and how long they had played with this particular team. Full results are to be published elsewhere.

Study Two: Coordination in an OzTag Team: Coach Interview

We conducted an in-depth, semi-structured interview with the coach of an adult women's representative Oz-Tag team. The team had previously played at a State level tournament, and

the interview questions focused on the team's performance at that tournament, in particular, the semi-final match where the team was defeated and consequently knocked-out of the competition. The interview is part of larger qualitative study involving players from the team, results of which are to be published elsewhere. Dr Cox was again an informed participant in this study, as she is a regular member of the team and a co-investigator.

The coach had been the coaching this particular team for one year at the time of this interview, and the team itself had played together for approximately 4 years. The coach provided us with written informed consent, and he was compensated financially for his time.

The results of this interview and the wider study are to be published elsewhere. The quote featured in Section Four of Chapter Four was taken from the coach's answer to Interview Question 3.

The following questions were comprised the interview and prompts were when necessary:

1. To begin, we'd like you to tell us what you think each player's role in the team should be?
    a. Justine Bonner
    b. Rebecca Briancourt
    c. Kate Cass
    d. Deborah Cooper
    e. Rochelle Cox
    f. Denise Evans
    g. Felicity Goodwin
    h. Amy Helm
    i. Georgie Israel
    j. Belinda Lester
    k. Leanne Marsh
    l. Jane Stanley
    m. Lisa Taliana
    n. Francine Walker
    o. Angela Wilcox

p. Melanie Locke
q. Corrine

2. Can you just briefly talk us through the state cup semi-final against ACT and describe how the game unfolded.
3. What do you think went wrong for the Souths team during the state cup semi-final against ACT?

   Prompts:

   a. Can you tell us about the Souths team's communication during this game?
   b. Can you think of a good and a bad key moment during the game and describe what happened?
   c. What factors contributed to this moment being good/bad?

4. Who usually plays well together?
   a. Why do you think they play well together?
   b. How did they play in the game against ACT?
   c. Why did they play well/poorly together in the game against ACT?

Now we'd like to ask you about groups of players in the team.

5. Are there any players who socialise together (that you are aware of)?
6. A lot of players in the team mentioned that they play well with Georgie. What are your thoughts about this?
   a. Why do you think they said this?
7. A lot of players in the team mentioned that they play well with Rochelle. What are your thoughts about this?
   a. Why do you think they said this?
8. Are there any players in the team who do not play well together? If so, who are these players?
   a. Why do you think they don't play well together?
   b. Did they play well together in the game against ACT?
9. And finally, what do you think the team needs to work on most to avoid the kind of problems that arose during the game against ACT?
   a. How might this be achieved?
10. Is there anything else you would like to add?

Ethics Approval

Gmail by Google

Kellie Williamson <kellie.m.williamson@gmail.com>

Fwd: Amendment Request for Project REF HE23FEB2007-R05039

John Sutton <john.sutton@mq.edu.au> Wed, Apr 13, 2011 at 4:17 PM
To: Kellie Williamson <kellie.m.williamson@gmail.com>, Rochelle Cox <rochelle.cox@mq.edu.au>, Amanda Barnier <amanda.barnier@mq.edu.au>

---------- Forwarded message ----------
From: Ethics Secretariat <ethics.secretariat@mq.edu.au>
Date: Wed, Apr 13, 2011 at 3:19 PM
Subject: Re: Amendment Request for Project REF HE23FEB2007-R05039
To: John Sutton <john.sutton@mq.edu.au>

Dear John

RE: "From autobiographical memory to collective memory: An interdisciplinary study of individual and group cognition" (REF: HE23FEB2007-R05039)

Thank you for your correspondence. The following amendment to the above study has been approved:

1. To amend your procedure by video-recording a game played by the South East Oztag Team and the Northern Suburbs Touch Team. For both teams you will code the footage and then conduct individual interviews with the players while they watch selected parts of the footage. The footage will be stored for use and used in research presentations.

Please accept this email as formal correspondence approving your amendment request.

Regards

Nicola Myton
[Quoted text hidden]
--
Office of the Deputy Vice Chancellor (Research)

Ethics Secretariat

Research Office
Level 3, Research HUB, Building C5C
Macquarie University
NSW 2109

Ph: +61 2 9850 6848
Fax: +61 2 9850 4465

Email: ethics.secretariat@mq.edu.au

[Quoted text hidden]

References

Abernethy, B. (1993) Searching for the Minimal Essential Information for Skilled Perception and Action, *Psychological Research*, 55: 131-138.

