

In support of action

Catherine Rebecca Dale

UCL

A thesis submitted for the degree of Ph.D. in Philosophy

I, Catherine Rebecca Dale, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

23rd July 2023

For my father, Paul Dale

The mutual dependence of men is so great in all societies that scarce any human action is entirely complete in itself, or is performed without some reference to the actions of others, which are requisite to make it answer fully to the intentions of the agent.

(Hume [1748] 2007, Section VIII, Part 1, §17, p89)

Abstract

Needing another agent to support one's action does not undermine our agency, it often facilitates it. We are not isolated entities, but agents living with and amongst other agents. This thesis spells out how it is that our reliance on others can enable agency and how supporting action can enhance rather than detract from another's ability to act on their intentions. At its core is a characterization of a familiar and commonplace kind of action, but one that has not been the focus of sustained philosophical investigation. This is what I call dependent intentional action, action in which an agent's ability to realize their ends relies on another agent directing their agency at those ends, in order to help them. The thesis begins with an exploration of the nature of individual and joint intentional action, offering criteria for them that are compatible with the dominant approaches in the existing literature. I then argue that our moral responsibility for actions relies on being the agent of the action, further showing the importance of bearing this particular relation to an action. Having set out this background, I identify dependent intentional action, characterizing it and arguing for the claim that this is the action of the supported agent, the one that is being helped. In the final chapters of the thesis, I consider the agency of young children. First, I argue that there are no insurmountable barriers to young children engaging in intentional action, either individually or with others. Then I argue that dependent intentional action extends the abilities of young children, allowing them to engage in intentional action even if they cannot do this alone. Taken as a whole, this thesis shows that being helped can enable human agents to do far more than they can alone.

Impact Statement

As of April 2018, UCL requires an impact statement to be included in all PhD theses, which should describe how the expertise, knowledge, analysis, discovery or insight presented in the thesis could be put to beneficial use.

The philosophy of action tends to focus on either individual intentional action – the things that agents do alone – or joint intentional action – the things that agents do when acting with another to realize a shared aim. This thesis identifies a kind of action that has not been addressed by this literature, intentional action in which one agent is supported by the help of another to realize their individual aim – dependent intentional action. Identifying the existence of this kind of action carves out new space in this field, resisting this standard picture. This has the potential to prompt future work on dependent intentional action from philosophers interested in exploring the nature of this kind of action, as well as those interested in action more generally. Chapters 6 and 7 engage with work from developmental psychology, and the ideas therein could also prompt further work in both philosophy and psychology. I hope to publish versions of several of the chapters of this thesis in academic journals, which would enable them to have this influence.

The idea of dependent intentional action also has a host of interdisciplinary implications. It corroborates and extends work in disability studies and within the disability movement that argues for the importance of access to particular kinds of support, and to ensuring society is organized so as to provide them. It has consequences for discussions around accomplice liability in jurisprudence, which themselves propose potential changes to criminal law.

Acknowledgements

My PhD journey has been longer and more complicated than I could have expected at the outset. The idea that I would reach submission has felt implausible at many points, and even now it seems a little unlikely that I am really here. However, I am, and I owe thanks to many people for getting me here, and to the LAHP, RIP, and FAGI at Universität Leipzig for financial support that allowed me to complete this research.

Rory Madden saw me through this thesis, reminding me that I could not make the argument all at once, and helping me to make far it more deftly than I could have otherwise. He is a thoughtful and insightful philosopher for whose influence and input I am endlessly grateful. The care he takes over his own ideas and the ideas of others has helped me to slow down in ways I needed, and his supervision has drawn from me better and clearer work than I could have dreamt of producing without his help. Often, I have found myself returning to work and realizing with reflection that the problems he was nudging me to see get to the heart of an issue that I was glossing over. His support and guidance were crucial, and I hope I can embody a modicum of the philosophical care and precision that he has modelled to me. I am still learning the lessons he has taught me, but I am so thankful for them, and for the kindness he has shown me and the support he has given me.

Lucy O'Brien supervised me for the majority of my MPhilStud and was my secondary supervisor for this thesis, and her influence can be felt throughout. She guided me on the journey from student to researcher, encouraging me to read widely and to explore and critically engage with ideas from a variety of places. Much of how I approach what it is to do philosophy is due to her; her openness to work across schools and subfields, and to the presence of philosophical insight in literature are some of things I hope to take with me. Mark Kalderon supervised me for a term early in the project, instilling in me the importance of not only having a view that one wants to state, but a question that one wants to answer, a problem one wants to solve. Seminars run by John Vorhaus and Sarah Richmond influenced this project, and I am thankful for the opportunities they granted me to present my work. Rob Simpson and Fiona Leigh both provided support for which I am extremely grateful.

The length of time I spent at UCL means that I got to know many generations of graduate students. Discussions in WIP, over coffees and lunches, during breaks in the graduate study space and

bootcamp sessions (both in person and online), have helped me shape these ideas and have been crucial to how enjoyable my time at UCL was. I want to thank Naomi Alderson, Victor Braga Weber, Julian Bacharach, Vanessa Carr, Nicholas Currie, Pete Faulconbridge, Lisa Grant, Alice Harberd, Alec Hinshelwood, Charles Jansen, Kirstine La Cour, Rowan Mellor, Alex Murphy, Polly Mitchell, Edgar Phillips, Ashley Shaw and Nikhil Venkatesh. I owe especial thanks to Lizzy Holt, Michael Markunas and Rebecca Rowson for their friendship and support.

The role that my time in Leipzig played in producing this project outstrips the briefness of my stays. It was a place that led to closer connections with people I knew from other times and places, but also for the building of new relationships. Thanks are owed to Kristina Musholt for allowing me to attend her seminars, without which this project would have not been what it is, and to the many other staff, students and visiting fellows who made it such a fertile philosophical environment. I am especially grateful to have met Jana Thesing there, and to have built a friendship with her that now spans an ocean.

Knowing where to place Zara Bain in these acknowledgements has been difficult, as she is a dear friend who happens to be an excellent philosopher, but I have settled for here. Thanks are owed to her for reminding me both that writing a thesis is hard because it's trying to put the octopus in the box, and that there is so much more that matters.

Claire, Ellie, Ellen and Esther – thank you for your friendship and love in the face of my tendency to retreat into a thesis-hole, I am lucky to know you.

The biggest thanks are owed to my family. The certainty of my father, Paul, that this was a thing that I not only could do, but would do, kept me going with the thesis at times when nothing else could. My mother, Cheryl, taught me to think philosophically long before either of us realized that was what she was doing. The support they have both shown me is immeasurable, I am so fortunate to have parents that encourage my curiosity and who so firmly believe in the value of learning and of inquiry. My brother, Adam, has been a huge support throughout this project and I am so thankful to call him a friend. Despite philosophy being firmly outside his comfort zone, he and my mother have between them read the entirety of this thesis. In the final months and weeks of this project he has listened as I have worked through difficulties, and has helped me to see the woods for the trees time and time again.

A note on pronouns

Throughout this thesis, I have used singular ‘they’ to refer to a generic agent. Use of ‘they/them/their/themselves’ to refer to individuals of unknown gender (‘whoever gets to the event first will be all by themselves, I hope they have their phone with them’) is not unusual. This choice avoids unnecessarily gendering elements of the discussion that apply to all agents.

In examples, I have used ‘they/them/their/themself’ alongside ‘she/her/hers/herself’ and ‘he/him/his/himself’. ‘Themself’ is the generally accepted reflexive form when an individual uses they/them pronouns. I use a variety of pronouns to as to allow easier differentiation between agents in examples involving multiple agents, and include they/them to acknowledge a variety of gender identities.

Contents

Abstract.....	5
Impact Statement.....	6
Acknowledgements.....	7
A note on pronouns.....	9
Contents.....	10
Introduction.....	12
The structure.....	18
Chapter 1 – Individual intentional action.....	21
1.1 Intentional agency.....	22
1.2 Intentional action.....	24
1.2.1 Bodily movements.....	28
1.2.2 Failures to act and mere attempts.....	31
1.3 Intention.....	35
1.3.1 Conditions on intention.....	42
1.4 Practical reason.....	45
1.4.1 Transformative accounts.....	48
1.4.2 Additive accounts.....	49
Chapter 2 – Joint intentional action.....	53
2.1 Joint intentional action.....	53
2.2 Intentions and joint action.....	56
2.2.1 Plural intention.....	57
2.2.2 Participatory intention.....	58
2.2.3 Conditional intention.....	59
2.2.4 Anscombean approaches to joint action.....	61
2.3 A general model of joint intentional action.....	65
Chapter 3 – Moral responsibility and personal identity.....	70
3.1 The ordinary view.....	71
3.2 Owning an action.....	74
3.3 Köhler’s defence of Slogan.....	79
3.4 Defending Slogan.....	85
3.5 Ownership and identity.....	88
Chapter 4 – Dependent intentional action.....	91
4.1 The nature of dependent intentional action.....	92
4.1.1 Examples of dependent intentional action.....	94
4.1.2 Characterizing dependent intentional actions.....	97
4.2 The role of guidance and responsiveness.....	100
4.3 The necessity of intentional helping.....	105
4.3.1 Non-agential support.....	107

4.3.2 Accidental help.....	109
4.3.3 Intentionally helping and intending to help.....	111
4.4 Dependency of actions and of act-types.....	117
Chapter 5 – Supported agents as agents.....	121
5.1 Dependent intentional action is action.....	122
5.2 Dependent intentional action is not joint action.....	126
5.2.1 The supported agent.....	127
5.2.2 The supporting agent.....	129
5.3 Dependent intentional action is not action of the supporting agent.....	131
5.4 Dependent intentional action and ways of not ϕ -ing.....	133
5.4.1 Acting on someone’s behalf.....	134
5.4.2 The supporting agent as not ϕ -ing.....	136
5.4.3 Intending to ϕ with help.....	138
5.5 Deviant causal chains.....	139
5.5.1 The relevance of deviant causal chains.....	140
5.5.2 Basic actions.....	143
5.5.3 Deviant causal chains and guidance.....	144
5.5.4 Kinds of heteromesy.....	147
Chapter 6 – Barriers to intentional agency in early childhood.....	150
6.1 Early childhood.....	151
6.2 Know-how.....	155
6.2.1 Early childhood and know-how.....	156
6.3 Intending to ϕ and knowing how to ϕ	159
6.3.1 Identifying intentional action.....	160
6.3.2 Semantic externalism.....	163
6.3.3 Learning to ϕ	165
6.3.4 Knowing what it is to ϕ	168
6.4 Joint intentional action.....	169
6.4.1 Understanding others as agents.....	170
6.4.2 Young children’s understanding of other agents.....	172
6.4.3 Intention, desire and belief.....	175
6.5 False-belief tests and theory of mind.....	176
6.5.1 Teleological and hybrid accounts.....	178
6.5.2 Minimal theory of mind.....	180
6.6 Joint action in early childhood.....	183
Chapter 7 – Dependent intentional action in early childhood.....	188
7.1 Supported intention.....	188
7.2 Scaffolding.....	191
7.3 Executive function.....	196
7.3.1 The development of executive function and its role in intentional action.....	199
7.3.2 Teaching executive function.....	201
7.3.3 Exogenous executive function.....	204
7.5 Scaffolded action as dependent action.....	207
Conclusion.....	213
Bibliography.....	215

Introduction

What is important is that disabled people have the right to choose. Independent living is about choice and control, it is not about doing everything by yourself. Nobody - whether they have an impairment or not - can do everything themselves.

When disabled people use P.A's it does not mean that they are dependent on others. If a P.A has to push a wheelchair, help a disabled person dress or reach for a book, it should be seen as enhancing the disabled persons ability to live independently.

(Barnes and British Council of Organizations of Disabled People 1993)¹

This thesis is in large part an attempt to resolve the apparent tension in this description. Needing another agent to support one's action does not undermine our agency, it often facilitates it. We are not isolated entities, but agents living with and amongst other agents. This thesis spells out how it is that our reliance on others can enable agency and how supporting action can enhance rather than detract from another's ability to act on their intentions. At its core is a characterization of a familiar and commonplace kind of action, but one that has not been the focus of sustained philosophical investigation. This is dependent intentional action, action in which an agent's ability to realize their ends relies on another agent directing their agency at those ends, in order to help them. In writing this thesis, my interest is in correctly describing the world — it is primarily a project in the philosophy of action from the perspective of metaphysics and the philosophy of mind — but this notion is one that is worth taking seriously if we are interested in doing right by other people. If there are ways of acting that allow us to facilitate others' agency, this is something with moral and political consequences. I am interested in showing that such a way of acting exists, and exploring the ways that it relates to other kinds of intentional action.

The idea that we act with and amongst others is not novel, as evidenced by the quotation from Hume with which this document opens. In this thesis, I am drawing out and characterizing a kind of action that has not been paid close attention in the existing literature. Before I can do this, though, I want to

¹ The notion of independent living, which is described in this quote, is staple of disability activism. It is enshrined in the United Nations Convention on the Rights of Persons with Disabilities (UNCRPD) (UN General Assembly 2007). It is explicitly stated in Article 19, 'Living independently and being included in the community', but it also runs through the convention

set out the ways in which philosophers have acknowledged the role that others can play in our actions, and the roles that we can play in the action of others. Joint action is a familiar consideration in the philosophy of action, and I will discuss this literature in Chapter 2. Unlike joint action, however, in which agents act together, the cases I am exploring are situations in which an agent contributes to another's action. There are several strands of work that acknowledge that there are kinds of involvement in intentional action that are not being an agent of an individual or joint intentional action. Primarily these are ways of supporting and enabling another's action. This is the kind of involvement that I will discuss in characterizing dependent intentional action, and at this point I want to identify strands of work that act as precursors to the ideas I will develop later in the thesis.

The first role that we play in others' action is as elements of a social context that can both limit and facilitate the authentic desires and preferences of an agent. Work in relational autonomy² argues that we should not understand autonomy as something that depends only on the individual agent, but something that arises out of social conditions:

The term "relational autonomy" does not refer to a single unified conception of autonomy but is rather an umbrella term, designating a range of related perspectives. These perspectives are premised on a shared conviction, the conviction that persons are socially embedded and that agents' identities are formed within the context of social relationships and shaped by a complex of intersecting social determinants, such as race, class, gender, and ethnicity. Thus the focus of relational approaches is to analyze the implications of the intersubjective and social dimensions of selfhood and identity for conceptions of individual autonomy and moral and political agency.

(Mackenzie and Stoljar 2000a, p4)

Relational autonomy theorists are interested in characterizing the way that oppressive social conditions impair and enable autonomy. Although the focus on autonomy situates the impact of other agents upstream of action, in the formation of the preferences and interests, these preferences and interests impact the choices an agent makes in acting. Which things they do, which things they aim at, how they behave on a day-to-day basis, are all consequences of their desires and preferences (amongst other things), and relational autonomy recognizes that other agents play a role in their formation.

² (Mackenzie and Stoljar 2000b) gives a good overview of this field.

The notion of the extended mind begins with idea that elements of our environment, such as a notebook, could form part of the material realization of our cognition, with Clark and Chalmers (1998) imagining an agent whose reliance on a notebook constitutes an extension of their mind. Later in this paper, they consider the possibility that extended cognition could extend to include another agent:

Could my mental states be partly constituted by the states of other thinkers? We see no reason why not, in principle. In an unusually interdependent couple, it is entirely possible that one partner's beliefs will play the same sort of role for the other as the notebook plays

(Clark and Chalmers 1998, p17)

As with relational autonomy, what this offers us is a picture of others as providing the conditions that facilitate agency. The other agent features as a tool, or as a part of the agent they help, rather than as themselves being an agent of the action. This kind of reliance on another expands what an agent is able to do, in the same way that a notebook allows an agent to access and act on information that they have forgotten. This is not an agent directing themselves at supporting a particular action of another, but being used by another as part of what realizes their cognition.

There is also literature that understands agents as supporting the actions of others in more direct and active ways, with agents acting to support others rather than forming part of the material reality that enables their agency. Baier (1997) argues that we are mistaken in conceiving of individual action as a precursor to joint action, and attends particularly to action in childhood. She claims:

We make unnecessary philosophical problems for ourselves if in our philosophy we forget that individual action was something we all had to learn, and that we learned it as a departure from common action.

(1997, p29)

In this essay, Baier centres the way that acting with others, and learning from others, is key to becoming able to act alone. This is an idea I will explore in the later sections of the thesis, looking in chapters 6 and 7 at what young children are able to do both individually and jointly, and the ways in which support from other agents can expand these abilities. In drawing attention to what she calls 'the commons of the mind', Baier reorients us to consider acting with others (although not necessarily

jointly) as central to our agency. I will follow her lead in looking at the ways that other agents are involved in our action.

There are two strands of literature that have developed in more detail the idea that an agent can be involved in the action of another. The first of these is work in jurisprudence, looking at the notion of accomplice liability.³ Crimes are committed by *principals*, those who actually do the thing in question. Accomplices to crimes are those that aid, abet, counsel or procure the principal of the crime. In UK and US law, accomplices of crimes, if charged, are charged with the same crime as the principal, but as a *secondary party*. This understands all those that are liable for a crime as doing the same thing, even if both the scale and nature of their contribution differs dramatically:

Recent case law from different jurisdictions has confirmed that D is guilty of theft if she holds P's baby while he steals cash from a register; of manufacturing bootleg liquor if she brings P a midday meal so that his work may continue uninterrupted; of supplying class A drugs if she points to the location of a bag of heroin during a drug transaction; of murder if, on hearing of P's plans to kill his wife, she utters the words "oh goody"; and of working without a permit if she attends a concert by a musician who has overstayed their visa.⁴

(Kaisermann 2021, p127)

Work on accomplice liability takes seriously that accomplices can and do play a role in the actions that constitute crimes, but considers whether the model that exists properly accounts for the agency involved. Kaisermann, for instance, argues that we should abandon the notion of accomplice liability, and instead 'introduce a new set of crimes, which explicitly proscribe partial responsibility for harms' (2021, p154), understanding aiding and abetting as ways of being partially responsible. We would look at the causal contribution of various agents in apportioning responsibility, rather than identifying who pulled the trigger, literally or metaphorically. Although I will argue in Chapter 3 that there is something particular and important about being an agent of an action when we consider moral

³ For instance, (Gardner 2007), (Kutz 2007), (Girgis 2013), (Kaisermann 2021), (Dyson 2022)

⁴ For these examples, Kaisermann cites in turn – State v. Duran 526 P.2D 188 (N.M. Ct. App. 1974), Alexander v. State 102 So. 597, 598 (Ala. Ct. App. 1925), United States v. Ortega, 44 F.3d 505 (7th Cir. 1995), Obiter in R v. Giannetto [1997] 1. Cr. App. R. 1, Wilcox v. Jeffrey [1951] 1 All E.R. 464 (K.B.).

responsibility for that action, this kind of approach opens up broader possibilities for how we might identify the agents of an action. Kaisermann considers a case in which

D hands P a gun, assuring him that it is filled with blanks. In fact it's filled with live ammunition, and D knows this. On D's encouragement, P fires the gun at V, intending to merely frighten her. V dies immediately.

(2021, p128)⁵

Our instinct here is likely that D has a far greater degree of responsibility for V's death than P, but if we are working within the bounds of accomplice liability, D cannot be charged with murder unless P also is, and we are likely to think him guilty of a lesser crime. Kaisermann suggests that his model could allow that 'D will be fully responsible for V's death, even though it was P, and not D, who administered the fatal blow' (2021, pp148-9). Although it is not framed in these terms, this makes space for the possibility that we understand D, not P, as the agent of the action. The work on accomplice liability is not work in the philosophy of action, and its interest is primarily what best serves the ends of our legal system (although Kaisermann is also engaging with the metaphysics of causation). Nonetheless, by looking at the differing roles that are played by principal and secondary parties, it recognizes and engages with the ways that agents can be involved with the actions of others, and how we might best understand and account for that.

The second strand of literature that looks more closely at the role that agents can play in the action of another agent uses the notion of scaffolding. This literature does not use the term 'scaffolding' in exactly the way that it is deployed in empirical work on education and development, which I will look at in Chapter 7, although it is clearly influenced by this. McGeer (2018) draws on Ryle's discussion of intelligent capacities, in which he says '[t]o be intelligent is not merely to satisfy criteria, but to apply them; to regulate one's actions and not merely to be well-regulated' (2009, p17). She is interested in the way that this kind of regulation could happen externally, proposing elsewhere that 'our attitudes

⁵ The challenges of this example echo *R v Cogan & Leak*, [1976] QB 217, in which 'The defendant L took the defendant C back to his home and told his wife that C wanted to have sexual intercourse with her and that he was going to see that she did. L's wife was not willing to have intercourse with C but she was frightened of L who made her go to the bedroom where C had sexual intercourse with her' *R v Cogan & Leak*, [1976] QB 217, p217. C successfully appealed his conviction of rape, since he believed she had consented. L's attempted appeal for his conviction of aiding and abetting rape was unsuccessful, however, 'on the ground that it was contrary to justice and common sense' *R v Cogan & Leak*, [1976] QB 217 p218 to quash it, since he knew that his wife had not consented. However, this leaves the situation as one in which L aided and abetted a rape, L's wife was raped, but no-one is guilty of that rape.

and practices of holding responsible play a critical role in developing and sustaining our capacity to recognize and respond to moral reasons' (McGeer 2019, p313). She argues that the reactive attitudes of others play a scaffolding role in one's own ability to engage in moral agency. She sees this scaffolding as having a developmental role, being how we learn these things, but drawing from Ryle she notes that '[i]t is of the essence of intelligent practices that one performance is modified by its predecessors. The agent is still learning' (Ryle 2009, p30). The kind of capacities she is interested in, of which moral agency is an example, are ones that are in a process of continual development, not ones that we acquire and then cease to learn. Given this, not only do we first learn these skills and capacities through scaffolding on her picture, but also we sustain them in virtue of scaffolding — being involved in ongoing practices of being held responsible, being the target of reactive attitudes related to our morally significant behaviour, is part of what enables our ongoing ability to be moral agents.

Kukla (2021) has recently applied a scaffolding model to sexual consent. They suggest that:

A partner can protect a vulnerable partner's agency by giving competent uptake across the board: being a good, skilled, and caring interpreter of their expressions of pleasure, pain, desire, embarrassment, comfort, fear, and the like. This means that one and the same person—someone with dementia, for instance—may well be able to have consensual sex with one partner, who has these skills and competencies and is committed to exercising them, and not with another.