Abernethy, B., Zawi, K. & Jackson, R. C. (2008) Expertise and Attunement to Kinematic Constraints, *Perception* 37: 931-948.

Adams, F. & Aizawa, K. (2001) The Bounds of Cognition, *Philosophical Psychology*, 14(1): 43-64.

Adams, F. & Aizawa, K. (2005) Defending Non-Derived Content, *Philosophical Psychology*, 18(6): 661-669.

Adams, F. & Aizawa, K. (2008) *The Bounds of Cognition*. Malden: Blackwell Publishing.

Araujo, D., Davids, K. & Hristovski, R. (2006) The Ecological Dynamics of Decision Making in Sport, *Psychology of Sport and Exercise*, 7: 653-676.

Bechtel, W. (2008) *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, New York, Routledge.

Bechtel, W. (2009a) Explanation: Mechanism, Modularity and Situated Cognition. In P. Robbins & M Aydede (eds), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press, 155-170.

Bechtel, W. (2009b) Looking Down, Around, and Up: Mechanistic Explanation in Psychology, *Philosophical Psychology*, 22(5): 543-564.

Beer, R. (2000) Dynamical Approaches to Cognitive Science, *Trends in Cognitive Science* **,** 4: 91-99.

Bermudez, J. L. (2010) *Cognitive Science: An Introduction to the Science of the Mind*. New York: Cambridge University Press.

Blickensderfer, E. L., Reynolds, R., Salas, E. & Cannon-Bowers, J. A. (2010) Shared

Expectations and Implicit Coordination in Tennis Doubles Teams, *Journal of Applied Sport Psychology*, 22(4): 486-499.

Boden, M. A. (2006) *Mind as Machine: A History of Cognitive Science*, Oxford: Oxford University Press.

Bohanek, J. G., Marin, K. A., Fivush, R., & Duke, M. P. (2006) Family Narrative Interaction and Children's Sense of Self, *Family Process*, 45(1), 39–54.

Bortolotti, L. & Cox, R. E. (2009) 'Faultless' Ignorance: Strenths and Limitations of Epistemic Definitions of Confabulation, *Consciousness and Cognition*, 18(4): 952-965.

Bratman, M. E. (1992) Shared Cooperative Activity, *The Philosophical Review*, 101(4): 327-341.

Bratman, M. E. (1993) Shared Intention, *Ethics*, 104: 97-113.

Bratman, M. E. (2009) Modest Sociality and the Distinctiveness of Intention, *Philosophical Studies*, 144(1): 149-165.

Bunge, M. (1977) Levels and Reduction, *American Journal of Physiology - Regulatory, Integrative and Comparative Physiology*, 233: 75-82.

Cannon-Bowers, J. A., Salas, E. & Converse, S. (1993) Shared Mental Models in Expert Team Decision-Making. In N. J. Castellan, Jr. (Ed.), *Individual and group decision making*, Hillsdale, NJ: Erlbaum, 221-246.

Cannon-Bowers J.A., Tannenbaum S. I., Salas E. & Volpe C. E. (1995) Defining competencies and establishing team training requirements. In R. A. Guzzo, & E. Salas (eds.), *Team Effectiveness and Decision Making in Organizations*. San Francisco: Jossey Basspp, 333–380.

Chemero, A. (2009) *Radical Embodied Cognitive Science*. Cambridge, MA: The MIT Press.

Clark, A. (1989) *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed*

Processing. Cambridge: The MIT Press.

Clark, A. (2001) *Mindware: an Introduction to the Philosophy of Cognitive Science.* New York: Oxford University Press.

Clark, A. (2005). "Intrinsic Content, Active Memory and the Extended Mind." Analysis **65**(1): 1-11.

Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension.* New York: Oxford University Press.

Clark, A. & Chalmers, D. (1998) The Extended Mind, *Analysis*, 58(1): 7-19.

Cordovil, R., Araújo, D., Davids, K., Gouveia, L., Barreiros, J., Fernandes, O., Sidónio, S. (2009) The Influence of Instructions and Body-Scaling as Constraints on Decision-Making Processes in Team Sports, *European Journal of Sport Science*, 9(3): 169-179.