(Kukla 2021, p285)

This comes alongside a broader argument that agency can be scaffolded, which draws on Lindemann's (2014) notion of 'holding another in personhood'. Lindemann introduces this idea in exploring the role that social relations play in personal identity, and to hold another in personhood is to contribute to the development and maintenance of their identity, through sustaining narratives about who they are. Holding another in personhood can 'support an individual in the creation and maintenance of a personal identity that allows her to flourish personally and in her interactions with others' (2014, *x*) but it can also 'hold people in invidious, destructive narratives' (2014, *x*). Holding someone in personhood impacts what they can do and who they can be.

Both Lindemann and Kukla are at times interested in contexts in which the individual who is being scaffolded is unable to act in isolation, with both attending to dementia. Although Kukla draws on the

literature on relational autonomy, these pictures are not only about the way in which structural and social features can impede or enable someone's autonomy, but rather on the way that close interpersonal relationships can provide the support needed to help someone realize their agency. This literature gestures at an idea that is at the core of this project — the agency of another, directed at helping one realize one's ends, can enable and extend one's agency.

These bodies of literature recognize and engage with ways that we can be involved in the action of another, and ways that another can be involved in our action. Through the notion of dependent intentional action, I will situate considerations of this kind of involvement squarely in the philosophy of action, offering a close examination of the nature of actions in which one agent directs their agency at supporting another achieve their ends.

The structure

Before I can characterize and discuss dependent intentional action, I need to set out the kinds of intentional action that are ordinarily discussed — individual intentional action and joint intentional action. Doing this enables me to clearly distinguish dependent intentional action from other kinds of action, alongside showing the lay of the land in the philosophy of action. Characterizing these kinds of action is the work of the first two chapters.

Chapter 1 sets out the understanding of intentional action and intentional agency that will run through the thesis. In large part, it does this by distinguishing it from proximate phenomena. The first distinction is between agency in general and intentional agency in particular. Once I have set out conditions for an event's being an intentional action, I look at further distinctions — between actions and mere bodily movements, and between actions, mere attempts at action, and failures to act. Drawing these distinctions serves to further explain the nature of intentional action, but also to characterize these proximate phenomena, which will be drawn on later in the thesis. I then discuss two central components of intentional action – intention and practical reason. Throughout this chapter, I draw on a variety of approaches to intentional action, surveying the way that these distinctions and concepts are understood in the philosophy of action. The picture of intentional action that emerges from this chapter is not one that is tied to a particular approach to the philosophy of action, but instead is designed to be compatible with these various approaches.

Chapter 2 approaches joint action in a similar way, although the focus here is slightly different. Chapter 1 offered a general picture of individual intentional action as part of a ground-clearing exercise, stipulating how certain terms would be used and understood. Chapter 2 prioritizes detailing the various existing approaches to joint action. These are set out in service of offering a picture of joint action that is compatible with all of them, understanding these actions as ways of realizing intentions to act together. Again, as well as serving to clarify the nature of joint action, this provides a picture that will be returned to later in the thesis in arguing in favour of understanding dependent intentional action in a particular way.

Chapter 3 continues to explore the ways in which agents can be involved in actions, considering the relationship that an agent must bear to an action in order to be morally responsible for said action. It argues that moral responsibility for an action requires identity with the agent of that action, in response to arguments that suggest a different relationship, that of ownership, is sufficient. This conclusion is important in and of itself, having consequences for how we apportion praise and blame for action. It also provides a valuable guide in situations in which it is unclear who the agent of an action is. For instance, we can find ourselves inclined to hold the agent that coerces, rather than the agent that is coerced, responsible for an action, and this chapter distinguishes the kind of responsibility that agents have for actions as actions from the kind of responsibility that agents have for the actions of those they coerce, arguing that only being the agent of an action can ground the former. This discussion acknowledges some of the ways in which agents can be involved in the actions of others, which will be a focus for much of the rest of the thesis.

Chapter 4 introduces dependent intentional action. The task here is twofold, both identifying this as a familiar kind of action and offering a characterization of this kind of action. This is partly stipulative — there are actions that are close to dependent intentional action, but that do not fall within the boundaries I am defining. This chapter spells out the intentions with which the agents involved are acting, the relationship between their contributions, and the nature of the helping that takes place in dependent intentional action.

Chapter 5 defends the remaining element of this characterization — the claim that dependent intentional actions are actions of the supported agent. This is done, in the main, by mapping out the potential characterizations of these events, and then arguing that none of the alternative explanations

can properly capture the nature of dependent intentional actions. Resisting potential challenges to this understanding also gives positive reason for understanding these as actions of the supported agent.

Chapter 6 considers the agency of young children. In particular, it explores whether there are barriers to young children engaging in either individual or joint intentional action. The idea that intentionally ϕ -ing requires that the agent know how to ϕ is considered and rejected, removing this as a reason to think young children cannot engage in individual intentional action. Joint intentional action is argued to rely on understanding others as intentional agents, and the rest of the chapter explores whether young children are able to so understand other agents. This chapter draws on empirical work in building its argument, looking both at evidence that young children can understand others as agents, and evidence that they do in fact engage in joint intentional action. At the end of this chapter, it is left possible that young children can act intentionally, but no argument has been offered that they do in fact do this.

Chapter 7 argues that young children can and do act intentionally, at least when engaging in dependent intentional action. As in chapter 6, the argument here is supported by empirical work, looking at the role that scaffolding plays in education. The scaffolding agent is shown to play the role of the supporting agent of a dependent intentional action. This is expanded from the pedagogical context in which scaffolding is ordinarily considered, looking at executive function as a kind of support that can be provided, enabling young children to act intentionally. In previous chapters, dependent intentional action is primarily considered as occurring between ordinary adult agents. This chapter sets out the possibility that dependent intentional action can facilitate the intentional action of an agent who could perhaps not intentionally act on their own, namely, a young child.

Chapter 1 – Individual intentional action

In this thesis, I am interested in kinds of active involvement in human agency and human actions. The primary kind of active involvement is, of course, being the agent of an action. In this first chapter, I will set out how I am going to understand individual intentional agency and individual intentional actions throughout the thesis. This use of ‘individual’ indicates that these are actions of a single agent, acting alone. I will not be arguing for a particular account of the nature of this kind of action, but rather giving the general model of this kind of action that I will be using throughout the project. I will begin by introducing the notion of intentional agency in §1.1. In §1.2, I will set out how I will be understanding intentional action. I will be offering the following requirements for individual intentional action (noting that sometimes $\psi=\phi$)

Intentional Action

An event X is an intentional ϕ -ing of an agent A only if

- (1) X is an action of A 's ϕ -ing
- (2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting
- (3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

I will then turn to a series of closely related considerations. I will begin by considering two questions that have each been central to a strand of the literature in the philosophy of action. The first, which I will discuss in §1.2.1, is how intentional bodily actions are to be distinguished from other kinds of human bodily movements. The second, which I will consider in §1.2.2, is how a successful intentional action differs from a mere attempt to engage in intentional action, and from a failure to engage in intentional action. I will not be fully settling these questions, since this would require choosing between various approaches to the nature of intentional action. Rather, I will be clarifying the questions that are being asked and identifying the phenomena that are being distinguished from intentional actions. I will then consider two elements of intentional action. In §1.3, I will discuss the nature of intention, and the role it plays in intentional action. Across this and §1.3.1, I will clarify the content of intention and the limitations on what an agent can intend. In §1.4, I will turn to the notion of ‘practical reason’, discussing how this term should be understood. I will then discuss approaches to the difference that

practical reason makes to the nature of agency, setting out the transformative approach in §1.4.1 and the additive approach in §1.4.2. At the end of this chapter, I will have cleared the ground with regards to the nature of individual intentional action and several proximate phenomena.

1.1 Intentional agency

Agency is the capacity to act. To act is to make things happen, by bringing about a change in oneself and thus, ordinarily, further changes in the world beyond oneself. An action is an event of one's acting — an action is not what an agent does, an action is an event of their doing something. In this thesis, I am interested in the kind of agency that humans have, and the kind of actions in which they can engage. There is a tendency to use 'agency' narrowly, 'to denote the performance of intentional actions' (Schlosser 2019), and it is *intentional agency* and *intentional action* in which I am primarily interested. This kind of agency is possessed by most ordinary adult human agents, and is the capacity to act intentionally. To clarify what this kind of agency is, I will begin by discussing other kinds of agency, as a precursor to distinguishing the kind of agency I am interested in from the broader category. I will then say a little about what it is to be an *intentional agent*, a term which I will use simply to denote beings with intentional agency, that is, with the capacity to engage in intentional action. I will then turn to two broad approaches to the difference between intentional agency and the agency possessed by other organisms, especially non-human animals. There are some non-human animals who are possibly intentional agents — great apes, corvids and dolphins, for instance.⁶ However, my interest, in keeping with the dominant tendency in the philosophy of action, is in human intentional agency, and throughout the thesis references to intentional action and intentional agency are to be taken as such.

The capacity to act is, of course, not isolated to humans. We describe the behaviour of our pets and other non-human animals using action-verbs (they chase, they hunt, they hide). We describe nations as 'acting unilaterally' and law firms as 'acting on behalf of' corporations (implying the corporation too may be capable of acting, something we also suggest when we, for instance, describe one as 'taking over' another). This language is in part a shorthand for claims about the actions of individual intentional agents, but there is also a real sense in which the agential power of a nation outstrips that

⁶ Whilst I was redrafting this thesis, there were a spate of orca attacks on sailing boats off the Iberian peninsula. The explanations offered for these attacks arguably involve an attribution of intentional agency to the orcas involved (Hoare 2023, for instance), either through the invention and repetition of a novel kind of play, or as a group response to individual orcas previous traumatic experiences with other boats in an attempt to prevent further occurrences.

of all of those whose decisions and actions contribute to the invasion of a foreign country, for instance. In environmental history, there are arguments that we should understand agency as a concept that extends beyond even non-human animals. Foltz states that '...humans were not solely responsible for the spread of [cotton] to areas beyond its native territory. Cotton itself was certainly a major actor in this story, since it either flourished or didn't and in doing so affected the fortunes of humans who had invested their money and energy in cultivating' (2003, p9) and that '[c]otton plants have also been direct competitors with humans for vital resources such as water...' (2003, p10) — suggesting that the best way to understand certain historical events and periods is to understand cotton as a historical agent. Marx and Engels describe the productive forces of the contemporary economy as if they have a kind of agency that far outstrips that of those that created them:

Modern bourgeois society with its relations of production, of exchange and of property, a society that has conjured up such gigantic means of production and of exchange, is like the sorcerer, who is no longer able to control the powers of the nether world which he has called up by his spells. For many a decade past the history of industry and commerce is but the history of the revolt of modern productive forces against modern conditions of production...

(Marx and Engels [1848] 2011, p225)

We also often use action-verbs in the description of events on the micro-scale — we say that amylase 'converts' starch to simple sugars, and that molecules 'bombard' others and cause them to move.

Unlike the action of cotton or of microscopic particles, the action of humans is guided by intentions, and is often the outcome of practical reasoning. When we think about persons as agents (that is, as things with the capacity to act) we are thinking about them as things that act in accordance with their will. I am not here reifying the will — I am not presuming it is some entity to be theorized, but instead using the term as shorthand for an individual's preferences, desires, intentions, plans, and other conative states and attitudes. Human action involves bringing about change in the world in the hopes of bringing things into accordance with one's will. This will might coincide with basic needs, but it need not, and here human action differs drastically from that of non-conscious agents, but also that of many non-human animal agents. It is not merely the altering of one's environment or of one's situation because that is what it is in one's nature to do, or even in order to better promote the chances of survival of oneself or things of one's kind. When human agents bring about states that we need, we

ordinarily do it intentionally – because we recognize that it would be better for us for things to be a certain way, and we act with an intention to have things be that way. Our relationship to the end state (of satiation or security, for instance) is one that often involves conceptualizing that as a desired state, and calculating how it is we can bring this to be.

Intentional agency also involves the bringing about of states of affairs that we *want* rather than merely need. Molecules move from areas of high pressure to areas of low pressure, plants produce seeds in locations well suited to the flourishing of their species, and predators build traps for their prey. Humans eat ice cream and sign up for Netflix subscriptions and buy knick-knacks, and they do these things because they *want* to do them. They do this not only when they don't need to do these things, but sometimes even when they arguably need not to do them — it would be better for someone in a cramped flat not to buy more curios with which to decorate it, and yet they intentionally buy them. They also act in the hope of far bigger goals — of passing an exam, securing a job, developing a relationship. Part of understanding intentional action is understanding what needs to be the case for a person to act on the world in this way, in a way that is aiming at bringing some aspect of the world in accordance with their will. Whether that happens will rely on many things that are not features of their action, but if things go as hoped, these actions bring things into accord with one's will. Before we can get there, we need first to be able to carry out these actions at all – we need intentional agency.

Intentional agency is the capacity to engage in intentional actions. To begin, then, I will explore what is involved an action being intentional, and how such an action is understood to differ from non-intentional action.

1.2 Intentional action

In this section, I will present the general understanding of individual intentional action I will be using throughout this thesis. If this were a project of a different kind, one in which I built a detailed analysis of the nature of individual intentional action, I would not be able to pass over detailed argument at this point. What I need, however, is not an argument in favour of a particular account of the nature of individual intentional action. Instead, I am describing the phenomenon that such accounts are designed to capture. I will thus make some stipulative moves in this section, and in the rest of this chapter, as part of a process of ground-clearing that will indicate how I am understanding things going forward.

Intentional action is a subcategory of action, which is itself, I will take it, a subcategory of event. There is some disagreement with this second claim,⁷ but the status of actions as events is generally accepted within the literature.⁸ What I need to do, then, is to say something about the difference between events in general and actions in particular, and between actions in general and intentional actions in particular. To act is to make something happen, an action is an event of an agent's making something happen. As embodied agents, we do this primarily by acting on and with the world in bodily ways. We act through our bodies on the world — we do not act *with* our bodies, since this is to present them as an instrument. Bodies are not things we act with, they are how we act. Not all human actions are bodily — we also engage in mental actions — but these are rarely purely mental, they go on to inform and impact bodily actions. I will, going forward, focus on *bodily action*, that is, *action that involves movements of (parts of) the body*. This is in keeping with much of the literature, and I will often use 'action' as a shorthand for 'bodily action.'

Following O'Brien, I will take it that the following are necessary conditions on some an agent's acting:

Agent Condition. I must exist when I act: I must exist if I raise my arm.

Change Condition. I change things when I act: I change my bodily position when I raise my arm, and change things caused by my changing.

Self-change Condition. I self-change when I act: I change, from having my arm down to having an arm up, when I raise my arm.

Active Condition. The self-change which is my action is up to me: The change from having an arm down to having an arm up is *up to me*.

(2017, p266)

The first of these captures the thought that acting must be the acting of someone. The second captures that acting is making something happen — if there is an action, something changes from how things would have been without that change. The third captures that the change involved in acting includes change of the agent. The fourth condition captures that the kind of change that acting involves is a change of an agent, by an agent. It is not enough that the agent is changed in some way, they must change themselves.

⁷ For instance, (Mourelatos 1978)

⁸ See (Davidson 2001c) for the seminal argument for this picture.

These conditions are by their nature minimal, but they are far from empty. They give us the following necessary condition on some event X 's being an action of agent A :

Action

An event X is an action of an agent A only if

- (A1) A exists at the time of X
- (A2) X involves a change in the world
- (A3) X involves A changing
- (A4) X is up to A

Not all ways of effecting change in the world through one's body are actions. Often, agents make a difference to the world by refraining from ϕ -ing, and this can be something they do intentionally. I might intentionally refrain from repeatedly refreshing my email whilst waiting for the outcome of an interview, and this might take a great deal of effort on my part. However, this refraining is not an action, it is inaction. It involves a change from how I and the world would be, if I were refreshing my email, but that counterfactual does not make it a change in me or in the world. I might ψ to help myself to refrain from ϕ -ing, but that does not make my not- ϕ -ing identical to my ψ -ing. If I distract myself by watching TV, the two are distinct. I might have distracted myself in some other way, but this would not change the nature of the refraining.⁹

As well as applying to certain absences of action, the term 'bodily action' also applies to cases in which we move more than just our bodies. Ford points out that:

Normally, what we are doing intentionally is not just raising our arms, or turning our heads, but transacting with things distinct from ourselves—we are filling or emptying something; closing or opening something; grabbing, pushing, pulling, bending, squeezing, flipping or pressing something; we are making or destroying something...

(Ford 2018, p699)

⁹ It would be remiss here not to cite Anscombe's (1966) brief dismissal of the idea that we can attribute what happens to what one does instead of that which one withholds from doing.

However, all of these are still things we do through bodily action — all of these things are done by moving one’s body. They are all bodily actions, it’s just that bodily actions are very rarely solely bodily — they also impact the things we act on, the bits of the world on which we act

These criteria, of course, apply to intentional actions, but there is also more to something’s being an intentional action in particular. Following O’Brien’s (2007b) approach to describing the nature of action, I will offer necessary conditions on some event’s being an intentional action, without suggesting that I can argue for the sufficiency of these conditions (or for the possibility of providing sufficient conditions). I propose the following requirements on intentional action:

Intentional Action

An event X is an intentional ϕ -ing of an agent A only if

- (I1) X is an action of A ’s ϕ -ing
- (I2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting
- (I3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

It is entirely possible that $\psi=\phi$ — agents often ϕ with an intention to ϕ , and this is intentional ϕ -ing. However, this is not the only intention with which an agent might ϕ (as I will discuss in §1.3) which is what this requirement captures — that an agent might intentionally ϕ whilst acting with an intention to ψ for some $\psi\neq\phi$. There are various approaches to understanding the key notions in both mine and O’Brien’s criteria, and I do not want to come down decisively on these, rather leaving it open that one could ‘fill out’ this picture however one sees fit. Despite this deliberate thinness, these criteria give substance to the notion of intentional action, and should help us to pick out the things to which this term refers. In the following sections, I will discuss several concepts that are closely related to action in general and intentional action in particular. I will discuss two proximate phenomena — mere bodily movements and mere attempts — to distinguish these from (intentional) action and to give an understanding of them that I will use going forward. Alongside this, I will discuss the differentiation of act-types. I will then discuss the nature of intention, and its role in intentional action. This will clarify the criteria I have just offered, giving sense to what it means to describe someone as ‘acting with an intention to ψ ’.

1.2.1 Bodily movements

In this section, I will discuss kinds of bodily movement that are not bodily actions, and touch on several approaches to this distinction. I will draw out the notions of movement_T and movement_I , as found in Hornsby (1980) and the idea of Wittgensteinian arithmetic.

Bodily actions are not the only ways that bodies move. My knee moves by reflex when hit, I blink and sneeze and cough, and parts of my body might be moved by external forces.

An agent may guide her paralyzed left arm along a certain path by using her active right arm to shove it through the relevant trajectory. The moving of her right arm, activated as it is by the normal exercise of her system of motor control, is a genuine action, but the movement of her left arm is not.

(Wilson and Shpall 2022)

The way in which the left and right arms move in this example are clearly different. The left arm's movement is what I will call a *mere bodily movement* — it is a movement of the body that is not an action. There are some movements of intentional agents that seem to sit between reflex movements and (fully) intentional actions. I might drum my fingers against the desk as I grasp for a word, or fiddle with my hair whilst anxiously waiting for a friend. If I do these things unknowingly, such that I could come to realize I am doing them, they are the kinds of things that O'Shaughnessy (2008) calls 'sub-intentional' actions. We might want to reserve the term 'mere bodily movement' for things like reflex actions, entirely removed from the agent's control, and understand there to be a category of non-intentional actions (including sub-intentional actions), like the absent-minded drumming of my fingers. For my purposes, litigating this boundary is unnecessary. Some movements are actions, they 'exhibit genuine agency' (O'Brien 2007a, p170), and some movements are not actions. Exactly which movements fall into each of these categories is not important, what is important is that this distinction exists, and that actions all fall on one side of it. What is up for grabs in determining whether the drumming of my fingers is an action is not what is involved in something's being an action, but whether this movement is an exhibition of genuine agency. We can set aside the questions about these boundary cases and recognize that there are some movements that are actions, and some movements that clearly are not.

As O'Brien puts it elsewhere (forthcoming), when I act, 'I, myself, move.' In mere bodily movements, my body moves, but I do not myself move. The difference between mere bodily movement and bodily action is often approached via a question given by Wittgenstein — 'what is left over if I subtract the fact that my arm goes up from that fact that I raise my arm?' (2009, p169e, §621). When taken as a coherent question warranting an answer (which is not how Wittgenstein presents it), it is a question about the difference between mere bodily movement and bodily action. The arm rising that is not an arm raising is a mere bodily movement, and 'what is left' is the difference between the two kinds of thing — mere bodily movements and bodily actions. One approach to the distinction between mere bodily movements and bodily actions, the one that takes the question drawn from Wittgenstein seriously, is that bodily action is bodily movement *plus some special factor*. Danto describes the problem as relying on the idea that 'an action [is] a movement of the body plus x... and the problem... is to solve in some philosophically interesting way for x' (1981, p5). This approach has sometimes been called 'Wittgensteinian arithmetic'¹⁰ and it relies on an assumption that when we talk of 'mere' bodily movement, we are speaking of bodily movement that would be part of an action, if only this additional factor were present.

A useful distinction can be drawn from Hornsby's (1980) discussion of the relationship between mere bodily movement and action. She draws on the grammatical notion of transitive and intransitive verbs, to distinguish between movements_T and movements_I. Transitive verbs are ones that have direct objects, and *to move* is clearly such a verb. Movements_T are movements that move something else, which moves_I. If I change desk in the library, I move_T my laptop, which moves_I. At first pass, we might think of the movement of the active right arm in (Wilson and Shpall 2022) as a movement_T, and the movement of the paralysed left arm as a movement_I. However, it would be more in keeping with the approaches that adopt this distinction to think of both of these as movements_I — the difference lies in what causes the movement_I of the right arm, such that it is an action, rather than the mere bodily movement of the movement_I of the left arm. Hornsby herself argues that actions caused by are movements_T of the agent who acts, which are over before the arm moves_I or the muscles contract_I. For Hornsby, movements_T are over before anything we would ordinarily recognize as movement occurs — they are whatever happens internally that causes those contractions and movements. These movements_T are the actions on Hornsby's picture — actions cause bodily movements_I, and mere bodily movements are bodily movements_I that are not caused by actions. Regardless of whether we

¹⁰ The earliest use of this term of which I am aware is in (Velleman 2000), it also appears in (Lavin 2015) (O'Brien 2017)

adopt Hornsby's model we can take from her the distinction between movements_I and movements_T. Actions will be a species of movement_T, whereas mere bodily movements need not be.

This is of course not the only way to understand the difference between mere bodily movement and bodily action through a causal lens. Some approaches adopt an event-causal approach, understanding actions as those bodily movements that are appropriately causally related to events and states involving the agent. On Davidson's (2001a, p5) account, for instance, actions are casually related to beliefs and desires of the agent.¹¹ Other approaches¹² understand actions as those movements that are appropriately causally related to the agent themselves, which does not entail understanding actions as caused by agents. Steward (2012) does not frame her account as one on which causation is central to explaining action. Nonetheless, she describes agents as having bodies, and an action as 'an exercise of [the agent's] power to make the body (or particular parts of the body) move' (Steward 2012, p32). The idea that the body which is moved is separate from the agent that moves it is in keeping with the framing we get from the various causal approaches, and with the contrast we find in Hornsby — Steward understands actions as movements_T (2012, p33), contrasted with the movements_I that are the results of actions.

Accepting that there is a difference between mere bodily movements and bodily actions does not require endorsing Wittgensteinian arithmetic or the idea that the body that moves is moved by me, where what I do happens at a degree of remove from that movement. Lavin (2013) resists this picture, arguing that the causal approach to action (alongside the idea of basic action), implies a Marxian alienation of the agent from their actions that does not exist. O'Brien (forthcoming) tells us that when I act, 'I, myself, move.' Here, she is drawing on Evans, who writes:

It is true that I manifest self-conscious thought... in action; but I manifest it, not in knowing which object to act upon, but in acting. (I do not move myself; I myself move.)

(Evans 1982, p207)

This understanding of how it is that I move in action clearly rejects the idea that what happens is the causing or bringing about of some movement of my body by something I do. Instead, I simply move

¹¹ Bratman's planning agency model (Bratman 1987, for instance) is also a version of event-causalism

¹² For instance (Alvarez and Hyman 1998)

— my acting is not reducible to a mere bodily movement plus something else. Haddock (2005, p164) understands bodily actions and mere bodily movements as determinates of the determinable ‘bodily movement’ — they are both kinds of bodily movement, but not in virtue of bodily movement being a common element of them.

In §1.2, I drew on O’Brien to give the following necessary condition on action: ‘(A4) X is up to A .’ All of the approaches I have touched on in this section are consistent with this. Where these accounts differ is in how they understand what it means for the change in the agent (that is, the movement of their body) to be up to them. To sum up — *mere bodily movements are those movements that are not up to the agent whose body is moving*.

1.2.2 Failures to act and mere attempts

In this section, I want to clarify two phenomena from which we must distinguish action — failures to act and mere attempts at action. These are both ways in which an agent that intends to ϕ might not succeed in ϕ -ing. I will set out the idea that it is only right to describe an agent as acting if that which they are doing actually happens, and then I will identify two key ways in which it might not happen. The central difference between the two ways in which an agent might not succeed is whether the action does not occur due to the circumstances surrounding the agent, or due to the agent themselves. This first kind of case is what I will call a *mere attempt*. It is what I also might call a *failed action*, but I will stick with the language of *mere attempt*, to leave open the possibility that (non-mere) attempts also play a role in successful actions. The second kind of case is a *failure to act*. These are both ways in which an agent might not successfully act. This discussion will also allow me to spell out how I am understanding act-types.

To describe someone as acting is to indicate that they did what they were trying to do.

‘Acting’ is a success verb. Just as you cannot be said to have seen an object unless the object exists, or know something unless it is true, you cannot be said to have acted unless there is something that has been done.

(O’Brien 2007a, p138)

When we act, we do something, something happens. But we don't do just anything. We do some particular thing, and it is only if that particular thing happens that we act — 'what has been done must match the description of what the agent is doing' (O'Brien 2007a, p138).

I am understanding actions to be events. In particular, they are events of agents acting. Actions are not things that agents do; what agents do, to paraphrase Anscombe, is what happens. I will, throughout this thesis, follow the convention of identifying act-types by Greek letters, primarily ϕ and ψ . When I act, I ϕ , and my action is the event of my ϕ -ing. Whether or not I have successfully ϕ -d will depend on what it is, on this occasion, to ϕ . If I make a cup of coffee, whether or not this counts as my ϕ -ing will vary depending on what ϕ -ing is. If to ϕ is to make coffee, or to make a hot drink, then I have successfully ϕ -d. These two act-types have different levels of precision, and this isn't a problem — sometimes what we set out to do is something very precise, sometimes it's much more vague. What matters in assessing whether someone has succeeded in acting is how they would describe that which they were doing. If I described myself as making a cup of tea, and I make a cup of coffee, what is done does not match what the agent described themselves as doing, and so I did not succeed. I did not make a cup of tea. What is not clear is whether I merely attempted to make a cup of tea or failed to make a cup of tea.

Many actions exemplify what is known as the imperfective paradox.¹³ This is the term for the perhaps surprising illegitimacy of the inference from a claim with imperfective aspect about what someone was doing to a claim with perfective aspect about what they have done. When we claim that someone was running we can infer that they ran — the truth of the former entails the truth of the latter. We cannot, however, make a similar inference from the truth of 'A was running-a-mile.' Running-a-mile is not homeomeric, that is, not all parts of it, taken in isolation, are themselves running-a-mile. In contrast, all parts of running are themselves running. Given this, that someone is running-a-mile at some moment does not entail that they succeed, it does not entail that they do it. Someone can have been running-a-mile, that is, roughly, running such that if things had continued as normal then they would have run a mile, but never succeed in reaching that goal. In some cases, the most we will be able to say is that they merely attempted to run a mile. In this section, I want to identify the cases I call *mere attempts*. A *mere attempt* (to ϕ) is what happens when an agent tries to ϕ but fails. Mere attempts are the events of an agent's attempting to act, just as actions are the events of an agent's acting. They are

¹³ See (Vendler 1967) and (Mourelatos 1978) for the classic formulation of this phenomenon

often also actions of a different kind, but they are not ϕ -ings, and this is what matters for the purposes of identifying if they are successful.

I want to distinguish between a *mere attempt* and a *failure to act* as a way of identifying the former category. Mere attempts happen because agents attempt to act. Imagine a case in which the train I am aiming for is cancelled as I walk down the hill to the station. I attempt to catch the train, but it leaves before I reach the platform, and this attempting involves my leaving the house, walking to the station, calculating when and where is best to cross the road etc. This is why attempts are not *failures to act* — they are not passive — what they are is *failed actions*. Although they often involve the agent's acting, instances of acting are, as I have noted, not interchangeable. It is not enough to act in some way, what matters is that the agent acts in the right way. In mere attempts, there is no event of the requisite act-type. There is nothing that constitutes a catching of a train. There is, however, an attempt to catch a train. I have described a case here in which it is very clear to the agent that they are attempting to catch the train — the possibility of failure is something that is present for them. What makes it a mere attempt, though, is not that they understood that they were trying hard to do something, but that they failed to do it. Hornsby (1980, p34) argues that we should not infer from the fact that we describe some situations as an agent trying to do something, that they only try in those situations. It might be that there are attempts in successful actions and in failed ones, but only in failed actions are there *mere attempts*.

Broadly, there are two kinds of ways I could not succeed in ϕ -ing. I could not succeed to ϕ in virtue of my (lack of) commitment to ϕ -ing, or in virtue of something preventing me from ϕ -ing. Either I fail to do something I need to do to carry out that action which is available to me to do, or something prevents me from carrying out that action — the circumstances necessary for my ϕ -ing do not obtain. I might fail to catch the train because I do not do something necessary to catching it. This might happen early on, I might not set my alarm the night before, I might ignore it when it goes off, or I might get sidetracked by playing a video game and lose track of time. Alternatively, I might make it out the door, cutting it a little fine, and then choose to stop and chat to a friend. My failure to ϕ is down to how I reacted to my circumstances. In at least some of these cases, we will be reticent to describe me as having made an attempt to ϕ . At least, if we can describe me as changing my mind about ϕ -ing, or deciding not to ϕ after all, this does not seem to be an attempt to ϕ . Rather, it seems to be not- ϕ -ing-at-all. These are cases where the fact that it was not an event of ϕ -ing in the end seems to track backwards – I was never really ϕ -ing. Of course, the scale of the act-type that ' ϕ -ing' is will

determine how likely we are to make these claims. If ‘ ϕ -ing’ is ‘working on a PhD thesis’, and I decide to leave the programme after 3 years, it seems hard to say that I was never really doing it at all. If, however, ‘ ϕ -ing’ is making a cup of tea, and I decide I want a cold drink by the time I get to the kitchen, it would seem completely implausible to say I attempted to make a cup of tea. If I leave the PhD programme because the cost of living makes continuing unfeasible, or I do not make a cup of tea because it transpires I have no tea bags left, then it seems more obviously correct to say that an attempt was made than in the cases where I simply change my mind about what I want to do. I did everything I could to successfully ϕ , but circumstances were such that I could not in fact ϕ . This is how we understand what happens when the train is cancelled at the last minute. In this case, the agent did everything they would do in successfully catching the train, except for catch the train. A mere attempt is a case in which an agent does what they need to in order to ϕ , but they fail to ϕ because they are prevented from ϕ -ing. The agent may well do everything that would otherwise be needed for them to ϕ , but they do not ϕ . A mere attempt is what happens when an agent is prevented from acting, rather than when the agent fails to act. The distinction is that *failures to act are when an agent’s not ϕ -ing is in virtue of their not making an effort to ϕ , mere attempts are when the world does not cooperate with an agent’s efforts to ϕ .*

What I call *mere attempts* are a central consideration for those that Grünbaum (2008) dubs the ‘The New Volitionalists.’ He identifies a set of core claims to which these approaches are committed, including that ‘[e]very time an agent intentionally moves her body (and she does so every time she is acting intentionally), her bodily movements are caused by her trying to do them’ (Grünbaum 2008, p68). The consideration of mere attempts (or, as Grünbaum puts it, total failure) leads to accounts on which the same thing happens in these and in successful actions. Armstrong (1973, p5) suggests ‘that trying is involved in all cases of intentional action,’ and Hornsby (1980, p33) argues that ‘[e]very action is an event of *trying* or attempting to act.’ O’Shaughnessy explicitly states that trying is ‘the same kind of event both on the occasion of success and on the occasion of failure’ (1973, p373). Others dissent from this image of attempts as featuring in just the same way in failed and successful action, with Brewer, for instance, arguing that:

[f]ailed attempts by an agent to move in various ways... are to be understood only as derivative of [successful attempts], not as psychologically identical with an independent physical consequence unfortunately missing.... The connection is not a shared, inner mental happening.

(Brewer 1993, pp310-11)

Despite the substantial differences in how they understand the relationship between the two kinds of event, Brewer and the New Volitionalists nonetheless agree that we can distinguish between successful actions and mere attempts at action. Of course, when an agent merely attempts to act in some way, they likely act in some other way. There is an event of their trying, and that event is the mere attempt. It is not enough that an agent acts, they need to act in the aimed for way. Conditions (A2) and (A3) give us the importance of an agent's actually changing (both the world and themselves) in their acting, but the concept of a mere attempt brings with it that there is some particular thing that the agent set out to do. This need not be something they intend to do, since it is clear that agents with more minimal kinds of goal-directedness can try and fail, but nonetheless, the first of my conditions on intentional action captures something important – it needs to be an event of the right type, it needs to be a ϕ -ing for the agent to successfully ϕ . An agent can only be said to have acted if the change they bring about is of the desired type — a mere attempt might well produce change, but not change of the right kind. Given this, *mere attempts are attempts to ϕ that do not culminate in the agent's ϕ -ing.*

1.3 Intention

In the background of the previous section was the idea that there are particular things agents set out to do, and in this section I want to consider a key form that this takes – intention. As Anscombe famously notes, we use the concept of intention in a variety of ways — we speak of ‘expression[s] of intention¹⁴... of an action as intentional.. [and] ask with what intention the thing was done’ ([1957] 2000, p1, §1) and these uses of the concept are not equivocal.¹⁵ The second and third of Anscombe's categories are my primary focus. In this section, I will discuss what it means for an agent to be acting with an intention, and how that relates to their acting intentionally. I will consider several approaches to the nature of intention, and clarify how I will be understanding the content of intentions.

There are two ways that an action can be intentional. It can be done with an intention to so act, or it can be done intentionally without such an intention. I will start with the first of these. I will not make

¹⁴ ‘Expressions of intentions’ are prior intentions — those formed and articulated before the action begins. I will be setting these aside, focussing instead on the intentions with which agents act, that is, on intentions-in-action. McDowell argues that ‘when acting is executing a prior intention, an intention in action is what the prior intention *becomes* when the time it determines for action arrives (2015, p148). Of course, intentional action is not necessarily dependent on prior intention, but the intention in action (the intention with which the agent acts) of those actions that do not depend on prior intentions should have the same character as the intentions in action of those actions that do so depend.

¹⁵ Although, as Lavin (2015, p614) reminds us, Anscombe does later remark that ‘[t]o a certain extent, the three divisions of the subject made in §1 are simply equivalent.’ ([1957] 2000, p40, §23)

any assumptions about the ontological status of intentions. We might want to insist that intentions are concrete and distinct mental states, but we might want something much more minimal. Whatever the metaphysics of intentions, we can talk about them as aspects of action. I will not presume that the intentions with which agents act must be prior intentions — they need not be formed in advance of the action. What I will stipulate at this point is the content of intentions. If they are mental states, this is the content of those mental states, if they are descriptions of actions then these are the content of those descriptions. Intentions, I maintain, should be understood as having infinitival content — an intention is an intention *to* ϕ . This is the ‘act-directed’ approach to the content of intentions. The alternative to this is to treat intentions as having propositional content — an intention is an intention that P , where P is some proposition.¹⁶ P might be the proposition ‘ $I \phi$ ’ or it might be about some state of affairs that my ϕ -ing will bring about, such that my intention would be an intention *that* $I \phi$ or *that* X *is the case*.

The claim that the content of intentions is an act-type, which is ordinarily expressed as the infinitival form of a verb, more neatly fits the way that we ordinarily articulate our intentions. Agents intend to do things, not to bring about states of affairs, except by doing something that brings about that state of affairs. Madden (2011, 2014) argues that the act-directed approach has further merits. It allows for identity between the content of intentions of different agents, such that we can speak of agents as acting with the same (type) intention. Both the propositional and act-directed approaches can allow that the same person has the same intention at different times. If I intended to go for a walk yesterday and intend so again today, the relevant proposition in both contexts is ‘that I go for a walk.’ The act-directed view, however, also easily allows for different agents holding the same intention. If I intend to go for a walk, and so too does my brother, then we both intend ‘to go for a walk.’ If our intentions are propositional attitudes, it is not straightforward that my ‘I will go for a walk’ and his ‘I will go for a walk’ are the same proposition, since the ‘I’ refers to me in the first and him in the second. If we both go for a walk, it seems right to say that we have done the same thing, and so it seems we should think of us as both realizing the same intention.¹⁷ Given the precedent for and arguments in favour of this move, and its value for arguments I will make in Chapter 2, I will take intentions to be act-directed throughout the thesis.

¹⁶ This view is often assumed rather than argued for — Davidson, for instance, includes intending in a list of verbs expressing propositional attitudes (Davidson 2001b, p 310)

¹⁷ (Madden 2011) and (2014) offer further arguments against taking intentions to have propositional content and in favour of the act-directed view, as do (Baier 1970), (Rumfitt 1994), (Thompson 2008) and (Campbell 2019)

The first way of intentionally ϕ -ing, then, is ϕ -ing with an intention to ϕ . Intentions are *realized* by intentional actions, and when an intention is realized, what was intended has been achieved. The following conditions holds:

Realization

An event X realizes an agent A's intention to ϕ only if X is A's ϕ -ing

The *only if* here is important, since there are many cases of someone's ϕ -ing that nonetheless fails to realize their intention to ϕ ,¹⁸ or when they have no such intention. This principle is also compatible with someone's intentionally ϕ -ing not requiring them to act with an intention to ϕ — this ϕ -ing contributes to their realizing their intention to ϕ , but it does not in and of itself realize that intention.

There is a second way for a ϕ -ing to be intentional. Bratman (1984) draws out the ways in which intentionally ϕ -ing does not require that an agent acts with an intention to ϕ . Firstly, he shows that acting with an intention to ϕ is more demanding than intentionally ϕ -ing. He imagines someone simultaneously playing two video games, in each of which they are aiming at a target, and that are linked such that if 'both targets are about to be hit simultaneously the machines just shut down' (ibid, p382). If hitting the target in each game is sufficiently difficult, it could nonetheless increase the player's chances of success to play both at once, meaning that they are trying to hit both targets. If the player successfully hits one of the targets, they do so intentionally. However, this cannot require that they are acting with an intention to hit that target, since this would require that they are also acting with an intention to hit the other target, since if they hit that one instead, they would have done so intentionally. This would leave the player holding inconsistent intentions, and thus being irrational, since they would be intending to do two incompatible things. Given that in cases such as this, an agent can intentionally ϕ by trying to ϕ or χ , an agent need not be acting with an intention to ϕ in order to intentionally ϕ . Instead, I can intentionally ϕ in the course of realizing an intention to ϕ .¹⁹ Bratman (ibid, pp399-400) also discusses the way in which 'wearing down my sneakers' (ibid, p400) might be intentional, if done in the course of running a marathon. This is especially clear, he suggests, if we take it that

- (a) 'I consciously note' (ibid) that I am wearing down my shoes whilst running and
- (b) 'wearing them down has some independent significance to me' (p400).

¹⁸ See §5.5 for further discussion of this

¹⁹ Where ϕ might be ' ϕ or χ ' — i.e. hit target 1 or hit target 2

In this context, wearing them down is something I do, and I do it in service of my end of completing the marathon, but I do not do it *with an intention to wear down my sneakers*. Rather, I do it by acting with an intention, and I do it *because* it contributes to the realization of the intention with which I am acting. It is not necessary to the realization of my intention. I could avoid it by wearing different shoes or perhaps changing my gait. It is, however, not merely incidental either, in the way that it would be if I have never considered that this something I would do in the course of running the marathon. If I had not considered the impact of this endeavour on my footwear, or if I perhaps believed the soles to be especially hard-wearing, it seems wrong to describe me as intentionally wearing down my sneakers. Intentionally ϕ -ing in the course of realizing my intention to ϕ involves what I will call *condoning* ϕ -ing — recognizing that I am likely to ϕ , and not taking steps to avoid it.

Whether or not I am ϕ -ing with an intention to ϕ , for my ϕ -ing to be intentional I must be acting with an intention. What it means to say that an agent is acting with an intention is understood in various ways, best characterized as different models of the kind of explanation an intention gives of an action. On causal views, the explanation is causal — when we identify the intention with which an agent acts, we identify something that plays a causal role in its occurrence. For Davidson (2001a), intentions are the ‘primary reason’ for which an action is done, and that primary reason is its cause. He went on to modify this view, understanding intentions as having an evaluative component — ‘in so far as a person acts intentionally he acts in the light of what he imagines (judges) to be the better’ (2001e, p22). This picture of intention as a judgement in favour of a particular state of affairs is still a picture of an explanation in terms of intention as a causal explanation — that it is judged to be better is what causes the agent to so act. Bratman understands intentions as features of plans, which are ‘mental states involving an appropriate sort of commitment to action: I have a plan to A only if it is true of me that I plan to A’ (1987, p85). His approach is functionalist (ibid, p34) and so intentions are to be understood in terms of the particular functional roles they play in plans. Asarnow (2020, p6) draws out three particular functional roles of intentions in Bratman’s picture:

- Intentions ‘settle practical questions’ - once you have arrived at an intention to act in a given way, you cease to deliberate about whether to act in that way and assume that you will so act. This is associated with the norm of ‘stability of intention’.
- Intentions ‘filter out inconsistent plans’ - once one has arrived at an intention to act in a given way, you ignore alternatives to acting in that way. This is associated with the norm of ‘consistency or agglomeration of intentions’.

- Intentions ‘pose means-end’ problems – once you have arrived at an intention to act in a given way, you will determine how to act in that way, and form intentions to carry out the actions you believe are means to satisfying this intention. This is associated with the norm of ‘means-end coherence’.

An agent acting with an intention is one with a mental state that serves such a role. Identifying the intention with which an agent is acting on this picture is identifying something that plays a causal role in the action’s occurrence.

Anscombe offers a radically different picture of the kind of explanation we give when we identify the intention with which an agent is acting. She tells us that intentional actions

are the actions to which a certain sense of the question ‘Why?’ is given application; the sense is of course that in which the answer, if positive, gives a reason for acting.

(Anscombe [1957] 2000, p9, §5)

There are two key things to note about this claim – that the question being given application does not require that the agent answers it, and that these reasons are not causes. For the question to be given application is for it to be appropriate to ask it of the action. It does not have application if the agent answers ‘I was not aware I was doing that’ (ibid, p11, §6) or ‘It was involuntary,’ (ibid p12, §7). The question is also refused application even by certain kinds of answers, namely ones that ‘impl[y] I *observed* I was doing that’ (ibid, p25, §17). However, if the agent tells us that they are doing it ‘just-because’, the question has application –

[A] possible answer to the question ‘Why?’ is one like ‘I just thought I would’ or ‘It was an impulse’ or ‘For no particular reason’ or ‘It was an idle action—I was just doodling’. I do not call an answer of this sort a rejection of the question. The question is not refused application because the answer to it says that that there is *no* reason, any more than the question how much money I have in my pocket is refused application by the answer ‘None’.