Craver, C. & Bechtel, W. (2007) Top-Down Causation Without Top-Down Causes, *Biology and Philosophy,* 22: 547-563.

Davids, K. W., Button, C., & Bennett, S. J. (2008) *Dynamics of Skill Acquisition*. Champaign, Il: Human Kinetics.

Dennett, D. C. (1996) *Kinds of Minds: Toward and Understanding of Consciousness*. New York: Basic Books.

Dreyfus, H. L. (2002) Intelligence without Representation: the Relevance of Phenomenology to Scientific Explanation, *Phenomenology and the Cognitive Sciences*, 1(4): 367-383.

Dreyfus, H. L. (2002) Refocusing the Question: Can There Be Skillfull Coping without Propositional Representations or Brain Representations? *Phenomenology and the Cognitive Sciences* 1(4): 413-425.

Eccles, D.W. & Tenenbaum, G. (2004) Why an Expert Team is More Than a Team of Experts: A Socio-Cognitive Conceptualization of Team Coordination and Communication in

Sport, *Journal of Sport and Exercise Psychology*, 26: 542-560.

Eccles, D. W. & Johnson, M. B. (2009) Letting the Social and Cognitive Merge. In S. D. Mellalieu & S. Hanton (eds.), *Advances in Applied Sports Psychology*, London: Routledge, 281-316.

Ericsson, K. A. (2006) The influence of experience and deliberate practice on the development of superior expert performance. In K. A. Ericsson, N. Charness, P. Feltovich and R.R. Hoffman (eds), *Cambridge Handbook of Expertise and Expert Performance*, Cambridge, UK: Cambridge University Press, 685-706.

Ericsson, K. A. & Lehmann A. C. (1996) Expert and exceptional performance: evidence on maximal adaptations on task constraints, *Annual Review of Psychology*, 47:273–305.

Ericsson, K. A. & Simon, H. A. (1980) Verbal Reports as Data, *Psychological Review*, 87: 215–259

Fajen, B. R., Riley, M. A., & Turvey, M. T. (2008) Information, Affordances, and the Control of Action in Sport, *International Journal of Sport Psychology*, 40: 79-107.

Farrow, D., & Abernethy, B. (2003) Do Expertise and the Degree of Perception-Action Coupling Affect Natural Anticipatory Performance?, *Perception*, 32: 1127-1139.

Fodor, J. (1968). *Psychological Explanation*. New York: Random House.

Fodor, J. (1980) Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology, *Behavioral and Brain Sciences*, 3(01): 63-73.

Gibson, J. J. (1986) *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Erlbaum.

Giere, R. N. (2002) Scientific Cognition as Distributed Cognition. In: P. Carruthers, S. Stich & M. Siegal (eds), *Cognitive Bases of Science*, Cambridge: Cambridge University Press, 285-299.

Giere, R. N. & Moffatt, B. (2003) Distributed Cognition: Where the Cognitive and the Social Merge, *Social Studies of Science* 33(2): 1-10.

Gilbert, M. (1987) Modelling Collective Belief, *Synthese,* 73(1): 185-204.

Gilbert, M. (1992) *On Social Facts*. New Jersey: Princeton University Press.

Halbwachs M (1992) [1925] *On Collective Memory*. A.C. Lewis (ed. & trans.) Chicago, IL: The University of Chicago Press.

Hinsz, V. B., Vollrath, D. A. & Tindale, R. S. (1997) The Emerging Conceptualisation of Groups as Information Processors, *Psychological Bulletin* 121(1): 43-64.

Huebner, B. (2008) Do You See What We See? An Investigation of an Argument Against Collective Representation, *Philosophical Psychology* 21(1): 91-112.

Huebner, B. (2011) Genuinely Collective Emotions, *European Journal for the Philosophy of Science*, 1(1): 89-118.

Hurley, S. L. (1998) *Consciousness in Action*. Cambridge, MA: Harvard University Press.

Hutchins, E. (1995a) *Cognition in the Wild.* Cambridge: The MIT Press.

Hutchins, E. (1995b) How A Cockpit Remembers Its Speeds, *Cognitive Science*, 19: 265-288.

Kirsh, D. (2010) Thinking with the Body. In S. Ohlsson & R. Catrambone (eds), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, Austin, TX: Cognitive Science Society, 2864-2869.

Klimoski, R. & Mohammed, S. (1994) Team mental model: Construct or metaphor? *Journal of Management*, 20:403-437.

List, C. & Pettit, P. (2011) *Group Agency: The Possibility, Design and Status of Group Agents.* New York, Oxford University Press.

Loftus, E. F. (1979). Reactions to Blatantly Contradictory Information, *Memory & Cognition*, 7:368-374.

Machamer, P., Darden, L. & Craver, C. (2000) Thinking about Mechanisms, *Philosophy of Science* 67(1): 1-25.

Mann, D. L., Abernethy, B. & Farrow, D. (2010) Action Specificity Increases Anticipatory Performance and the Expert Adavantage in Natural Interceptive Tasks, *Acta Psychologica* 135(1): 17-23.

McPherson, S. L. & MacMahon, C. (2008) How Baseball Players Prepare to Bat: Tactical Knowledge as a Mediator of Expert Performance in Baseball, *Journal of Sport and Exercise Psychology*, 30: 755-778.

McPherson, S. L. & Thomas, J. R. (1989) Relation of Knowledge and Performance in Boys' Tennis: Age and Expertise, *Journal of Experimental Child Psychology*, 48(2): 190-211.

McRobert, A. P., Ward, P., Eccles, D. W. & Willams, A. M. (2011) The Effect of Manipulating Context-Specific Information on Perceptual-Cognitive Processes During a Simulated Anticipation Task, *British Journal of Psychology*, 102(3): 519-534.

Montero, B (2010) Does Bodily Awareness Interfere with Highly Skilled Movement? *Inquiry*, 53(2): 105-122.

Montero, B. & Evans, (2011) Intuitions without Concepts Loses the Game: Mindedness in the Art of Chess, *Phenomenology and the Cognitive Sciences*, 10(2): 175-194.

Nisbett, R. E. & Wilson, T. D. (1977) Telling more than we can know: Verbal reports on mental processes, *Psychological Review*, 84 (3): 231–259

O'Brien, G. & Opie, J. (2006). How Do Connectionist Newtworks Compute? *Cognitive Processing* **7**: 30-41.

Pacherie, E. & Haggard, P. (2010) What are Intentions? In L.Nadel & W. Sinnott-Armstrong (eds), *Conscious Will and Responsibility: a Tribute to Benjamin Libet.* New York, Oxford University Press**:** 70-84.

Pettit, P. (2003) Groups with Minds of Their Own. In F. F. Schmitt (ed), *Socializing Metaphysics*. Lanham, MD: Rowan & Littlefield Publishers Inc.**:** 167-193.

Pettit, P. & Schweikard, D. (2006) Joint Actions and Group Agents, *Philosophy of Social Sciences*, 36: 18-39.

Preston, B. (2012) *A Philosophy of Material Culture: Action, Function and Mind*. Routledge.

Pylyshyn, Z. W. (1986) *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge, MA: The MIT Press.

Reimer, T. R., Park, E. S. & Hinsz, V. B. (2006) Shared and Coordinated Cognition in Competitive and Dynamic Task Environments: An Information-Processing Perspective for Team Sports, *International Journal of Sport and Exercise Psychology*, 4(4): 376-400.

Richardson, D. C., & Dale, R. (2005) Looking to Understand: The Coupling Between Speakers' and Listeners' Eye Movements and Its Relationship to Discourse Comprehension, *Cognitive Science*, 29(6): 1045-1060.

Richardson, M. J., Marsh, K. L., Isenhower, R. W., Goodman, J. R. L, & Schmidt, R. C. (2007) Rocking Together: Dynamics of Intentional and Unintentional Interpersonal Coordination, *Human Movement Science*, 26: 867-891.

Roediger III, H. L., Meade, M. L. & Bergman, E. T. (2001) Social Contagion of Memory, *Psychonomic Bulletin & Review*, 8(2): 365-371.

Rupert, R. (2005) Minding One's Cognitive Systems: When Does a Group of Minds Constitute a Single Cognitive Unit? *Episteme: A Journal of Social Epistemology*, 1: 177–88.

Rupert, R. (2011) Empirical Arguments for Group Minds: A Critical Appraisal, *Philosophy Compass*, 6(9): 630-639.