(Anscombe [1957] 2000, p25, §17)

Answers to Anscombe's question are reasons for action in that they are explanations that the agent offers of their own behaviour. These point to the action in which the agent is engaging through acting in this way. This is the intention with which you are acting — you are ψ -ing because you are acting with an intention to ϕ . 'Why,' I might ask you, 'are you breaking that egg?' '*Because I am combining wet ingredients*' 'Why are you combining wet ingredients?' '*Because I am baking a cake.*' Why are you baking a cake?' '*Because I am making pudding.*' Here, it seems, the questions would stop, although of course you could instead have answered '*Because it is early afternoon,*' '*Because it's Cheryl's birthday party tomorrow*' or '*Because I fancy some cake.*' These do not give an answer to Anscombe's question, in that they do not give the intention with which you are acting, but they can easily be refigured as '*Because I am making sure the cake is ready on time,*' '*Because I am bringing Cheryl's birthday cake to the party*' or '*Just because.*'

These answers are all reasons, but these reasons are also all further actions in which the agent is engaged. This does not entail that these answers all identify distinct actions; each of these is a redescription of what the agent is doing. The cracking-an-egg is the combining-wet-ingredients is the baking-a-cake is the making-pudding. What Anscombe's question gives us is the 'A-D series' of descriptions of an action, which place these explanations in an order 'in which each is dependent on the previous one, though independent of the following one.' (Anscombe [1957] 2000, p45, §26). Taking the examples I gave in the previous paragraph, we have:

- A - I am breaking that egg
- B - I am combining wet ingredients
- C - I am baking a cake
- D - I am making pudding

That what I am doing is combining wet ingredients depends on it being the case that I am breaking an egg, for instance — I cannot combine wet ingredients without doing this. However, there are other things I could be doing by combining wet ingredients, such that my doing this is independent of my making a cake. Each stage of the series gives us an intention with which the agent is acting, and the last gives us 'the intention *with* which the act in each of its other descriptions was done, and this intention so to speak swallows up all the preceding intentions *with* which the earlier members of the series were done' (ibid, p46, §26). The intentions that we find in this way are themselves actions, things which the agent is themselves doing, or which they aim to do. It might turn out that I am merely attempting to act, but assuming it comes off, baking a cake is in fact what I am doing in cracking the

egg. When looking for the intention with which an agent is acting we are not looking for something that plays a causal role, but something the agent is doing — we are looking for a course of action. When someone gives us a ‘just-because’ answer, it seems that we have hit the final explanation — the end for which they are acting is the end specified in the description of the action they have given. I am baking a cake just because I want to have baked a cake, or just because I want to engage in a process of cake baking. The baking done for this reason needs no further explanation, in much the same way as the baking that is the bringing of the cake to the party or the preparation of pudding needs no further explanation.

These two kinds of pictures give us very different understandings of how intentional action operates, and how to understand ‘intention’. From these arise two very different pictures of how intentions relate to the actions that realize them. For the first kind of approach, there is a *separation* between the intention and the action, with the former *bringing about* the latter, or at least the things that we ordinarily use to identify the latter. On the second kind of approach, though, we find ourselves with a very different image of this relationship. Lavin (2015) argues that we should read Anscombe as resisting the move that seeks to separate out the intention and the action and understand them in terms of their relation to one another. The intention with which the agent acts does not come apart from their action. Instead, intentional action is a particular kind of unity, one that is bound together by the teleology of the ongoing event of acting. This teleology is what the A-D series illuminates, it shows us the ends at which the agent’s acting is directed. On this view the intention with which the agent acts is not distinct from this unity or some causal element in it, but rather is a characterization *of* this teleological unity — it is how we identify that which binds the whole together. Whether we understand intentions as offering causal explanations or explaining what it is the agent is doing, we understand intentional actions as requiring that an agent is acting with an intention. It need not be an intention to do just that, but there is an intention involved.

There are two broad ways that an action might be intentional — either the agent ϕ s with an intention to ϕ , or they intentionally ϕ in the course of realizing an intention to ψ . Conditions (I2) and (I3) capture this — the agent must be acting with an intention, and they must be doing that which they are doing because it contributes to the realization of their intention. We might mean by this that their ϕ -ing plays a causal role in bringing about their ψ -ing, or that their ϕ -ing is explained by the fact that they are ψ -ing, but on either kind of understanding (I2) and (I3) will be true. Intentions are act-

directed — an agent intends to act in some particular way. Before I move on to discussing the role of practical reason I want to say a little more about the conditions on intentions.

1.3.1 Conditions on intention

Having said a little bit about the nature of intention and the role that it plays in intentional action, I want to say something about the kind of things an agent can intend. As I set out in the previous section, the contents of intentions are act types, but more can be said than this. I will detail the kind of act-types I understand agents as intending to engage in, and some conditions and limitations that exist for the content of intentions.

Imagine that I have an intention to bake a cake. It is an intention to engage in a non-basic action, and it contains no reference to the means by which it will be realized. A basic action is something an agent can do ‘just like that’, such as moving a limb, without needing to perform any further actions. ‘Baking a cake’ is given by Hornsby (1980, p68) as an example of a non-basic action. It consists of many sub-actions, each contributing to the overall action of cake-baking. This is the kind of act type in which agents usually intend to engage. It is rare for an agent to intend to carry out a basic action, and even more so to intend it as a basic action. I do not simply mean that ‘basic action’ is not a familiar concept outside of philosophy, which it is not, but also that we do not intend them under the descriptions that frame them as basic. I might intend to move my fingers across the keyboard, but that is not how I understand what I am intending or doing — I intend to type these sentences. The description of an action under which it is intended is the one under which the agent conceives of it. It is because I am baking a cake that I engage in the sub-actions that are necessary precursors to and components of the overall action. It seems unproblematic to assume that at least some cake-bakings are understood by their agents as cake-bakings, and that when they intend them, they intend to bake a cake. The content of an intention is that which the agent intends to do.

We can intend to do something without yet intending detailed means of doing it, and what we are intending does not change even if the means we pick change during the course of the action. Intending to bake a cake specifies nothing about how this will be achieved. The baking is the end, rather than a means to an end of having a cake. There will be conditions on how I expect to realize my intention — I am unlikely to be intending to bake a cake in the kitchens at Buckingham Palace — but these are not part of the intention. When I intend to do something that involves interactions with

my environment — such as kitchen equipment — that I intend to use the means available to me is not specified in my intention, because it does not need to be — this is the only way I could bake a cake. I can only intend to do something using the means available to me, but it does not follow that my intention includes the means. If the means were part of the intention, and I needed to make a small change, such as substituting kinds of sugar, either I would no longer be realizing my intention or my intention would change. Neither of these would reflect my belief that I had done what I set out to do. Interacting with the tools and obstacles that enable and impede our action is a feature of intentional action by an embodied human agent, but these particular tools and obstacles are not part of the content of intention. This means that my intention to bake a cake and your intention to bake a cake can be intentions with the same content even if we realize this intention by following different recipes, in different kitchens, and using different ingredients, and even if we expect to do those different things. They are intentions to engage in the same act-type.

There will be some cases in which the means are part of the content of the intention. If I intend to cross America by Route 66 but discover that large chunks of it are closed for maintenance, in finding an alternative route I will set a new intention, having failed to realize my initial one. If I intend to test the oven by baking a cake, I cannot realize this intention by baking a cake in a different oven.²⁰ This latter example fits the model of ‘by’ statements as ways of ‘form[ing] verbs out of verbs and verb phrases’ (Hornsby 1980, p7) — the intention is to carry out an oven-testing-by-cake-baking, where the oven testing is the end and the cake-baking the planned way of doing it. This contrasts with the ordinary case of cake-baking, as the particular oven is crucial to realizing the intention, but this is because it is not a cake-baking but an oven-testing. If the cake-baking is also part of the oven-testing intention, and it transpires I have no sugar, then in testing-the-oven-by-making-a-quiche I do not realize the original intention. A different, related, intention has been set and realized. If my intention was to test the oven, or to test the oven by cooking something, then a quiche will do just as well as a cake. If my intention is to cross-America-by-driving-Route-66, then with Route 66 closed I simply cannot realize this intention. The cases where the means are built into the intention are ones in which those means are part of what is intended. In saying that the means do not form part of the content of intention, I am not denying that these cases exist, but acknowledging that in these cases, these details are best understood as part of the ends. They do not feature in the intentions as means.

²⁰ Thanks to Pete Faulconbridge for these examples

There are two related limitations on what can be intended — I cannot intend something that I do not take to be realizable, and I can only intend to do something myself. The only actions that can realize my intentions are my own actions. The first of these does not mean that my intention has to, in fact, be realizable, I might be mistaken. However, if I believe that something is impossible for me to do, I cannot intend to do it. I can desire to do it, wish to do it, dream of doing it, but I cannot intend to do it. Intentions involve a commitment to doing that which they are about that is not a feature of these other conative states, and so these other states can be about things that are not possible in a way that intentions cannot. Relatedly, an intention can only be realized by the agent that has this intention. If someone else were to bake a cake in just the way that I intended to, it would not realize my intention. It might allow me to realize my intention to eat a cake of that kind, but my intention to bake the cake can only be realized by me baking it myself.²¹ What is intended is a ϕ -ing by the agent that so intends — we can only intend first-personally. This ‘*own-action*’ condition, which states that an agent can only intend to do something themselves, rather than intend for other agents, is built in to the content of intention.²² On the act-directed view, intending for another needs to go through some action of my own just in virtue of how they are constructed. Acting myself is what is intended – I can only intend to make (or coerce or encourage or induce or...) someone else (to) ϕ . In later chapters, when I discuss dependent intentional action, we will see how an agent’s actions can be guided by another agent’s intentions, but this is not the same as their being intended by another agent. Even in these cases, each agent’s intention is directed only at themselves and their own acting. I cannot simply intend someone else to ϕ . I can only intend to act, where this intention is an intention to act *myself*. This first-personal element is not explicit in the intention because it need not be. As Madden puts it, intentions are “‘directed” or “sent” to oneself” (2014, p82), and this involves no reference to oneself in the content of the intention, as I am the only agent on whose behalf I could ever be intending. That the intention must be carried out by the agent is part of intending, so it need not be part of the content of intention. Nonetheless, it is a requirement on the content of intention that an agent does not intend what they do not believe possible for them to do, and this includes only intending to do something themselves.

²¹ I will discuss in §2.1 why joint intentional actions are not counterexamples to the claims made in this paragraph

²² Bratman (2009, p148) uses this term for the condition that one can intend only one's own actions. He also considers weaker alternatives, the ‘control condition’ and the ‘settle condition’, for contexts in which one can in fact induce someone else to act such that one can intend someone else’s actions more directly. None of the cases I will consider in later chapters involve the giving of orders or similar directing of others, and in cases where you cannot control or determine someone else’s action, they will amount to the same thing.

1.4 Practical reason

There is a final important concept for the explanation of intentional action. Although it does not feature in any of the conditions on some event's being an intentional action, it is generally understood as necessary for an agent's intentionally acting. Intentional action requires practical reasoning. As with the preceding sections, I will not be arguing in favour of a particular picture of practical reasoning, but instead will be identifying fairly loosely the phenomenon to which this term refers, and which various approaches will spell out differently. Having done this, I will turn to detailing two broad approaches to the role that the capacity for practical reasoning — that is, *practical rationality* or *practical reason* — plays.

As I noted in §1.1, not all action is intentional. This is true not only of the action of things like molecules and cotton, but also of sentient organisms. My cats are agents, but their actions are (arguably) not intentional. This is because they are not acting with an intention, but also, at least in their case, because they cannot act with an intention. They may well be acting in ways that are purposive and goal-directed — they sit outside the kitchen door with the goal of getting more food — but they are not acting with intentions. To do this requires a capacity that they lack, namely practical reason, although how this is spelled out would differ depending on the picture of intention at work. For instance, practical reason allows agents to form the all-out judgements in favour of a particular course of action which are intentions on Davidson's picture. The disposition to engage in means-end reasoning, that is, practical reason, is part of the functional role that Bratman attributes to planning states. We can contrast practical reasoning with theoretical reason, which is the capacity to reason about what is the case. Anscombe describes ordinary reasoning (that is, theoretical reasoning) as 'reasoning towards the truth of a proposition, which is supposedly shewn to be true by the premises.' ([1957] 2000, §33, p58). When I reason theoretically, I draw conclusions from the things I know and believe about the world, and those conclusions are things I come to know or believe. Deductive, inductive and abductive reasoning are all kinds of theoretical reasoning. That I am calling this reasoning 'theoretical' does not mean that the subject matter must be abstract, but rather that what it issues in are beliefs. My cats lack the capacity for theoretical reason, but this is not why they are incapable of intentional action. Practical reason is the capacity to reason about what to do (rather than what to believe). Practical reason enables me to weight preferences, and settle on the thing that, all things considered, I want to do. It involves reflectiveness on reasons, thinking about how my interests and aims interact, and planning how to bring about those ends. I can think about how I will

do that which I want to do — I do not merely identify a goal and hope for the best. Once I have worked out what I want to do, I can work out how to do it. Practical reason enables me to do this, to reason from my intention to ϕ to the value of my ϕ -ing in order to enable me to realize this intention.

All of this was perhaps a little vague, but that is in part because different models of intentional action also involve different models of practical reasoning. In particular, they disagree on the nature of the outcome of practical reasoning, which is bound up with different models of the nature of practical reasoning more generally. I will thus say a little about how some philosophers understand practical reasoning, so as to illuminate what is meant by this term.

Davidson gives the following picture of intentional action:

When a person acts with an intention, the following seems to be a true, if rough and incomplete, description of what goes on: he sets a positive value on some state of affairs (an end, or the performance by himself of an action satisfying certain conditions); he believes (or knows or perceives) that an action, of a kind open to him to perform, will promote or produce or realize the valued state of affairs; and so he acts (that is, he acts because of his value or desire and his belief).

(Davidson 2001e, p31)

This gives us a model of practical reasoning as the assessment and comparison of various competing considerations, weighing up various desires and looking at relevant beliefs. This goes with a picture of practical reasoning as akin to theoretical reasoning, taking beliefs and desires rather than propositions as its premises, and using these to arrive at a judgement about what one wants to do. Practical reasoning, on this picture, issues in intentions, which are judgements in favour of a course of action, which cause that course of action.

Bratman, as I have discussed, understands the kind of agency that involves intentional actions as planning agency. Unsurprisingly, then, his picture of practical reasoning is bound up with this notion of planning. He tells us:

Practical reasoning... has two levels: prior intentions and plans pose problems and provide a filter on options that are potential solutions to those problems; desire-belief

reasons enter as considerations to be weighed in deliberating between relevant and admissible options.

(Bratman 1987, p35)

This is a picture of practical reasoning as solving the problem of what to do, with the ‘personal level psychological phenomena’ (PPPs) (O’Brien 2017, p274) that make up plans acting as constraints on the reasoning. Plans (and thus intentions) put particularly strong constraints on practical reasoning, since these are things to which one is already committed. Bratman’s picture is of practical reasoning as means-end reasoning — it is reasoning about the means one should adopt to achieve one’s ends. These means might be necessary, that is, the only way of realizing one’s end, or they might be one among several options; but in either case ‘one settles on an option so as to avoid a threatened incoherence in one’s plans’ (Bratman 1987, p35). Practical reasoning issues in a choice about how to act, which is bound up with plans such that we should understand it as a mental state.

Anscombe’s picture of practical reasoning resists the tendency to think of it as akin to theoretical reasoning that appears to be present in both of the foregoing approaches. She understands practical reasoning as issuing in practical knowledge, which is knowledge of what it is one is doing. Drawing on Aquinas ([1485] 1920, Ia IIae, Q3, preprint 5, obj. 1.) she writes:

[p]ractical knowledge is ‘the cause of what it understands’, unlike ‘speculative’ knowledge, which ‘is received from the objects known’. This means more than that practical knowledge is observed to be a necessary condition of the production of various results; or that an idea of doing such-and-such in ways us such a condition. It means that without it what happens does not come under the description —execution of intentions...

(Anscombe [1957] 2000, §48, pp87-8)

As Lavin (2015, p614) notes, the reference to causation in the above passage might seem at odds with the picture of Anscombe as rejecting causal analyses of action that I gave in §1.2.1. However, that oddness relies on the idea that she is offering an alternative suggestion as to what connection intention and action, with ‘causation’ as one of several options to fulfil this role. We should instead understand her as rejecting the idea of a separation between the two. Hinshelwood (2018) argues that practical thought is identical with intentional action and McDowell suggests ‘an interpretation of the phrase “intention in action” according to which the intentions so characterizable are not distinct from the

actions in which they are said to be' (2015, p145). Practical reasoning issues in practical knowledge, which is knowledge of what one is doing, where what one is doing is not something we can cleave apart from what one intends.²³ It is knowledge of the intention with which one is acting and knowledge of what one is doing in acting. When Anscombe, drawing on Aquinas, tells us that practical knowledge is the cause of what it understands, she is contrasting it with knowledge that is caused by what it understands — practical knowledge brings about the things known, because it is knowledge of what one is doing and thus what one is making happen.

Regardless of which of these pictures of practical reasoning we align ourselves with, we can understand practical reasoning as reasoning about what to do, rather than what to believe. Intentional action requires practical reason, since it requires intention. As I noted earlier in this section, the exact role played by practical reason will differ between approaches. Nonetheless, across all of these approaches, practical reason is necessary for deliberating about how to act in the ways involved in intention. Before I close this chapter, I want to say a little about the difference that practical reason makes, and thus the difference between the agency of those agents that can and cannot engage in intentional action. These approaches mark different understandings of the kind of agents that humans are, and this discussion will resurface in the final chapters of the thesis, as I consider the agency of young children.

1.4.1 Transformative accounts

In the previous section, I discussed the role that practical reason plays in intentional action. One way of describing humans as having practical reason to is to describe them as having 'rational animal agency.' This is a perhaps slightly unusual phrase, but it is not an unmotivated one — humans are animals, they are agents, and they have practical reason, the kind of rationality relevant to agency. This phrasing is also valuable as it allows us to pull apart ambiguities in how we understand the role of practical reason in human agency. There are two ways of understanding the statement that humans have rational animal agency. We can either understand 'rational' as modifying 'animal' or as modifying 'animal agency'. That is, the claim is either that we have '(rational animal) agency' or 'rational (animal agency)'. The first of these is what Boyle presents as the view that "rational"

²³ This of course raises questions about how one can intend to ϕ without in fact ϕ -ing. This is outside the scope of this project, but answering this will be helped by remembering the distinction I drew between mere attempts and failures to act in §1.2.2. There's nothing puzzling in the case of a mere attempt to ϕ , since in these cases it's right to say that the agent is ϕ -ing. In failures to act, we might be inclined to deny that the agent really intends to ϕ , if they do not commit to ϕ -ing.

designates a characteristic that *differentiates* the genus “animal” (2012, p399). The rationality that we speak of in rational animal agency is the rationality of the animal whose agency it is, rather than a rationality of the animal agency. This might seem at first glance a mere linguistic distinction, but it is more than this. To take ‘rational’ to be descriptive of the kind of animal we are is to adopt what Boyle (2016) terms a ‘transformative’ approach to rationality — rationality is something that runs through the very nature of the kind of animals that we are, it is part of the essence of the kind of which we are members. On this approach:

“‘Rational’ counts as a differentiating predicate of “animal”, rather than merely as the name of a trait that certain animals exhibit, in virtue of the fact that what is rational differs in its way of being an animal from what is not.’

(Boyle 2012, p409)

Accepting this entails accepting that all human agency is rational agency, not in virtue of any particular feature of the agency or actions of that particular human, but because of the kind of being that they are. We do not need to identify exercises of practical reason to understand them as rational agents. In following this kind of approach, Rödl tells us that ‘[t]he legs of a human newborn are rational legs, personal legs, we may say, even as he moves them randomly during his first weeks of life.’ (2016, p91). The movements of the legs are movements of a rational being, despite the extremely limited abilities of a newborn child, because the child is human. Although the kicking does not rely on reasoning, it is the action of a rational being. We need not assess the actions of a human agent to determine whether any of them are intentional. That someone is capable of intentional action follows from their being an intentional agent, rather than their being an intentional agent following from their engaging in intentional action. In the next section, alongside setting out the alternative, additive, model, I will further clarify how the transformative theorist understands the actions of intentional agents that do not involve the use of practical reason.

1.4.2 Additive accounts

The alternative to a transformative approach to rationality is an ‘additive’ one (Boyle 2016). On this understanding, ‘rational’ does not mark out a kind of animal, but a kind of animal agency. Rational animal agency is not the agency of rational animals, but rational agency of animals. Some non-human animals are rightly described as having beliefs and desires, and as acting on the basis of these. On this, the transformative and additive approaches agree. Where they differ is how they understand

the ‘register’ (Boyle 2012, p419) in which we use these terms — which things they are being used to pick out on each occasion of their use. When we speak of rational action as action, and also speak of the action of a cat as action, we use the term ‘action’ in a different register in the two cases. On this, it seems that additive and transformative theorists will agree. The cat’s action is not rational, and thus both additive and transformative theorists are using ‘action’ with one register in describing this as an action and with another in describing rational action as an action. Where they will disagree, however, is whether we use the term in the same register when we speak of actions that we might otherwise think of as non-rational actions of rational agents (such as the kicking of the newborns legs.) Similarly, they will differ on whether we should understand the use of terms like ‘desire’ and ‘belief’ as having the same register when they are applied to the mental states of rational and non rational agents. Boyle describes the position of additive theorists as involving the assumption that ‘a psychological or epistemic concept which applies to both rational and to non rational animals must be susceptible of a single, undifferentiating account that covers both sorts of application’ (Boyle 2012, p417). He draws from this what he calls the univocal assumption.

Univocal assumption: concepts such as beliefs and desires ‘must be treated as univocal in their application to rational and [non-rational]²⁴ animals.’ (ibid)

Although the description of this view as ‘additive’ originates in the work of transformative theorists, it’s an approach that is familiar. For instance, Burge states:

Children and higher non-human animals do not have *reasons* for their perceptual beliefs. They lack concepts like *reliable*, *normal condition*, *perceptual state*, *individuation*, *defeating condition*, that are necessary for having such reasons. Yet they have perceptual beliefs.