Searle, J. R. (1990) Collective Intentions and Actions, In P. Cohen, J. Morgan & M. E. Pollack (eds), *Intentions in Communication*. Cambridge: The MIT Press.

Shore, B. (2009) Making Time for Family: Schemas for Long-Term Family Memory, *Social Indicators Research*, 93: 95–103.

Smeeton, N. J., Williams, A. M., Hodges, N. J., and Ward, P. (2005) The relative effectiveness of explicit instruction, guided-discovery and discovery learning techniques in enhancing perceptual skill in sport, *Journal of Experimental Psychology: Applied*, 11:98-110.

Steel, D. (2004) Can a Reductionist be a Pluralist? *Biology and Philosophy*, 19: 55-73.

Steel, K. A., Adams, R. D. & Canning, C. G. (2006) Identifying Runners as Football Teammates From 400msec Video Clips, *Perceptual and Motor Skills*, 103: 901-911.

Steel, K. A., Adams, R. D., Canning C. G. (2007) Identifying Swimmers as Water-Polo or Swim Team-mates from Visuak Displays of Less than One Second, *Journal of Sports Sciences*, 25(11): 1251-1258.

Stevens, C., Malloch, S., McKechnie, S., & Steven, N. (2003) Choreographic cognition: The time-course and phenomenology of creating a dance, *Pragmatics & Cognition*, 11(2) 297-326.

Sutton, J. (2007) Batting, Habit and Memory, *Sport in Society*, 10(5): 763-786.

Sutton, J., Harris, C. B., Keil, P. G. and Barnier, A. J. (2010) The Psychology of Memory, Extended Cognition, and Socially Distributed Remembering, *Phenomenology and the Cognitive Sciences*, 9: 521-560.

Theiner, G. (2010a) Making Sense of Group Cognition: The Curious Case of Transactive Memory Systems, *ASCS09: Proceedings of the 9th Conference of the Australasian Society for Cognitive Science*, Sydney, Australia: Macquarie Centre for Cognitive Science.

Theiner, G., Allen, C. & Goldstone, R. L. (2010b). Recognizing Group Cognition, *Cognitive Systems Research*, 11(4): 378-395.

Theiner, G. & O'Connor, T. (2010c) The emergence of group cognition. In A. Corradini & T. O'Connor (ed.), *Emergence in science and philosophy*, New York: Routledge, 78-117.

Theureau, J. (2003) Course-of-Action Analysis and Course-of-Action Centered Design. In E. Hollnagel (ed) *Handbook of Cognitive Task Design*. Mahwah, NJ: Lawrence Erlbaum Inc.**:** 55-81.

Tollefsen, D., & Dale, R. (2011) Naturalizing Joint Action: A Process Based Approach. *Philosophical Psychology*, iFirst: 1-23.

Totterdell, P. (2000) Catching Moods and Hitting Runs: Mood Linkage and Subjective Performance in Professional Sport Teams, *Journal of Applied Psychology*, 85(6):848-859.

van Gelder, T. (1998) The Dynamical Hypothesis in Cognitive Science, *Behavioral and Brain Sciences*, 21: 1-14.

Wegner, D. M. (1986) Transactive Memory: A Contemporary Analysis of the Group Mind. In B. Mullen, & G. R. Goethals, G.R., *Theories of Group Behaviour*. New York: Springer-Verlag 185-208.

Williams, M. A. & Ericsson, K. A. (2005) Perceptual-cognitive Expertise in Sport: Some Considerations when Applying the Expert Performance Approach, *Human Movement Science*, 24: 283-307.

Wilson, R. A. (2004) *Boundaries of the Mind: The Individual in the Fragile Sciences*. Cambridge, Cambridge University Press.

Wilson, R. A. (2005) Collective Memory, Group Minds, and the Extended Mind Thesis, *Cognitive Processing*, 6: 227-236.

Wilson, R. A., & Clark, A. (2008) How to Situate Cognition: Letting Nature Take its Course. In P. Robbins & M Aydede, *The Cambridge Handbook of Situated Cognition*. New York, Cambridge University Press**:** 55-77.

Wimsatt, W.C. (2007) *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge: Harvard University Press.

Von Eckardt, B. (1993) *What is Cognitive Science?* Cambridge, MA: The MIT Press.