(Burge 2003, p528)

The thought expressed by the additive theorist here is twofold — children and non-human animals lack rational capacities, and thus perceptual beliefs must not be rational. This, then, implicitly relies on the univocal assumption, taking it that however perceptual beliefs function for animals lacking rational capacities, so too must they function for those that possess such capacities. On the additive approach, the beliefs and desires of rational agents are not rational, if they are also beliefs of a kind

²⁴ Boyle uses ‘nonrational’

that a non-rational agent could have. The actions of rational agents are not rational if they are also actions in which a non-rational agent could engage. On this picture, intentional agents are intentional in virtue of engaging in intentional actions. No action is intentional solely in virtue of the kind of agent that carries it out, it is a feature of the action itself. Nonetheless, most actions of intentional agents will be intentional, because it is rare for an agent with the capacity for intentional action to do things in just the same way as a non-intentional agent. For instance, only if you describe things very loosely are my cats and I doing the same thing when we eat our dinner. My eating of my dinner is the outcome of a deliberative process, likely preceded by shopping for ingredients and cooking the meal, and throughout the eating I may well be engaged in conversation, and using my rational capacities in determining which elements to eat together, for instance. My cats eat their dinner by hunkering over their bowls to gobble down the food that we have given them. On this picture, my eating my dinner is an intentional action, not because I am an intentional agent, but because it does involve the use of rational capacities. On an additive approach, being an intentional agent is something that follows from my engaging in intentional actions, rather than my engaging in intentional actions following from my nature as intentional.

In this chapter, I discussed individual intentional action. I began in §1.1 by distinguishing intentional agency from other kinds of agency. In §1.2, I presented conditions on an event's being an action in general, and an intentional action in particular. I then clarified how I was understanding several concepts that are closely related to intentional action. In §1.2.1, I discussed bodily movements, concluding that *mere bodily movements are those movements that are not up to the agent whose body is moving*. In §1.2.2, I distinguished intentionally ϕ -ing from two ways an agent might not ϕ – by failing to act, and by merely attempting to act. *A failure to act is when an agent's not ϕ -ing is in virtue of their not doing something that they need to in order to ϕ* , whereas *mere attempts are attempts to ϕ that do not culminate in the agent's ϕ -ing*. Mere attempts are cases of an agent not ϕ -ing, where the fault does not lie with the agent. In §1.3, I considered the role of intention in intentional action, considering various accounts of the nature of intention. I clarified how I will be understanding the content of intention throughout the thesis – treating intentions as act-directed – and what it is for an agent to act with an intention. In §1.3.1, I set out conditions on the content of intention, including the role that the intended means plays in the content of intention, and the requirement that an agent intends only their own action. Finally, in §1.4, I discussed the notion of practical reason. After setting out what practical reason is, I presented two approaches to the difference that practical reason makes to agency, considering a transformative approach in §1.4.1, and an additive approach in §1.4.2. This chapter has served to clear the ground

with regards to individual intentional action, clarifying how I will be using central notions throughout the thesis.

Chapter 2 – Joint intentional action

In this chapter, I will consider the other kind of intentional action that is central to the literature in the philosophy of action — joint intentional action. Across this chapter, I will argue that the following criteria are compatible with the prominent approaches to joint intentional action:

Joint Intentional Action

An event X is a joint intentional action of agents A_1 - A_n if

- (J1) X is an event of A_1 - A_n 's ϕ -ing
- (J2) A_1 - A_n 's ϕ -ing consists of each of A_1 - A_n acting intentionally such that A_i is ϕ_i -ing
- (J3) A_i is ϕ_i -ing because their ϕ_i -ing contributes to realizing their intention to jointly ϕ
- (J4) A_1 - A_n are ϕ -ing to realize their intention to ϕ

As with individual intentional action, I am not offering a detailed analysis of joint intentional action, but rather a model of it that is compatible with the approaches I survey, and can be filled out in more detail on any of them. In §2.1, I will introduce the phenomenon of joint intentional action, distinguishing it from other kinds of joint activity. In §2.2, I will discuss the role of intentions in joint intentional action. I will then survey four approaches to joint intentional action – plural intention (§2.2.1), participatory intention (§2.2.2), conditional intention (§2.2.3) and Anscombean joint action (§2.2.4). I will draw on these approaches in §2.3 to give a general set of conditions on an event's being a joint intentional action.

2.1 Joint intentional action

In §1.3, I introduced the principle I am calling *Realization*, which is as follows:

Realization

An event X realizes an agent A 's intention to ψ only if X is A 's ψ -ing

An agent cannot form intentions that they believe themselves unable to realize. In §1.3.1, I discussed the own-action condition — that an agent can only intend to do something themselves, they cannot intend something be done by someone else. Intending to ϕ is intending to ϕ *myself*, not intending ϕ -ing to occur or that some other agent ϕ s. Despite this, we often engage in joint intentional actions, which seem to involve intending ϕ -ing that is not one's own. Is joint intentional action a counterexample to the own-action condition? Joint actions are those in which two or more agents act *together*, and joint intentional actions are those in which that acting together involves acting with an intention, or acting to realize an intention. Some joint actions are spontaneous (we might find ourselves walking together without either of us having intended to) and others are based on prior intentions (we might carefully coordinate plans to find a time to play a board game together). For the sake of simplicity, I will focus on joint actions of two agents, but there is nothing significantly structurally different when more agents are involved, although of course there is a greater complexity of coordination since there are more 'moving pieces'. I will focus on the intention with which a joint action occurs, which may or may not be a prior intention.

What is the problem posed by joint action for the claim that one can only intend one's own action? If you and I intend to bake a cake together, then there is a sense in which I intend (that) you bake a cake. Realizing my intention cannot but involve you baking a cake with me. If this were an individual action, it would not be possible for me to intend that you bake a cake, although I might have a lot of pro-attitudes directed towards your potential cake-baking, or intend to induce you to bake a cake. In the context of joint action, though, I intend to do something that requires you to do something with me, in a way that involves intending that you do something. This intending of your actions does not go via intending some way of making you do something, at least once it is agreed that we will act together. That is, I might induce you to agree to the joint action, but in so agreeing you set an intention. At that point, then, we each hold an intention that can be realized only by the other doing something alongside us. The problem is how such an intention is possible. The answer, which I will explore in the following sections, is that I only intend you bake a cake as a part of our intending we bake a cake. We have a shared intention to bake a cake together. My intending of your action is my intending *our* action, as part of our intending our action.

In the following sections, I will survey the predominant approaches to joint action. As with individual intentional action, I will present a generalized description of joint action, intended to be compatible with all the approaches I discuss. There are two, related, reasons for considering joint action in the

context of this project. I am putting forward the claim that dependent intentional action is a recognizable and distinct kind of action. In particular, I am arguing that one of the ways of realizing an intention to ϕ is dependently ϕ -ing. A potential hurdle to this is that intending to ϕ dependently involves intending that someone else support my ϕ -ing. By considering joint action, I will show that this kind of claim has the potential to be true, that there are contexts in which one can intend something that involves the action of another, without this going via inducing them to act. One of the ways of realizing an intention to ϕ is by jointly ϕ -ing (for some values of ϕ). Firstly, then, the consideration of joint action helps make plausible that dependent intentional action is a legitimate category of action. Although this similarity is valuable to motivating my claim, the second purpose of this discussion is to make clear that dependent intentional action is not a species of joint action. It is a kind of *joint activity*, as Smith (2011) uses the term, but this is a broader category. Joint activity includes joint action, but it also includes the cases in which agents jointly do something without engaging in a joint action. A joint action of ϕ -ing is the action of two or more agents who are acting together in order to ϕ . They intentionally jointly act, it is not merely a consequence of their individual intentional actions when considered jointly that a cake is baked. Smith draws out this distinction with the following cases:

Suppose that Löwenheim and Skolem collaborated on a proof of the [Löwenheim-Skolem] theorem... Now suppose they worked in isolation, ignorant of each other's existence, that Löwenheim proved a lemma of the theorem, and that Skolem deduced the theorem from it... In the first case, some action of Löwenheim and Skolem's exists to serve as their instance of proving the theorem, but in the second case, nothing that we would ordinarily call an action of theirs exists to play this role.

(2011, p217)

When they collaborate, there is a joint action. They share a goal, and they act intentionally in coordination with the other to achieve that goal. They can engage in joint activity without jointly acting, such as when their actions happen to coordinate to bring about something that they may not have been aiming at individually or jointly. I will draw on the characterization that I develop here in Chapter 5 to show that dependent intentional action is distinct from joint intentional action.

2.2 Intentions and joint action

As I have indicated, a central puzzle in understanding joint intentional action is making sense of the intentions with which agents engage in joint action. ‘Shared intention’ is the name that Bratman (1993) gives to the kind of intention that we have when we intend a joint action. It being shared does not imply that one single intention is held by all the agents, but it does indicate that there is a sense in which all the agents of a joint action have the same intention. Analysing shared intention can be understood as trying to identify the ‘match condition’ for the intentions of the individual agents — that is, it is about identifying in what sense the intentions of the agents can be described as matching. To capture what is intended in a joint action I will speak of the agents as all acting with the intention *to ϕ together* — in a case of a joint action of two agents, both agents would intend to ϕ together. This ‘together’ serves as a modification of the ϕ -ing. Just as we might intend to bake a cake carefully, and a slapdash baking cannot realize this intention, we can intend to bake a cake together and an individual cake baking cannot realize this intention. In speaking in this way I am not committing myself to a particular approach to shared intention, but instead pointing to the thing that such approaches aim to explain — that at some level the agents intend the same thing, which is to act together. Even if the agents do not have intentions with identical content we can understand these intentions as symmetrical — everyone intends something that mirrors what the others intend, even if there are differences between the specific intentions that each has. These symmetrical intentions will have the same structure, even if their content differs — if I intend to ϕ with you, and you intend to ϕ with me, our intentions have the same structure. This symmetry and structural similarity is the sense in which they are the same.

The claims I make about intending to ϕ together are compatible with what I take to be the main ways of analysing shared intentions. What I will suggest is that on all of these approaches we can think of the intentions of both agents as the same. Both agents of the action intend to ϕ together, and the accounts of joint action are accounts of that to which this intention to ϕ together amounts. There are four broad kinds of approach to joint intention that I will consider — *plural intention*, *participatory intentions*, *conditional intentions*, and *Anscombean joint intentions*. All four of these will be captured by the model of joint action as realizing the intention, held by both agents, to ϕ together. In the following sections, I will consider these in turn to show that they fit the model.

2.2.1 Plural intention

‘Plural intention’ is how shared intentions are analysed on non-reductive models, in particular in the work of Gilbert (e.g.1990, 2009). This approach is non-reductive in that it proposes that the intentions involved in joint action cannot be reduced to the intentions of individual agents. This does not preclude other kinds of reduction — these plural intentions might still be understood as analysable into beliefs and desires. What cannot be reduced is the subject of the intention, and thus the intention into further intentions. These intentions have a plural subject – the joint agent – and the sharing of an intention by the individual agents involves a ‘joint commitment *as far as possible to produce, by virtue of the actions of each, a single instance of intending to do that thing*’ (Gilbert 2009, p181). The plural intention is an intention to ϕ , held by all the agents jointly. Its being held by the agents jointly means more than it being the product of the intentions of the individual agents, since it would then be reducible to them. Rather, if there is a plural subject with a joint intention

each of a number of persons (two or more) has, in effect, offered his will to be part of a pool of wills which is dedicated, as one, to that goal. It is common knowledge that, when each has done this in conditions of common knowledge, the pool will have been set up. Thus what is achieved is a binding together of a set of individual wills so as to constitute a single, ‘plural will’ dedicated to a particular goal.

(Gilbert 1990, p7)

This plural subject holds a single intention to ϕ , where ϕ -ing is the joint action. This ‘should not be interpreted to imply that there is anything more “on the ground” in such situations than the people in question and their joint commitment’ (Gilbert 2017, p137) or to suggest a joint consciousness. However, the kind of commitment they have to acting together is more than just an agreement that can be easily shrugged off. Gilbert describes the agents of a plural intention as intending ‘as a body’ to do something

the parties are jointly committed as far as possible *to emulate, by virtue of the actions of each, a single body that intends* to do the thing in question.... the parties are jointly committed to emulate a single body with a certain intention.

(Gilbert 2009, p180)

They do not form a single overarching entity, but rather are committed to acting as if they are one, as a plurality that holds a single intention to ϕ . Although the individual agents will have intentions relating to the sub-actions of the ϕ -ing, the intention to bake a cake is held jointly. All the agents have the same intention alongside these individual intentions, not just in virtue of having the same content, but because there is one intention, jointly held. In individual intention, the intention is ‘sent to oneself’. That is, an agent intends that they themselves act, where this reflexivity is built in without any need to make it explicit. The same is true here, it is just that it is sent to the plural subject rather than the individual. On this approach, the worry about intending the actions of another cannot even get started. The plural subject intends the action of the plural subject, and so intends only their own action. Plural intentions are intentions *to ϕ* held by the plural subject.

2.2.2 Participatory intention

The next two kinds of account are reductive – they take shared intention to be analysable in terms of individual intentions held by the individual agents. The first of these is a ‘participatory intentions’ model, one that analyses shared intention in terms of intentions to *participate* in a joint action. Kutz describes a participatory intention as ‘an intention to do my part of a collective act, where my part is defined as the task I ought to perform if we are to be successful in realizing a shared goal.’ (2000, p10). Bratman’s ‘Shared Cooperative Activity’ account (1992; 1993; 2009) characterizes shared intention in terms of intentions that we do something together and that we coordinate our actions to bring this about. That is, ‘when you and I together intend to J we each intend that we J, and we each intend that we J in accordance with and because of meshing subplans of each of our intentions that we J.’ (2009, p143). What exactly is involved in doing one’s part will depend on the role one is playing in the joint action. Maybe doing my part is weighing and measuring ingredients, and you doing your part is combining ingredients and pouring them into a lined cake tin. Our baking a cake together does not require that we, for instance, each hold one side of the eggs as we crack them, or that we hold the spoon together. What it requires is that we act together to bake a cake, and that can involve us deciding together how to divide up the task. Perhaps we do this following a discussion of the division of labour, but also perhaps just through an exchange of looks, responding to what the other does, or through habit after baking many cakes together. Each of us having the intention to do our own part does not require each of us having an intention that requires us to do exactly the same thing. Since the intention of each agent is an intention to do their own part, there is no worry about their intending someone else’s action. Realizing this intention rests on the other agent realizing an intention, but the

model is more of these two slotting together, in a way that is planned, than of either agent intending the actions of the other.

Bratman formulates shared intentions as intentions *that* rather than intentions *to*, and includes in them explicit reference to the agents involved. Nonetheless, a reasonable approximation of these is that they are intentions to do one's part of a joint cake-baking. This alternative framing captures both the idea that this is something we do together and that we need some kind of coordination. Tuomela's (2005) approach is that joint action involves a 'we-intention', and that each individual agent has a 'slice' of that intention as their intention. He takes it that relevant part of the content of a 'we-intention' in ordinary cases of joint action is "to perform X together" entailing a participation intention for each participant' (ibid, p334), which could be captured as each having the intention to do one's part of a joint ϕ -ing. Every agent would have an identical 'slice' of the intention, which means that they would all have the intention to do their part of the cake-baking. Every agent has an intention with the same structure, and if we can formulate them to remove the explicitly self-referential element of 'do one's part' these intentions would also have the same content. As I argued in §1.3.1, following Madden (2014), all intentions have this kind of self-directed element built in. If my intending to ϕ and your intending to ϕ is each of us intending the same thing (where ϕ -ing is acting individually), then there is no problem here. Participatory intentions are no more self-referential than other intentions, and so there is no special problem to thinking of the agents that hold participatory intentions as holding the same intention. Participatory intentions are intentions *to do one's part of a joint ϕ -ing* held by each individual agent.

2.2.3 Conditional intention

The second kind of reductive account is one that analyses shared intention as *conditional intentions*. Velleman (1997) draws on Gilbert's claim in setting up her account, that the 'plural subject comes into being' (Velleman 1997, p31) once 'each person expresses a special form of *conditional commitment* such that (as is understood) only when *everyone* has done similarly is *anyone* committed' (Gilbert 1990, p7). The constitution of the plural subject rests on each of the individual subjects committing to 'constituting a plural subject' (Gilbert 2014, p67) if the other individual subjects make such a commitment. Only then will the plural subject be formed and the plural intention be held. Velleman expands on this model of conditional commitment, and understands the shared intention to act

together as consisting of two individual intentions held in conjunction, where each is conditional on the other. Agent *A* has an intention expressible as ‘I will if you will’, and agent *B* has the intention expressible as ‘Then I will.’ If the condition of the intentions is met, they will bring about the intended action. The condition is that the other agent also expresses such an intention. In saying ‘then I will’, I am saying ‘Given that you have willed likewise, I will it, too.’ (Velleman 1997, p46). *A*’s intention contains a reference to *B*’s intention, in that it is an intention to act together if *B* is willing, where *B*’s being willing is their having a conditional intention to act together if *A* is willing.

Velleman suggests that if the intention ‘I will if you will’ were ‘fully precise and explicit’ (1997, p46) it would be spoken as ‘I will, if you will and if I recognize you as willing’ (ibid). That is, it is not enough that we each hold the intention to act if the other will act, we need to realize that the other agent holds that intention, else the action will not get started. In saying ‘Then I will’, and thus ‘Given that you have willed likewise, I will it too’, what I am saying is ‘I will, if you will, and I recognize that you will’. That is, I am saying ‘I will if you will’, but with the recognition that you will. That I will if you will is not conditional on your being willing. However, that I am committing to the action, not merely conditionally but actually, is so conditional. We both hold the intention to act, if the other will act. When one of us speaks first, the other updates not their intention, but the information relevant to whether they will be able to act on it. If they do not express the intention as ‘Then I will’, then the action cannot get off the ground. Unless an agent intends to do something they can do ‘just like that’, part of intending to act is working out how to do what it is that one intends. If I purport to intend to go out for dinner with a friend, but make no effort to contact them, think about where we might eat, or how to get there, it is questionable whether I hold that intention at all. If Agent *B* believes that *A* holds the conditional intention, because *A* has expressed it, if *B* did not indicate that they recognized *A* holds that intention then they would not be taking steps to realize their intention. If *B* does not communicate to *A* that they recognize that *A* holds such an intention, the action cannot get going. *Someone* needs to say ‘then I will’, else the two will be trapped in stalemate forever, failing to act. The difference between the two agents is not the intentions that they hold, but the order in which they speak, and thus the information with which they speak. What *B* intends does not change once *A* has spoken, they just know that they are closer to realizing that intention, and part of what they need to do to ensure that they do in fact act together is voice that.

What both agents intend is to act jointly, on the condition that the other agent also intends to act jointly. They both intend to be an agent of a joint action. When I speak of ‘an agent of a joint action’,

I mean one of the two or more agents who, together, jointly act. This does not presume any kind of joint agent, as an ontological entity over and above the individual agents. It also does not attribute to any of the individual agents the joint action. It merely acknowledges that all of the individual agents are involved in the joint action as agents. The action is attributable to them, but not solely to them. There is no weighty claim here, it is merely a way of recognizing that joint actions come about through the acting of multiple agents, and so those agents are all agents of that joint action. Unlike on the participatory intentions model, the intention of the other agent features in the content of each agent's action. The intention of the other that features in the content of an intention is the conditional intention for a joint action – each has an intention for a joint action featuring the other's intention for a joint action. Importantly, though, this is not the actual intention. *A*'s intention does not refer to *B*'s real existent intention to act jointly, since it cannot — *A* does not know of the existence of this intention at the time of speaking. Rather, it is the hypothetical conditional intention, *A* intends to act, *if B holds an intention of the same kind*. *B* intends to act, *if A holds an intention of the same kind*, but when *B* expresses their intention, they know that *A* holds such an intention. *B* knows that the hypothetical intention of the other is actual. Since the intentions contain no particular reference to an actual intention of the other agent, this does not serve to give the two intentions different content, since they do not contain reference to distinct entities (the intentions). They are structurally the same, and they do not have unique distinguishing content. Much as with participatory intentions, we can think of agents holding conditional intentions to act jointly as holding the same intention. Conditional intentions are intentions *to ϕ together, if the other is willing*.

2.2.4 Anscombean approaches to joint action

Alongside these approaches, there have been recent attempts to develop an Anscombean model of joint action, and of the intentions involved therein. Examples of this kind of approach to joint action can be found in the work of Laurence (2011), Rödl, (2015, 2018), Satne (2020), and Schmid (2016). Those advocating for this approach take Anscombe's picture of individual intentional action over into their discussion of joint intentional action.

Laurence, Rödl, Satne and Schmid differ in their versions of this approach, with Satne explicitly setting herself in dialogue with and contrast to the others. However, there is a common core to all of these, such that it is right to think of them as offering the same kind of account, even if the details differ. These approaches extend the Anscombean conception of practical knowledge to cases of joint

action. As I discussed in §1.4, practical knowledge, for Anscombe, is the knowledge that an agent has of what they themselves are doing. One knows what one is doing without observation, and sometimes by arriving at the end of a process of practical reasoning. Practical reasoning does not issue in judgements that one should act in some way, but in action, and this action is what is known in practical knowledge. In the case of individual intentional action, this knowledge is displayed when we ask the agent Anscombe's question 'why?'. If the question is not refused application by the agent, then they have knowledge of that action, and if they tell us that they are ϕ -ing because they are ψ -ing, they demonstrate that knowledge. If their explanation terminates with ψ -ing — with the birthday-cake-bringing that explains the cake-baking that explains the wet-ingredient-combining that explains the egg-cracking, the agent has identified that which guides all of those sub-actions. That they are bringing a birthday cake to the party explains why they are doing all of these things; they are doing them because they are ways of bringing the birthday cake to the party, or some part of so doing. Satne says of the relationship between practical knowledge and intentional action that

[i]n both the solitary and the collective case what defines intentional action is the practical knowledge that agents have of what they are doing.

(Satne 2020, p7)

She then offers the following hallmarks of practical knowledge:

- (a) is not based on evidence or observation,
- (b) assertions expressing it are simultaneously expressions of intentions, and
- (c) its possession is intelligible in the light of the practical reasoning informing the corresponding intention.'

(ibid)

What underlies the various Anscombean approaches to joint action is the idea that this kind of knowledge is present in joint intentional action. It is not solely present in virtue of the individual agents having practical knowledge of what they are individually doing, but as practical knowledge of the joint intentional action.

The details of what it means for agents to have practical knowledge of joint action are spelled out differently on each of these approaches. Laurence does not make the role of practical knowledge as

explicit as the others, at least in terms of using this particular language. His central definition of joint intentional action is that

[a]n individual is acting as part of some group if and only if his action is subject to the special collective action sense of the question “Why?” A group of people are acting together if and only if (1) their individual, first-person-singular actions are all subject to the special collective action sense of the question “Why?” and (2) the same answer holds in each case, consisting of an appeal to an action with a first- person-plural subject.

(Laurence 2011, p289)

If an individual is engaged in a joint action, then their answer to ‘why are you ϕ -ing?’ is ‘because *we* are ϕ -ing’. That is, the agent will offer the explanation ‘I am ϕ -ing because we are ϕ -ing.’ This ‘because’ is of just the same kind as the one that is used in answering ‘why?’ when asked about an individual intentional action. In the case of joint intentional action, the overarching action, the intention that explains all the individual agents’ actions, is a joint one. Satne, although disagreeing with Laurence on some details, makes a similar claim, suggesting that the knowledge individuals have in joint action can be expressed in the following way:

We are doing J (by means of me doing F and you doing G), I am doing F and you are doing G. I am doing F because we are doing J, you are doing G because we are doing J.

(Satne 2020, p21)

This, she tells us, has the hallmarks of practical knowledge that I quoted above, and is an answer to Anscombe’s ‘why?’. Schmid’s (2016) focus is not specifically on the nature of joint intentional action, but rather the nature of this kind of joint practical knowledge. Implicit in arguing for the existence of joint practical knowledge is the idea that joint intentional action has the kind of character that Laurence and Satne are suggesting. Moreover, Schmid’s paper proceeds in part as an attempt to identify what it is ‘that provides the form of unity that is necessary for an aggregate of individual acts to be components of a joint intentional activity’ (2016, p59). This principle of unity is that which means that the actions of the individual agents have the right kind of harmony with each other. In individual intentional action, this unity is provided by the intention with which the agent acts, which is

known in individual practical knowledge. In intentional action, ‘the unity of acts in virtue of which they belong together as parts of one act has a special character: *it is a unity that resides in [knowledge]*²⁵ of *this unity*’ (Rödl 2018, p128). Rödl here is identifying the ‘productive’ character of practical knowledge, that ‘this kind of knowledge produces its own object, it creates what is thus known: an action, a series of actions.’ (Satne 2020, p4). Rödl’s argument is that this applies to joint intentional action as well as individual intentional action. Just as the overarching action and the agent’s practical knowledge thereof provides the principle of unity for individual intentional action, so too do the overarching joint action and the agents’ practical knowledge thereof unify joint intentional action. Since practical knowledge is knowledge of the intention with which one is acting, in joint intentional acting, this practical knowledge is plural. There is, Rödl maintains ‘one act of intending’ (2018, p135) on which joint intentional action rests, just as there is in individual intentional action. In describing what it is that is being done in both an individual and joint tea making, he says ‘[w]e could have spoken of someone making tea; Helen was a variable. If English allowed it, we could have used “someone” in such a way that “Helen and Marc” is a value of it’ (2018, p131). His claim is that we do not ordinarily think of the two actions (individual and joint) of tea making as involving the same intention in the way that we would if we were thinking about two individual tea-makings, in virtue of the lack of a variable that can stand for both individual and plural subjects. This is not a commitment to a joint agent, but rather to agents acting together with ‘one act of intending, an act by those who are acting together’ (2018, p135).

On all of these pictures, joint practical knowledge, knowledge of what is being done held by the agents of the action, is that which underlies the joint action. This practical knowledge is the answer to the ‘why?’ that we might ask an individual agent about their action. In joint action, the answer of the individual agent to the this question is given in terms of the joint action. We are ϕ -ing because we are ϕ -ing. I am χ -ing because we are ϕ -ing. There are significant questions²⁶ to be answered about how we understand this kind of joint practical knowledge, and its status as knowledge, but there is something appealing about this kind of picture. It seems right to say that this is exactly the explanation I would and do offer for my part of a joint action. I am doing what I am doing because of what we are doing. I am doing it as part of what we are doing together, I am doing it because we are acting together. There might be individual reasons I can offer for my choosing to participate in the joint action — I

²⁵ Rödl uses the term ‘consciousness’ here, but earlier in the paper he writes ‘I have reached conviction that the considerations I shall bring forward to explain the consciousness of action contained in action suffice to reveal the relevant consciousness to be knowledge. But I shall not attempt to show this.’ (2018, p124), and so I am comfortable with this substitution.

²⁶ See Roessler (2020) for a discussion of some of these questions.

am meeting you for a walk because I am committed to getting in my 10k steps a day, or because I want to try to new bakery you have told me about that we will get to at the end of our journey — but these are reasons for my action, they are not the intention with which I am acting. When I put on my shoes and grab my keys, if someone asks me why I am doing these things, I will say that it is because *we* (you and I, not my questioner and I) are going for a walk. Whether we should render this as ‘because I am going for a walk with X ’ is a reasonable question, given concerns about the coherence of plural practical knowledge, but what does seem right is that the joint action is the answer to the ‘why?’. On these accounts, the intention with which all of the agents is acting is the intention to ϕ , where ϕ -ing is joint. For some values of ϕ , this will necessitate that the intention is one for a joint action. If ϕ -ing is ‘playing a board game’, for instance, the ϕ -ing will have to be joint. It is not unreasonable to answer the question ‘what are you doing?’ with ‘I am playing a board game’, but implicit in this is that I am doing it with others, because that is the only way of doing this. For other values of ϕ , like going for a walk, there is no guarantee that this intention is one for a joint action. Given that, then, it seems right to render this intention as an intention ‘to jointly ϕ ’, as this indicates that the intention can only be realized by multiple agents acting together — it can only be realized by a joint action. I am ϕ -ing because we are jointly ϕ -ing, the intention with which we are acting is the intention to jointly ϕ . Anscombean joint intentions are intentions *to jointly ϕ* .

2.3 A general model of joint intentional action

I now have four potential models of shared intention. The intention with which the agents are acting in a joint intentional ϕ -ing is an intention *to ϕ together*, and this is how I will refer to this intention going forward. What an intention to ϕ together amounts to will depend on which model of joint action is being used, but they can all be described as intention to ϕ together. If we move back to the cake-baking case, the intention ‘to bake a cake together’ is the intention with which agents jointly bake a cake, which might be analysed in the following ways:

Plural intention: to bake a cake (held by the joint agent)

Participatory intention: to do one’s part of a joint cake baking

Conditional intention: to bake a cake together, if the other is willing.

Anscombean intention: to jointly bake a cake

On any of these approaches, intending to ϕ together involves ineliminable reference to the other agent. The plural intention includes the other agent in the plural subject, the participatory and conditional intentions include the other agent in its content, and the Anscombean intention is what is known in joint practical knowledge. In plural intention, the agents jointly intend the action of the joint agent of which you and the other are parts. This is a first-person plural intention, as is the Anscombean intention. Participatory and conditional intentions are individual and about the intender's action, but realizing them necessarily involves another's acting. Although another features in the intention, they do not feature as intended-*for* – the agent does not intend the intentions of another agent – but as intended-*with*. They feature as a distinct intentional agent, *with whom* the agent wishes to act. If an object featured in my intention, it would feature only as something *on which* I wished to act, even if we speak of acting with it (such as 'beating the eggs with a whisk'.) I can substitute in a different whisk and different eggs and still bake a cake, but I cannot bake a cake with you without you – you are not means to my end but another agent of our shared end. The individual agents each intend a joint action, which involves another agent acting, but they intend only their part and that they coordinate and cooperate with the other agent in appropriate ways. The agents cannot realize their intentions without another acting, and another features in their intentions, but that is not the same as intending the actions of another. All of these analyses involve only first-personal intending and so do not violate the own-action condition.

However we frame joint intentions, they are intentions that (1) are in some sense the same for all agents of the action, at least structurally, and (2) contain ineliminable reference to the other agents as other agents of the action. Central to joint intentional action is that all the agents of the action are acting with the same intention. This 'same' points to token identity on the first and fourth models that I detailed, and type identity on the second and third. There need not be a single act of intending, but the intention being realized by the action is the same intention for all agents — an intention *to ϕ together*. Although the contributions of the various agents differ — 'Sally is chopping the onions and John is washing the lettuce' (Satne 2020, p6) — there is nonetheless a symmetry. They are both directing their agency at realizing the intention to ϕ together, in this case to cook dinner together, and their contributions are both directed at this. Sally might be better with a knife than John, and thus do all of the chopping, but this does not mean there is an asymmetry between the two agents from the perspective of what they are doing, where we understand that as making a salad. Both of them are participating in the joint salad-making as agents, both of them are engaged in making a salad jointly,

and so we can understand what they are doing as symmetrical, even if what one of them is doing is more arduous or complex.

I now want to give a general definition of joint intentional action, drawing on these approaches. As a first attempt, I can simply adapt the definition I called ***Intentional Action*** as follows, considering a group of agents A_1 - A_n :

Joint Intentional Action₁

An event X is a joint intentional action of agents A_1 - A_n if

- (1) X is an event of A_1 - A_n 's ϕ -ing
- (2) A_1 - A_n 's ϕ -ing contributes to the realization of an intention to ψ together with which A_1 - A_n are acting
- (3) A_1 - A_n are ϕ -ing because their ϕ -ing contributes to their ψ -ing together

This definition tells us little about what is involved in A_1 - A_n ϕ -ing — in particular, it does not tell us what each agent is doing. Together, A_1 - A_n are ϕ -ing, but this does not mean that each of A_1 - A_n are ϕ -ing. The issue is not that this definition entails that they are all individually ϕ -ing, but that it tells us nothing at all about what they are individually doing. All the agents are acting with the same intention, the intention to ϕ together, but they realize this intention through each doing their part of the joint action. If some agents are jointly ϕ -ing, it does not entail that they are severally ϕ -ing — the ϕ -ing need not distribute in this way. You and I are jointly baking a cake, but we are are not individually baking a cake²⁷ — I am combining the wet ingredients, you are weighing the flour. In the general case, for A_1 - A_n to be jointly ϕ -ing, A_1 - A_n must each be acting — A_1 is ϕ_1 -ing, A_2 is ϕ_2 -ing, ... A_n is ϕ_n -ing. Some of the agents may be doing the same thing — ϕ_1 - ϕ_n need not be distinct act-types — but they may all be doing something different. Together these actions form a joint ϕ -ing. This is not enough, of course, since the ϕ -ing could be an unintended consequence of the agents actions, a serendipitous outcome of the agents acting in parallel. In joint intentional action, the agents are all acting with the same intention, the intention to jointly ϕ . Given this, I have reformulated the definition as follows:

²⁷ See (T. H. Smith 2011) for a careful discussion of this phenomenon.

Joint Intentional Action₂

An event X is a joint intentional action of agents A_1 - A_n if

- (1) X is an event of A_1 - A_n 's ϕ -ing
- (2) A_1 - A_n 's ϕ -ing consists of each of A_1 - A_n acting intentionally such that A_i is ϕ_i -ing
- (3) A_i is ϕ_i -ing because their ϕ_i -ing contributes to realizing their intention to ϕ together
- (4) The joint ϕ -ing contributes to the realization of an intention to ψ together with which A_1 - A_n are acting
- (5) A_1 - A_n are ϕ -ing because their ϕ -ing contributes to their ψ -ing together

In joint intentional action, the transition in (5) from jointly ϕ -ing to jointly ψ -ing is, I think, an unnecessary complication. In individual intentional action, it was important to keep separate what one is intentionally doing and the intention with which one is acting — I am intentionally ϕ -ing and acting with an intention to ψ . Including this distinction in the definition of joint intentional action is unnecessarily confusing. We may well be engaged in a joint ϕ -ing because we are jointly ψ -ing — we are robbing the bank because we are collecting start-up funds for our new venture — but this does not necessarily work as a transitive explanation. I might in fact be looking over plans for vault locks because we are setting up an artisan cheese shop, but the explanation that is most relevant to this is that we are engaged in a bank heist.²⁸ If I just tell you about our desire to sell Stichelton, you might think I've failed to understand the question. This is perhaps also the case if I tell you that I am cracking eggs because I am taking the cake to Cheryl's birthday party, but there seems to be something different about this. In the individual action case, the egg cracking and wet ingredient combining are both understandable as sub-actions of taking the cake to the party. These are all actions of which I am the agent. In the context of joint action, we should distinguish between the explanation of the individual agent's actions by the joint action, and the explanation of the joint action by a further joint action. It is the former that I am interested in here, the latter will echo the explanation of individual action by individual action. The latter is, in fact, what ***Joint Intentional Action₁*** characterizes, since it is framed in terms of joint actions all the way down. To capture what it is for agents to be acting

²⁸ Moreover, although these overarching intentions are the kinds of explanation that the Anscombean might look for, since they feature in the A-D series of explanations, the other kinds of approach that I have surveyed are likely even less interested in the opening of the cheese shop. In explaining the individual action, these approaches look for what it is that caused this particular action — it may be that the agent is acting in service of some broader intention, but they may not be, and either way this is not the kind of explanation that is being looked for.

jointly, we need to characterize the transition from individual action to joint action. Given this, I will characterize joint intentional action as follows:

Joint Intentional Action

An event X is a joint intentional action of agents A_1 - A_n if

- (J1) X is an event of A_1 - A_n 's ϕ -ing
- (J2) A_1 - A_n 's ϕ -ing consists of each of A_1 - A_n acting intentionally such that A_i is ϕ_i -ing
- (J3) A_i is ϕ_i -ing because their ϕ_i -ing contributes to realizing their intention to jointly ϕ
- (J4) A_1 - A_n are ϕ -ing to realize their intention to ϕ

This explains the individual action of each agent in terms of the joint action in which they are engaged, and recognizes that the agents of a joint intentional action may not all be doing exactly the same thing.

In this chapter, I considered the nature of joint intentional action. In §2.1 I clarified this notion, distinguishing it from other kinds of joint activity. In §2.2, I set out the sense in which all the agents of a joint intentional action are acting with the same intention — they are symmetrical and structurally similar. I then discussed four approaches to joint intentional action — plural intention (§2.2.1), participatory intention (§2.2.2), conditional intention (§2.2.3) and Anscombean joint action (§2.2.4). I closed the chapter in §2.3 by giving a general definition of joint intentional action as acting in which the agents act with an intention *to ϕ together*. This definition is compatible with the definitions surveyed in the previous sections.

Across this and the previous chapter, I have discussed the two main kinds of action that are discussed in the philosophy of action — individual and joint intentional action. In chapter 4, I will turn to a further kind of intentional action that has not been properly described or accounted for in the literature thus far. Before doing this, I will argue for the centrality of being the agent of an action to a key set of questions that we ask in relation to action — those regarding moral responsibility.

Chapter 3 – Moral responsibility and personal identity

In this chapter I will argue that being morally responsible for an action presupposes that one is identical with an agent of that action. I will be using the language of identity, since that is what is used in the literature with which I am engaging, but being identical with an agent of an action just is being an agent of the action. In the previous chapter, I touched on the ways in which an agent can be involved in the action of another agent. It will follow from this chapter, that, if we understand someone as morally responsible for an action, e.g. of harm as an action (rather than the action of tricking or coercing someone into harming another), then they are the agent of that action.

I will begin in §3.1 by setting out why we should believe that moral responsibility for an action presupposes identity with the agent of the action, suggesting that it represents a common and common sense approach. Parfit and his successors have raised doubts about the adequacy of personal identity to capture the relationship that matters in certain practical contexts. In particular, I will turn in §3.2 to Shoemaker's arguments for the existence of moral responsibility in the absence of personal identity. Having set them out, I will first consider Köhler's response in §3.3, before rejecting this and moving to present my own argument against Shoemaker in §3.4, defending the view that we can only be morally responsible for actions if we are identical with their agent. This will be supported in §3.5 by considerations about what is involved in ownership of an action.

I am using 'moral responsibility' in what I take to be a fairly commonplace way, which can be approximated as

Moral responsibility: A person P is morally responsible for E iff it is appropriate to hold certain moral attitudes towards them in relation to E²⁹

²⁹ This points only to extensional equivalence, rather than purporting to offer a meta-ethical analysis of moral responsibility

This understanding pitches moral responsibility in terms of its consequences, rather than its conditions, but that is sufficient for my purposes — identifying when it is that we are holding someone morally responsible enables us to identify the necessary conditions for so doing. The attitudes are those like praise and blame, which are often accompanied by reward and punishment, respectively. I will not assume at the outset that moral responsibility implies causal responsibility, nor will I make any assumptions throughout about how we should analyse any of the associated attitudes. I'll also use 'accountability' throughout to mean this same thing.

3.1 The ordinary view

That moral responsibility presupposes personal identity seems to fall naturally out of two principles that I take to be in the background of our ordinary thinking. The first of these is what Shoemaker dubs 'Platitude', which states:

Platitude: One can be morally responsible only for one's own actions.

(2011, p488) (2012, p109)

I will keep Shoemaker's formulation, but this could also be stated as 'One can be morally responsible for an action only if that action is one's own.'

The second principle is what I will call 'Ownership':

Ownership: What makes an action one's own is that one is its agent

Platitude is the claim that moral responsibility presupposes action ownership, and Ownership is the claim that action ownership presupposes identity with the agent of the action. Since personal identity is the identity relation, restricted to the domain of persons, and assuming that only persons can be morally responsible for actions, the following principle, or what Shoemaker calls 'Slogan', follows straightforwardly from Platitude and Ownership

Slogan: Moral responsibility presupposes personal identity. (2012, p109)

If Platitude is true, and Ownership characterizes what it is for an action to be someone's own, then we have to conclude that moral responsibility presupposes personal identity. Before I move to consider arguments against Slogan, I want to consider two potential objections to Platitude.

Platitude states that 'One can be morally responsible only for one's own actions.' There are (at least) two readings of this claim, depending on the scope of 'only'. This could be the claim that, of all the occurrences, the only ones for which a given individual can be morally responsible are those that are their own actions. Alternatively, it could be the narrower claim that, given the set of all actions, the only ones for which a given individual can be morally responsible are their own.

The first reading of 'Platitude' is clearly false. One can be morally responsible for things other than actions — for instance, if a given state of affairs is a foreseeable consequence of my action, and if I foresaw its occurrence before the time of acting, it is at least possible that I be held morally responsible for that consequence. If you take the last seat on the tube, leaving a pregnant person standing, then you may be morally responsible for their discomfort whilst standing. Your responsibility is not guaranteed — you might not have seen them moving towards the seat, and not noticed them whilst sitting — and you may not be solely responsible — if others sat at the same time. What this kind of example indicates is that it is possible that someone be morally responsible for the consequences of their actions, and thus that their own actions are not the only thing for which they can be responsible. If we turn to the second reading, we have a plausible claim. If we know who is potentially responsible for an action, we may consider whether there are mitigating factors that mean they are not in fact responsible. However, another common kind of situation is one in which we know an action has occurred — we recognize that a state of affairs must be the outcome of someone's action — and we want to work out who is potentially responsible for that action. We do that by working out 'whodunnit' — only once we know that do we know whether anyone can be blamed. When a parent comes into a room and finds crayon all over the walls, they ask 'who did this?', because they need to know who the agent is before they can determine whether someone is responsible and should be punished. "It wasn't me" serves not as an excuse, but so as to remove the possibility that one need offer an excuse. Once it is clear that you are not the agent of an action, there is no way that you could be responsible for it, and so Platitude looks plausible.

Let 'action-responsibility' be the kind of moral responsibility we have for *actions as actions*. This differs from the moral responsibility we have for states of affairs and non-action events (including actions

when they are not being framed or considered as actions). We should then understand Platitude as the statement that ‘One can be action-responsible only for one’s own actions.’ Platitude is true only if no one can be *action-responsible* for actions of which they are not the agent. Cases of coercion or control raise a potential problem here. If I force someone to ϕ against their will, then our instinct may be that I am morally responsible for the ϕ -ing and its consequences (if anyone is). Hyman discusses a case of this kind as follows:

[I]n *R v. Bourne*, a man who had compelled his wife to [ϕ] was convicted of aiding and abetting an illegal sex act, but Lord Goddard stated that if the woman herself had been charged with the offence, she could have successfully pleaded duress. And if it is conceded that the wife in *Bourne* did not [ϕ] voluntarily, presumably the same can be true in other cases of duress as well, for example, cases of perjury, criminal damage, theft, or receiving stolen goods.

(Hyman 2017, p.90)

The wife is the agent of the ϕ -ing, which is an involuntary, but intentional, action. The husband coerced the wife in order to bring about her ϕ -ing. Here we have an action of an agent, for which that agent is not responsible, and a further agent to whom we are inclined to attribute a degree of responsibility.³⁰ We only have a problem if the husband has *action-responsibility* for the ϕ -ing, so we should consider why we hold him responsible at all. The ϕ -ing is a foreseeable and foreseen consequence of the husband’s actions, and the primary motivation of his acting. He engaged in the coercion in order to bring about the ϕ -ing. As Schechtman puts it ‘we are holding the person responsible for actions that result from his prior actions, and so the assignment of moral responsibility for someone else’s actions is always via a more primary ascription of another action to the person who is being held [potentially] responsible’ (Schechtman 2007a, p14, fn15). They are responsible for the other’s actions in the way we are responsible for other consequences of our actions — they are responsible for their occurrence — but they do not have action-responsibility for them. That the kind of responsibility is of a different kind to action-responsibility is clear if we reflect on a parallel case. In this, the husband attempts to coerce the wife to ϕ , but he does so unnecessarily, since she already wanted to ϕ . If he was unaware of her desire to ϕ , we may well hold him responsible in just the same

³⁰ We can imagine other cases of this kind — the negligent parent whose child breaks fragile goods in a shop, or the injuries sustained after in the chaos someone shouts ‘fire’ in a crowded theatre (Schechtman 2007a, p.14 fn.15). In all of these, there are actions (the ϕ -ing, the breaking, the injuring) whose agents are not responsible for them, and there is a distinct agent who we are inclined to hold responsible.

way as when her ϕ -ing was involuntary, but we would hold her responsible too – and here she would have action-responsibility. Only she is responsible for ϕ -ing, he is responsible for causing-a- ϕ -ing. In the case where she is in fact coerced, she nonetheless is the only person who is potentially action-responsible for the ϕ -ing. It does not present a challenge to *Platitude* if we understand it, as we should, as a claim about action-responsibility. There are other kinds of moral responsibility — such as that for the consequences of our actions, or what we might call coercion-responsibility, for coercing someone with the aim of their acting in a particular way — but we can only have action-responsibility for our own actions.

3.2 Owning an action

In the two papers in which he names *Platitude* and *Slogan*, Shoemaker argues against the idea that the former principle entails the latter. He maintains that *Platitude* is true, but that *Slogan* is false, and thus cannot be entailed by it. He argues for this by considering various criteria of personal identity, ‘looking for whatever criterion of personal identity most plausibly fills in the following blank: *what makes an action one's own is* _____.’ (D. W. Shoemaker 2011, p488). He states that if *Platitude* is to entail *Slogan*, our criterion of personal identity must fill the blank, so that personal identity entails moral responsibility. He then argues that no criterion of personal identity can adequately fill this blank.

I will call the uncompleted sentence *Ownership**

Ownership*: *what makes an action one's own is* _____.’

My ‘Ownership’ is the claim what we get when we fill in this gap with ‘that one is its agent’. Note that this does not rely on any particular criteria of personal identity, but it will require that whatever makes someone one and the same person also preserves their being the agent of their actions — this does not seem an unusual demand on such criteria.

Shoemaker proceeds primarily by considering cases in which he believes there is or would be moral responsibility in the absence of personal identity. If these cases hold, and if we accept *Platitude*, then *Ownership** must be filled in with something other than criteria of personal identity. He provides both fantastical and ordinary cases in which he takes it that we would have moral responsibility without

personal identity. The first kind are thought experiments of the sort common in the personal identity literature, and he focusses on two cases that he sees as especially difficult for personal identity criteria to overcome:

Physical Fission: After performing some immoral action, I split down the middle. The two resultant halves re-grow to result in two people who are both psychologically and biologically exactly similar to me. (2012, pp118-119)

Branch Line: After performing some immoral action, I go to sleep, and am secretly duplicated by mad scientists. After a few days, they kill me and cremate my body, secretly inserting the duplicate into my life. (2012, p120)

Both cases have the same basic form – at time t_1 , there is an action of person P for which they are morally responsible. At a later time t_2 , some person or persons are created, such that at the moment of creation they are intrinsically just as P would have been if he had survived in the normal way. In the fission case, neither person at time t_2 can be identical to P — they cannot both be identical to P , because that would require their being identical to each other; and there is no way to choose between them. In the Branch-Line case the non-identity is even clearer, since person P does not go out of existence at the moment of creation of the duplicate — the original continues to exist at time t_2 and thus is identical to P . At a later time t_3 , however, P is dead, and although their duplicate remains alive, they are not identical to P . Shoemaker's claim is that, despite the clear lack of identity, it is right to hold the products of fission and duplication (henceforth collectively referred to as Products) morally responsible for the action that occurred at time t_1 .

He thinks this is plausible both from the third-person perspective on them, and from their own first-personal perspective. Wiggins comments that '[i]n a society where people occasionally divided... a malefactor could scarcely evade responsibility by contriving his own fission' (1976, p138) and this sort of thought would likely have some influence on us here. We may well feel that the fission products were 'getting away with something' if they continued to benefit from P 's immoral action without facing any of its consequences. Moreover, Shoemaker suggests that this would be exactly how they would feel about action X . Of the duplicate produced in the Branch-Line case, he says:

He will come into existence full-blown thinking he is me, and he will (quasi-)remember my actions, will delight in thinking that he's gotten away with the immorality, will carry out my intention to celebrate the immorality, and so forth. Once we imaginatively project ourselves into his mental terrain and appreciate that he is psychologically exactly similar to the *real* me, it becomes extraordinarily difficult to resist the verdict that my actions are properly attributable to him as well, that he is just as eligible for being held accountable as I am.

(D. Shoemaker 2012, pp120)

He thinks that not only do we have the inclination to punish them for X, they will have the inclination that they should be punished, and they will be relieved at avoiding this fate. If we hold them morally responsible, and we accept Platitude, then we need to find a way to fill in the gap in Ownership*. Since no reductionist theory of personal identity will take identity to hold in these cases, we must use something other than criteria of personal identity to do so. These cases suggest that moral responsibility does not presuppose personal identity.

We might not be convinced by these fantastical cases, but Shoemaker also offers ordinary cases in which there is moral responsibility in the absence of personal identity. I will call them *Joint Action* and *Giving Orders* and they go as follows.

Joint Action: A joint action X is carried out by some agents A_1-A_n . I am one of these agents, and thus it is correct to hold me morally responsible for it.³¹ (D. W. Shoemaker 2011, p499)

Giving Orders: A general orders his soldiers to 'take the bridge'. If they are successful, it is right to praise him, if they are unsuccessful, it is right to blame him. The general is morally responsible for the taking of the bridge. (ibid, p500)

In both cases, we hold someone morally responsible for an action despite their non-identity with the agent of the action. In *Joint Action*, we hold several people responsible, but none of these individuals is identical with the agent of the action 'simply because that original agent was a *we*, not an *I*' (ibid,

³¹ Shoemaker frames this as my being one of the 'members of the joint agent' (2011, p499) of the action, but I am avoiding this needless commitment to the existence of such an entity.

p499). Even if we deny that there is an overarching joint agent, it remains clear that I am not *the* agent of this action, since nothing I do is the joint action. In *Giving Orders*, the taking of the bridge is something that happens on the bridge. The general, comfortable in his quarters at the time of the attack, ‘performed *none* of the individual actions contributing to the action’ (ibid, p500). In both of these cases, Shoemaker describes the agent or agents that we hold responsible as owning the action in question. He describes the fantastical cases in the same way — the Products own the action of person *P*. This ownership is what Shoemaker takes it means it is right to hold them responsible, and what ensures the truth of Platitude.

We know that Shoemaker does not take the ownership of an action to depend on personal identity, so he must be filling in the blank in *Ownership** with something else. How is Shoemaker filling this gap? The second kind of argument he offers against Slogan helps here. He argues that the logic of the relations of identity and action ownership leave space for a ‘conceptual wedge’ (D. Shoemaker 2016, p318) between being the agent of an action and being its owner. He states:

...there is no contradiction in supposing that some action *A* belongs to me where I am not identical with the agent of *A*. The source of this non-contradiction is that the two relations—identity and ownership— just have two different sets of relata. The relata of the personal identity relation are persons (or human individuals) at different times; the relata of the ownership relation are persons and actions. It's unsurprising, then, that one might have the latter without the former with no contradiction.

(D. W. Shoemaker 2011, p498)

Two entities standing in one of these relations to one another can tell us nothing about which things they stand in the other relation to (that is, two persons being identical tells us nothing about which actions they own, and a person owning an action tells us nothing about who they are identical to). Action ownership is clearly a relation that can be one-many, since one person owns many actions. Given *Joint Action*, it seems it can also be many-one.

Shoemaker understands action ownership as akin to property ownership. He repeatedly points to this, discussing the way that a house can be jointly owned, and noting that the fact that you can only be taxed on your own property ‘doesn't imply that all property that's mine is mine exclusively.’ (2011, p498) Our starting point should be that property ownership is the model for action ownership as

Shoemaker understands it. That means it is a relation that (a) multiple individuals can bear to the same entity, and (b) that is transferable. In co-authored work, he argues that the ownership we have of the products of our labour stems from the ownership we have of that labour, and the status of that labour as an intentional action ‘that has its ultimate causal source in an agent’s psychology’ (Jaworski and Shoemaker 2018, p12). In the same piece, Jaworski and Shoemaker suggest that attributability – that which makes an action mine – requires of an action its ‘meeting the basic requirements for the agent to be open to moral appraisal for it’ and ‘its expressing some essential feature of the agent... [such that] attributable actions are agential fingerprints on the window of the world, a residue of practical identity’(2018, p5). Alongside its resemblance to property ownership, then, ownership of actions also involves its expressing the psychology and practical identity of the agent.

Shoemaker (2016) considers what he calls responsibility-ownership to be one amongst several kinds of ownership. In this context, he characterizes responsibility-ownership as something that ‘...consists in something like the preservation of the psychological elements contributing to one’s volitional network, but this may be very different from psychological continuity full stop.’³² The other kinds of ownership he discusses help clarify what that preservation could require. When he discusses self-concern, for instance, he suggests that the concern I have for my friends, family, acquaintances and neighbourhood are grounded in ownership (2016, p320). They are ‘mine’ in a way that stems from ownership, and that ownership ‘could be preserved in [a] duplicate despite the loss of identity, and despite the fact that the duplicate didn’t *make* any of those friends or acquaintances, say, or wasn’t around for the genesis of the family or the original self’ (2016, pp320-1). We need to tweak the idea that ownership of an action requires its ‘having its ultimate causal source in an agent’s psychology,’ since in this kind of duplication case, that psychology will not come into existence until after the action has taken place, and so this cannot account of the kind of transferability for which Shoemaker is arguing. Instead, it seems like it is enough that the owner’s psychology and the action have a common cause of the appropriate kind. I will take it, then, that Shoemaker fills out Ownership* as follows:

Ownership: what makes an action one's own is its being appropriately causally related to one’s psychology such that one can (a) be thought of as owning it in a way akin to property and (b) it expresses one’s practical identity

³²This, he suggests in (D. Shoemaker 2016, p318, fn23) is what (D. Shoemaker 2012) fleshes out.

This means that Platitude no longer entails Slogan, since although “[s]ometimes attributability may seem to ride on the coattails of a theory of personal identity... this is just an illusion’ (Jaworski and Shoemaker 2018, p.17). If we look back at the fantastical cases, Ownerships accounts for the relationship that the Products have to the action in question. They have inherited ownership of the action from their predecessor, its agent *P*, and it expresses their practical identity, since it expressed that of the person from who they are produced. The action and their psychology are appropriately causally related, since both are caused by the psychology of the action’s original agent. In *Joint Action*, the individual agent’s psychology caused the action, they own it because they were involved in doing it, and it expresses their practical identity because that is (part of) what brought the action about. *Giving Orders* works in a similar way — the general’s psychology caused the taking of the bridge since he gave the order, he owns it because he brought it about, and it expresses his practical identity since the strategy was his.

If we take Ownerships to characterize the relationship that a person bears to the actions that they own, then we must accept that we have no grounds to accept the truth of Slogan. When we want to know who is responsible for a given action, we should not look to identity with its agent. If this delivers up the right person, it does so by coincidence. Rejecting Slogan is not a move without precedent³³ but as I noted in §3.1, Slogan seems to coincide with our ordinary practices of holding responsible. That is not grounds to reject Shoemaker’s argument, but it is grounds to at least attempt to resist it. In the following sections, I will try to do that. First, I will consider Köhler’s response to Shoemaker, before rejecting it. Doing this will give some insight in to how to proceed, however, and in the final sections I will argue for a return to Slogan, and in favour of Ownership over Ownerships.

3.3 Köhler’s defence of Slogan

In response to Shoemaker, Köhler argues for an alternative approach to personal identity, one that would not fall foul of Shoemaker’s cases. The format of these cases is such that they can in principle be tweaked to cover any reductionist³⁴ criteria of personal identity that we might offer. However, a non-reductionist approach is itself unappealing. As Köhler puts it:

³³ As well as Shoemaker and Parfit, the view that personal identity is not what matters for responsibility is defended by Khoury and Matheson, e.g. (Khoury 2013), (Matheson 2014) (Khoury and Matheson 2018),

³⁴ That is, any account on which facts about personal identity are not *sui generis* but are instead reducible to facts of another kind, for instance those about bodily or psychological continuity

On the one hand, reductionists about personal identity have to deny the plausible connection articulated by [Slogan].³⁵ Hence, reductionists miss something distinctive about PERSONAL IDENTITY. This suggests that PERSONAL IDENTITY is a non-reductionist concept. But, while non-reductionism preserves [Slogan], it faces epistemological (and probably also metaphysical) problems that undermine the legitimacy of the practice invoking personal identity.

(Köhler 2021, p53)

Given this, he relies on neither reductionist nor non-reductionist criteria of personal identity. He suggests we move to *non-representationalism*, which is an approach most familiar in ethics. It involves rejecting *representationalism*, which is the idea that ‘we account for the nature of [a] concept in terms of what it *represents*, i.e. what personal identity *consists in*’ (ibid, p52). Rather than looking to the metaphysics of persons, he suggests we should look for a ‘non-representational function’ (ibid) served by judgements of personal identity. We should characterize personal identity in terms of how it is used, and the practices in which it plays a role. Rather than offering an answer to questions like ‘what is required for two persons to be identical?’, the approach will answer ones like ‘what do statements of personal identity express?’ — the first kind of question, which is about the nature of persons and their persistence, is just not the kind of thing that this sort of theory is concerned with. Answers to the second kind of question will be given in terms of the uses that we can put the statements to, and the way they interact with other things we may say or do. Offering a non-representationalist account of a concept involves the implicit claim that its ‘functional role doesn’t primarily involve representing the world as being a certain way’ (ibid, p54). Judgements of personal identity clearly serve a role in structuring and guiding our beliefs and behaviour, even if they also clearly purport to represent the world. Given that representationalist approaches have failed to protect Slogan, Köhler’s alternative seems worth entertaining if it can avoid this problem.

The practices he looks to are those involved in ‘*making judgements about certain questions about attributability over time*’ (ibid). Here, he says, personal identity plays a central role, suggesting that

PERSONAL IDENTITY functions to allow us to make judgements about which being is the *inheritor* of the attributability of a particular set of actions, mental states,

³⁵ Köhler renames ‘Slogan’ as ‘ENTAILMENT’ and ‘Platitude’ as ‘ACCOUNTABILITY & OWNERSHIP’. I have preserved Shoemaker’s terminology throughout to avoid unnecessary complication.

and mental events, i.e. who to be disposed to hold to account for that set of actions, mental states, and mental events.

(Köhler 2021, p54)

Köhler's analysis of what judgements of personal identity, attributability, and accountability express depends on his use of the notion of 'plans,' as defined by Gibbard (2008). This term is not to be understood in the ordinary way but as a technical term for things that 'structure, control, govern, and organize intentions in distinctive ways' (Köhler 2021, p54). This '*regulating role*' (ibid) is played by a variety of mental states. Judgements of personal identity, attributability, and accountability, are all plans in this sense, and can be characterized as follows:

Accountability_K: judgements of accountability are '*plans to hold some being to account for some particular [X]*'³⁶ (ibid, p55)

Attributability_K: judgements of attributability are '*plans that play a structuring, controlling, governing, and organizing role regarding the plans that constitute judgements about accountability*' (ibid)

Köhler elaborates attributability as follows:

to think that [X] is attributable to some being, *just is* to plan to hold it accountable for [X], depending on one's further judgements about [X] and the being, such as e.g. whether one judges that the being had sufficient control over and knowledge of [X].

(Köhler 2021, p55, emphasis mine)

From this, we can derive:

Personal Identity_K: '*To think of A and B as the same person is to settle that those actions, mental states, and mental events attributable to A are attributable to B... to think of B as the inheritor of the attributability for a particular set of actions, mental states, and mental events...*'

(Köhler 2021, p55)

³⁶ Köhler uses 'φ' here, but since in my notation, 'φ' refers to act-types rather than actions, I have substituted this for 'X'.

Judgements of personal identity are plans disposing one to make particular judgements about attributability — they are judgements about the inheritance of the attributability of a set of actions, mental states and events. Since judgements about attributability are plans that dispose one to hold someone accountable, judgements of personal identity nest these two dispositions, and are plans that dispose one to be disposed to hold someone accountable. When I say ‘*A* is the same person as *B*’, I express my disposition to be disposed to hold *B* accountable for *A*’s actions and mental states and mental events (or vice versa, depending on the temporal location of *A* and *B*). In this way, then, judgements of accountability are the ones that are basic in this context — judgements of attributability are built out of judgements of accountability, and judgements of personal identity are built out of judgements of attributability. This means that judgements of personal identity are, at their core, judgements about moral responsibility, and so the truth of Slogan falls naturally from the definition of these terms. Moral responsibility presupposes personal identity because a judgement about personal identity has a built in judgement about moral responsibility — any judgement of personal identity just is a judgement of moral responsibility. There could be no moral responsibility without personal identity, because from any moral responsibility judgement we could construct a personal identity judgement about the person we are holding responsible. Saying that *A* is identical to the owner of an action just is saying that they are (potentially) accountable for that action.

This does save Slogan, but at too great a cost. I have both internal and external critiques of the account that Köhler develops. I will take these in turn. Firstly, the account fails by its own standards. A non-representationalist account of a concept is one that characterizes it in terms of a non-representational role played by that concept. Although Köhler is right to say that judgements of personal identity play a role in judgements about attributability over time, this is far from their only functional role. Characterizing what judgements of personal identity express solely in terms of judgements about attributability and accountability leaves it unclear what many of the things we say that invoke personal identity could express. Consider the following, uttered when a friend mentions plans to have coffee with a mutual acquaintance later in the week:

“I saw him yesterday”

Underlying this kind of utterance is a personal identity judgement — that the person my friend is seeing is the same person that I saw yesterday. On Köhler’s account, what this expresses is my disposition to be disposed to hold the person that my friend is having coffee with accountable for the

actions, mental states, and mental events, of the person I saw yesterday. Even if this is technically true, this does no justice to what I am expressing when I utter this, or to why I might say it. I might be inviting my conversational partner to ask how our acquaintance was, indicating my fondness for them, or pre-empting a suggestion that I join their planned coffee (in much the same way that you might say “I ate earlier” when offered food). If the function of the personal identity judgement is to express what Köhler proposes, it becomes hard to make sense of why I would say this at all. Whatever I’m trying to convey, it has little to do with attributability and accountability.

Even if we focus solely on explicit personal identity judgements, moral responsibility is not the only practical attitude for which they play a functional role. The following is a non-exhaustive list of attitudes for which personal identity judgements play a role:

- prudential concern (there is particular kind of concern we have only for our own future states)
- compensation (we want the person that is compensated for some sacrifice to be the person that made the sacrifice)
- survival³⁷
- anticipation (of future experiences)
- intra-personal maximization (prioritizing maximizing one’s own eg wellbeing over maximizing overall wellbeing)
- self-conscious emotions (eg pride, embarrassment)
- patterns of third- and first-personal reidentification (such as recognizing oneself in a photo or an acquaintance on meeting them again.)³⁸

If Köhler’s account is of the functional role that judgements of personal identity play, then it should accommodate these kinds of patterns of use. This would require, however, characterizing what each express using accountability and attributability. Imagine a discussion of parallel propositions, *Platitude_A* and *Slogan_A*, as follows:

Platitude_A: One can anticipate only one’s own experiences

Slogan_A: Anticipation presupposes personal identity

³⁷ Items 1-3 taken from (Schechtman 2007b, p164)

³⁸ Items 4-7 taken from (D. W. Shoemaker 2007, pp317-8)

How would we even begin to assess them? A characterization of anticipation of experiences built out of this conception of personal identity would be, at best, gerrymandered — we would have to find a way to understand experiences, which fall outside the domain of the things for which we might be accountable, through the lens of accountability and attributability, and the forward-looking aspect of anticipation would need to be rendered through dispositions that are partly backward-looking.

Even if we can refine Köhler's non-representationalism to resolve these issues, there remain further problems with this approach. These are what I referred to as an external problem — it is not a problem with it as a non-representationalist approach, but with it as an approach at all. Slogan is the claim that 'Moral responsibility presupposes personal identity'. Köhler's non-representationalism renders Slogan true. However, it fails to accommodate our belief in the truth of Slogan. It does not make it true for the right reason. We think Slogan true because we think that personal identity involves (or guarantees) a particular kind of relationship to the action in question. The existence of that relationship in turn means that one is potentially responsible for the action. Non-identity with the agent of an action removes the possibility of accountability for an action. Köhler's non-representationalism flips this on its head. Rather than being potentially accountable because one is identical, one is identical *because* of potential accountability. On this picture, not being potentially responsible for an action removes the possibility of being identical with its agent. It doesn't do this because it shows that the precondition of being its agent isn't met, but because that is what it is not to be its agent. This objection parallels Butler's objection to Locke's account of personal identity on the grounds that 'consciousness of personal identity presupposes, and therefore cannot constitute, personal identity' (2008, p100). The 'presupposes' of Slogan should be read not as meaning 'entails' but rather, as Butler uses this term, to suggest a direction of explanatory priority — personal identity explains memory (that is, consciousness of personal identity) and it explains accountability.

Moreover, this picture does not have built into it any way of preserving the logic of identity. To say that *B* is the same person as *A* is to say that I am disposed to hold *A* accountable for *B*'s actions, and to say that *C* is the same person as *A* is to say that I am disposed to hold *A* accountable for *C*'s actions. The logic of identity entails that if *A* is the same person as *B*, and *B* is the same person as *C*, then *A* is the same person as *C*. However, the disposition to hold someone accountable does not necessarily have this kind of transitive structure. There is nothing about holding *A* accountable for *B*'s actions and holding *A* accountable for *C*'s actions that entails holding *B* accountable for *C*'s actions or vice versa.

This change to the logical structure of personal identity is so revisionary that we should resist it if there is another way to preserve Slogan.

3.4 Defending Slogan

In trying to preserve Slogan, we are not merely interested in its expressing a truth. We are interested in its expressing a particular truth. If Slogan becomes a claim about the dependency of judgements of personal identity on judgements of moral responsibility, it ceases to be the thing whose truth we were interested in preserving. Given that this is the only strong reason that we have for considering Köhler's non-representationalist approach, it does not matter whether it can resolve the internal issues I pointed to — we have no reason to adopt it.

Shoemaker falls foul of a similar problem to Köhler — preserving the letter of a principle at the cost of its spirit. The principle in question is Platitude. Before I turn to this, though, I want to throw some doubt on the detail of Shoemaker's argument. Let us recall that his central approach is to provide a selection of cases which he believes to be compatible with Platitude, but incompatible with Slogan. In each of these, we have agents who are morally responsible for actions that they own, but they are not identical to the agent of the action for which they are responsible. If this description of these cases is correct, then moral responsibility cannot presuppose personal identity. That these descriptions are correct is what I now want to bring into question.

Shoemaker considers these cases as instances where the agents who are action-responsible for an action are not identical with the agent of that action. In both of the 'ordinary' cases, I think that this is straightforwardly false. The first of these cases is *Joint Action*. Shoemaker argues that the individual members of a joint agent share ownership of the joint action, but none of those individual agents is identical with the agent of the action, because that agent is plural. This is true, but it's also not enough to undermine Slogan. Shoemaker interprets Slogan as a claim about individual action — that the individual agent that is action-responsible for an action must be identical with the individual agent of that action. There is no need to interpret it as such. We could formulate it instead as

Slogan': the agents action-responsible for an action must be identical with the agents of that action

In the case of individual action, this just is Slogan. In the case of joint action, this approach is neutral on whether we need to posit an additional joint agent over and above the individual members of that agent, but on both understandings it will be true in the case that Shoemaker presents. Each of the individual responsible agents will also be amongst agents of the action, and if there is an irreducible joint agent, they will also be amongst those agents. Shoemaker describes all of the responsible agents as ‘individual members’ of the joint agent — this seems to amount to their contributing a subagent of the overall joint agent. They will have responsibility for the action as a whole, but they will also be one of the agents of the action as a whole. That is, moral responsibility in this case will presuppose personal identity.

In *Giving Orders*, the general has responsibility for the action, even though he is not its agent. Here we need to pull apart two strands of responsibility. The general has responsibility for his own actions, such as the orders that he gives. This is straightforwardly action-responsibility, and since the general is the one that gives the orders, it presents no problems for Slogan. The second strand of responsibility is the one that is supposed to cause us difficulties. Not only is he responsible for giving orders, he is responsible for the outcome of these orders — for the taking of the bridge by the soldiers. If it goes well, the general is lauded, if it goes badly, he is blamed.³⁹ This kind of structure should be familiar from §3.1. The responsibility that the general has for the taking of the bridge is not action-responsibility. It is what I tentatively called coercion-responsibility. It is responsibility for intentionally causing someone else to act in a particular way in order to bring about a particular end. It is likely also the kind of responsibility we have for the consequences of an action. Because the general is responsible for the consequences of the bridge being taken, understood as consequences of their orders and the actions of others they intentionally brought about, Slogan remains intact here, and it seems likely that those involved in the taking of the bridge would also have (joint) action-responsibility for it.

In neither of the ordinary cases do we have any reason to reject Slogan. What about the fantastical cases? There are two ways we could resist these — by denying that the identity relation is missing, or by denying that the relevant responsibility relation is present. In favour of the first option, it is worth mentioning that, as Köhler notes, these cases ‘don’t merely speak against the biological and psychological views’ ability to capture our intuitions about attributability. They also speak against their

³⁹ It seems worth noting that this is what it is to hold this kind of role — it is to agree that you have this kind of responsibility for the actions that you order others to carry out.

ability to capture our intuitions about personal identity’ (2021, p43). This need not threaten those accounts of personal identity, it could instead threaten those intuitions, but it does indicate that these cases are ones where questions about personal identity are not straightforwardly answered. Nonetheless, I think the second option has a better chance of allowing us to defend Slogan. Remember that the kind of responsibility that is relevant to Slogan and Platitude is action-responsibility — responsibility for actions as actions. Is this what the Products have in the cases in question? We may seem to think so — if we are persuaded by Shoemaker’s cases, we take it that they are responsible in just the same way that the agent would have been. Why do we think this? Primarily, we think it because they take themselves to be so responsible — they have the same attitudes towards the action as its agent would. Shoemaker maintains that the intuition that the actions are attributable to the Duplicate ‘has its most compelling source in a vivid representation of the phenomenology of the duplicate.’ (2012, p120) If we look at what this phenomenology involves, it’s that ‘[h]e will come into existence full-blown thinking he is me, and he will (quasi-)remember my actions...’ (ibid). He holds the attitudes that he does towards the action because he thinks he is the agent of the action. What we know of his phenomenology is that it is indistinguishable from the phenomenology of the agent of the action, which involves in part thinking about the action as something for which he is accountable.

Of this, Köhler says ‘surely murders wouldn’t suddenly become attributable to people if they are brainwashed to have the psychological life of a murderer’ (2021, p.46). We cannot directly infer from someone’s thinking that they are the agent of an action to their being its agent. This is not quite what Shoemaker is doing, however. The Product’s belief that the action is their own is something that motivates moving away from Slogan. Here we should reflect on what it is to refer to something as ‘one’s own.’ This is at the core of Platitude. When I describe something as ‘my own’, I can be using ‘own’ in several ways. Shoemaker’s reading of Platitude, and his construction of *Ownership_s* clearly relies on understanding this in a way that makes it about ownership in the ordinary sense.⁴⁰ This is not, however, the only way that ‘own’ can function after a possessive adjective – it can also be used ‘to emphasize the identity of the subject: for or to oneself.’ (Oxford English Dictionary 2023) This use conveys no implications about ownership, it serves primarily to secure the reference of the possessive adjective — ordinarily, we use it in this way if there are multiple potential referents. It is to make clear that the subject of a sentence is who is being picked out. If I say “Ben was tackling James when he

⁴⁰ That is, he treats it as this kind of use:

‘Used after a possessive adjective, or a noun in the genitive, to emphasize possession or ownership: of or belonging to the specified person or thing; individual, peculiar.’ (Oxford English Dictionary 2023)

broke his leg,” it is not clear whose leg was broken. If I say “Ben was tackling James when he broke his own leg,” it is fairly clear it was Ben’s leg that was broken.⁴¹ I readily concede that this is neither the only nor the primary use of ‘own,’ but that is not what is needed. All I need is that this is the use at work in the sense of ‘my own’ that is at play in Platitude, and in Shoemaker’s thought experiments.

3.5 Ownership and identity

That we are using ‘own’ in the emphatic way when we assent to Platitude seems clear. It may be that an action’s being mine involves an ownership relation, but when I describe it as ‘my own’ what I am doing is emphasizing that it is mine. The primary relation that we bear to actions is that of being their agent. Actions are events of a particular kind, ones that come into existence only through the acting of an agent. The most relevant agent to any action will be the agent of that action, and the most relevant actions to any agent will be those of which they are the agent. This is the relation that the Products believe themselves to bear to the actions of the person from whom they are produced. They think of those actions as their own, in the sense that they are their agent. That belief leads them to believe they are accountable for them. Shoemaker takes their belief that they are so accountable and works in the other direction — he knows that whatever relationship the fission product bears to the actions, it is not agency. Nonetheless, they think themselves accountable for them, and that they are their own. Thus, their being their own cannot be about their being their agent, but instead a relation akin to property ownership — something that can be transferred, in the appropriate circumstances. And *Ownership_S* is supposed to capture that relation. However, *Ownership_S* is not the relation that they take themselves to bear to the action in question, and it does not capture why it is they think themselves accountable. If this is all they have, there is no reason to get the project of explaining their accountability off the ground. We thought they were accountable because they thought they were, and they only thought they were because they thought they were the agent. They were mistaken. They are not action-responsible, and we have no reason to doubt Slogan.

In the ordinary cases, every instance of action-responsibility is also an instance of identity. In the fantastical cases, we have no action-responsibility at all. Shoemaker’s argument is not just that we should understand these cases in a particular way, but that they show that Platitude does not entail Slogan. This rests on his having preserved the truth of Platitude, and this is the last thing on which I

⁴¹ We would say “James was being tackled by Ben when he broke his own leg” if we wanted to convey that James broke his leg at the same time as Ben tackling him, but that it was not caused by the tackle.

want to cast doubt. Platitude, as he renders it, is the claim that one can only be morally responsible for actions that one owns. This is not what we take Platitude to mean, or why we take it to be true. We take it to be true because we understand ‘one's own actions’ to be the actions of which one is the agent. In the kind of cases Shoemaker refers to, the pull to hold people responsible for actions as actions comes from conceptualizing them as the agent. With this in sight, it becomes clearer that whatever else is the case, Platitude does not apply in these cases, and so they pose no threat to Slogan. If we accepted Shoemaker’s version of Platitude, much as with Köhler’s version of Slogan, we would be preserving the truth of a principle at the cost of its spirit, and of what we think it means. We should not be bothered by any lingering feeling that the fission products and duplicates might have a kind of responsibility for the actions in question. I have discussed kinds of moral responsibility that do not rely on being the agent of an action in this chapter, and here we might simply have another kind. We could think of it as inherited-responsibility, a kind of moral responsibility we have in virtue of inheriting the benefits of someone’s actions at the cost of someone else (who may also have inherited that cost). This seems to capture what is at work in these cases — what is left when we strip away their mistaken phenomenology is the feeling that they are reaping the rewards of the agent’s actions, and someone else is suffering because of those same actions. This calls for rectifying, but not because the action is in some way owned by the person who benefits. This isn’t an unfamiliar thought, we see it in discussions of reparations or the return of seized lands and property, it’s simply an unfamiliar way of inheriting. The unfamiliarity of this way of inheriting points to a broader issue, one touched on in the Wiggins quote I referred to in §3.2. Shoemaker is trying to make our practices responsive to circumstances for which they are not designed. If there was a pattern of fission, then we would adjust our approach to moral responsibility accordingly. Trying to make the practices we have now work for that context is itself misguided, but it becomes even more so once we afford it priority over metaphysics. The relationship between an agent and an action that we should take as key is the one that the metaphysics centres, not the one that makes practical attitudes make most sense. We should not be doing what both Shoemaker and Köhler attempt, and prioritizing answering questions about accountability over questions about personal identity or the nature of agency, because those questions of accountability are conventional in a way that those around agency and personal identity are not. Being the agent of an action is not a matter of practice or convention, it’s a matter of metaphysics. With this kind of possibility on the table, it no longer seems that we need to bend Platitude to capture these cases in order to explain our intuitions.

Moral responsibility presupposes personal identity, because one can only be responsible for actions as actions if one is their agent. Shoemaker's argument against this, which I set out in §3.2, rests on a conflation of this kind of responsibility with other kinds, and on a mischaracterization of various cases, as I argued in §3.4. Both his approach and Köhler's attempt to salvage the relationship between personal identity and moral responsibility, which I discussed in §3.3, fail to capture our understanding of the relevant concepts, or why it is we would think the principles at hand to be true. When we think about responsibility for actions, we primarily care about responsibility for them as actions, about action-responsibility. When we think about the relationships we might bear to actions, we primarily care about being their agent. With this in view, Slogan is true, and Shoemaker's argument against it fails.

Chapter 4 – Dependent intentional action

In the first two chapters of this thesis, I discussed the nature of individual intentional action — intentional action of an individual agent realizing an individual intention — and joint intentional action — intentional action of two or more agents realizing a shared intention. I also discussed literature that considers kinds of action that do not fall neatly into either of these categories. In many of these cases, we found someone helping another act, resulting in an action that seems to fall between individual intentional action and joint intentional action. Although some of the literature I discussed in the introduction acknowledges the existence of these actions, it is not concerned with developing an account of it. That will be the aim of this chapter — to characterize this kind of action, which I call dependent intentional action. Dependent intentional action involves the agency of multiple agents, directed at a particular end, but it is distinct from joint action. I will give a general characterization of this kind of action in §4.1, illustrating it with examples in §4.1.1 that will show the familiarity of this phenomenon and applying that characterization to them in §4.1.2. I will then, across the following sections, explore several key elements of the characterization of dependent intentional action I have offered through drawing distinctions between genuine cases of this kind of action and proximate events that are not dependent intentional actions. In §4.2, I will argue for the necessity of the supported agent acting with an intention to ϕ , and of guidance given by the supported agent and followed by the supporting agent. In §4.3, I will argue that the supporting agent of a dependent intentional action must be acting with an intention to help. I will argue first that this support must be agential in §4.3.1 and that it cannot be accidental in §4.3.2. In §4.3.3, I will explore the relationship between intentionally helping and acting with an intention to help, arguing that in cases where an agent is helping a particular agent with a particular action, they cannot merely intentionally help. In §4.4 I will argue, using the characterization I have defended in this chapter, that the kind of dependency present in dependent intentional action is a feature of actions rather than of act-types. In this chapter, then, I will identify and characterize a familiar and widespread phenomenon, and one that accounts should attend to if they want to properly capture the varieties of human intentional action.

4.1 The nature of dependent intentional action

Although ‘dependent intentional action’ is a new term, the phenomenon that it describes is familiar and widespread. There is an extent to which any endeavour to identify a new category is stipulative and an activity of conceptual engineering, and the boundaries I draw around dependent intentional action are partly those that I think are most fruitful for our theoretical purposes, but I also take it that this category picks out something with which we are acquainted in ordinary experience. It is action that involves multiple agents, but, I will argue, it is the action of one of those agents in particular, who I will call the ‘supported agent’. To dependently intentionally act is to be helped, to support dependent intentional action is to help. At a first brush, this action can be described as follows:

Dependent intentional action is action in which one agent (the supporting agent) directs their agency at the ends of another (the supported agent), whose realizing their intention is dependent on the help of the supporting agent.

Although this is illuminating in beginning to identify the kind of action in which I am interested, it is clearly insufficiently precise. My realizing my intentions in individual action often relies on others. This morning, I got up, I brushed my teeth, I made a cup of tea, and I had a shower. All of these are actions during which I was acting with an intention and which involved me as the only agent. I doubt that anyone would deny that these were individual actions. However, other agents were crucial to my ability to carry out all of these actions. Developers made my alarm app, shop staff sold me my toothpaste, farmers grew my tea and plumbers installed my shower. These are just some of the many intentional actions by other agents that produced the circumstances of my intentional actions this morning. The developers had good reason to believe that the app would be used to wake people up in the morning, the shop staff that I would brush my teeth with that toothpaste, the farmers that the tea would be drunk and the plumbers that the shower would be used. However, although these agents are involved in creating what I will call the *infrastructure* necessary for my action (the objects, conditions and circumstances which enable my action to occur in the way that it does) they are not involved in my action. This is true even if their carrying out of their actions is necessary for the particular events which occurred this morning — without those particular app-developers, my alarm would have been different and thus so too would have been the nature of my getting up. What matters is that although they do these things with an awareness that they will help people act, they have no particular people and no particular actions in sight. What they are doing is over before I get started, but what matters about this is not the timing but its not being directed at my actions. I will discuss these ideas in §4.3.3

in more depth, but at this point what I want to flag is that, at most, those that I have mentioned are partially causally responsible for particular features of my action. They causally contributed to the occurrence of my action, but this is not the same as their being the supporting agents of my dependent intentional actions.

I will argue that dependent intentional action should be characterized as follows:

Dependent Intentional Action

An event X is a dependent intentional action of an agent A , supported by an agent B , iff

- (D1) X is an event of A 's ϕ -ing
- (D2) A is acting with an intention to ϕ
- (D3) B is acting with an intention to help A ϕ
- (D4) B is ψ -ing because their ψ ing contributes to the realization of their intention help A ϕ
- (D5) A guides X such that X occurs in keeping with their direction
- (D6) B is attuned to and responsive to A 's guidance and to A 's actions in realizing their intention to ϕ

Condition (D1), stating that a dependent intentional action is the action of the supported agent, will require significant argument. For this reason, I will set my defence of it aside for the time being, returning to it in the following chapter. I will in this chapter make claims that are related to (D1) — for instance, that the supported agent acts intentionally — but the claim that the overarching action is an action of the supported agent will be reserved for Chapter 5. In the current chapter, I will motivate and defend conditions (D2-D6). I will begin by giving examples of dependent intentional action, showing that they meet this characterization. This will serve to both show the existence of this kind of action, and to begin to justify this characterization of such action. Having done that, I will take conditions (D2-D6) in turn, illustrating the role they play in dependent intentional action. I will do this by drawing on cases that are close to my examples, but which do not satisfy these conditions. That is, I will begin by indicating that this characterization is sufficient to capture dependent intentional action, before turning to the claim that it is necessary.

4.1.1 Examples of dependent intentional action

What kinds of events do I have in mind when I speak about dependent intentional action? Here are four examples, some of which may be familiar, and hopefully all of which will strike the reader as plausible occurrences. I will use these examples throughout the following chapters:

Dressing:

Andrea is an elderly woman. She remains fiercely independent, and her health has not yet reached the point at which she feels she needs residential care. Nonetheless, her mobility and manual dexterity have deteriorated over the last few years, and she has visits from carers in the morning and the evening to support her in getting up and going to bed. She has always taken great pride in her appearance, and this morning she has picked out an outfit that includes her favourite blouse. This blouse has fiddly buttons, and so her carer Barbara helps her dress by doing up the buttons when Andrea, struggling with them, asks for her help. With this help, Andrea gets dressed.

Stacking Rings:

Charlie is a child of 3 years of age. He received a new toy from a visiting relative, a set of wooden rings of various sizes, which can be stacked on an upright pole. Although these rings can be stacked in any order, if stacked by decreasing size they form a tower in a rainbow sequence of colours. Charlie has seen the image on the packaging, and is trying to arrange them in this way, but is getting stuck with the last few rings. His parent, Derek, sets these last three rings out in front of Charlie, and asks him which colour he thinks come next. When Charlie makes a mistake, Derek asks him if that was the right choice. With this questioning, and help in laying out the pieces, Charlie completes the tower.

Driving Lesson:

Edith is learning to drive. They have had several lessons with Frankie at this point, and they thus know the basics of moving the car. In this lesson, however, they are attempting a new manoeuvre – reverse parking. When they first attempt this, they are following instructions given by Frankie throughout, but as the lesson goes on, they are

instead asked by Frankie to narrate their attempt themselves. They do so whilst responding to prompts as given and relying on cues from Frankie's body language to determine if they need to adapt anything that they are doing. Frankie has dual controls, meaning that if she needed to, she could stop the car herself, but this turns out never to be needed. With the support and guidance of Frankie, Edith successfully carries out the manoeuvre.

Yoga:

Gary has been practicing yoga for a long time, both on his own and in classes. Today's class, given by Henry, is challenging, but he knows how to do everything that they are guiding the class to do. Nonetheless, as Gary moves into a particularly tricky posture, Henry is moving around the room, and when they reach Gary they make some adjustments to his form, meaning that he can hold the posture with more ease and that he stretches the correct muscle group more fully. With Henry's help, Gary holds the pose for the desired duration.

We would, I believe, describe each of these as actions of individual agents. That these agents are helped in their actions does not undermine their being the agents of said actions. Andrea is dressing herself (with help), Charlie is stacking the rings (with help), Edith is reverse parking (with help), and Gary is holding the yoga pose (with help). If we attend to the act-types, this becomes even clearer. In 'Yoga', we might think of the act-type as 'stretching' or 'holding a position'. It is of course true that we can stretch and/or hold in particular positions things other than our own bodies. What is involved in yoga, however, is not simply that which we do to an elastic band or to two objects we are gluing together, translated to the body. It is a movement of one's body by one's body. The action going on in 'Yoga' is self-directed, it is an agent acting on themselves in order to modify their own body. The act-types involved are reflexive, in that the subject and object of the action are one and the same.

Dressing is slightly more complicated, as there are material objects involved in its successful commission, namely clothes. Nonetheless, it is something that must be done to someone, and we can understand it as primarily modification of someone's body. Dressing more generally is an action of manoeuvring the body such that it is clothed. Although we can and do dress other people, ordinarily, one dresses oneself. What occurs in a dressing of oneself is different to that which occurs in dressing another — the former is a kind of modification of one's own body, whereas in the latter the body of

the person who is being dressed is often more like an obstacle. Parents have to manipulate the limbs of their newborn children in ways that they find initially distressing in order to work them into their clothes, whereas when one dresses oneself, the movements of putting the clothes onto the body are the movements of having clothes put onto the body. The kind of action that is going on in *Dressing* is self-directed, it is acting on oneself with oneself – it is an action of modifying one’s own body.

In both *Dressing* and *Yoga*, we have act-types that involve an agent modifying their own body. Andrea dresses *herself*, Gary stretches *his body*. The final two act types do not fit this model, as they are clearly kinds of acting on a material object. Stacking rings requires the rings that are to be stacked, and driving a car requires the car that is to be driven. We cannot point to the person whose body is modified as a way of identifying the agent of the action in the way that we can look at Gary, holding a yoga pose, as evidence that he is the one who is stretching. What we can point at, however, in both of these cases, is the person whose body is moving the objects (or at least, doing the majority of the moving of the objects). Charlie is placing the rings on top of each other, his parent is simply helping him identify the correct rings. Edith is turning the steering wheel, looking in their rearview mirrors, pressing on each of the pedals and changing gear as appropriate, Frankie is simply helping them to do these things in the correct order. The agents that are helped in these cases are the agents who are acting. I do not want to commit myself to the idea that there is some particular bodily component of an action that identifies the agent of the action in all cases⁴², and I will return to this in discussing what is involved in the supported agent’s ‘guiding’ the action. However, in these cases, it seems uncontroversial to think of these as actions of the supported agent in virtue of their bodily contributions to the actions. In all four examples, then, we have actions of individual agents, acting with help.

Not only are these individual agents acting, they are acting intentionally. Recall the characterization of individual intentional action that I gave in §1.1:

Intentional Action

An event X is an intentional ϕ -ing of an agent A only if

- (I1) X is an action of A ’s ϕ -ing
- (I2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting

⁴² This possibility has been suggested by Glanville Williams with regard to driving, who states that ‘drive’ ‘must be taken to have such a strong bodily connotation that only the actual driver can be the perpetrator of a driving offense.’ (1983, p370)

(I3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

In *Dressing*, Andrea has selected the outfit in advance of dressing — it seems that we could describe her not only as acting with an intention, but also as acting on a prior intention to get dressed into those clothes. In *Stacking Rings*, we can imagine Charlie stating that he is trying to get them in this order, or we can imagine him appearing to follow the image on the package by glancing back and forth between the two, or expressing frustration when he cannot run his hand smoothly along the edge of the tower because some of the rings are out of order.⁴³ Edith booked their driving lesson, clearly keen to learn to drive, and we can imagine Frankie explaining this manoeuvre to them, detailing its value and challenges, and Edith then expressing desire to learn it, and listening to and following Frankie's instructions. Gary booked this yoga class, has a consistent yoga practice, and is attempting the movement to some degree of success when Henry comes to help him. The agents are acting with intentions, and they are doing the things that they are doing because they contribute to the realization of that intention. All of these, then, seem to be cases in which the action was intentional on the part of the agent to whom I am attributing it, but they are also all cases in which the help of another agent was part of the action being successful. These are all dependent intentional actions.

4.1.2 Characterizing dependent intentional actions

Having identified these as dependent intentional actions, I will now show that they satisfy my characterization of such actions. Although this is obviously insufficient to fully defend this as the correct characterization, it should help support its adoption. That these actions satisfy it lends support to the claim that these conditions correctly characterize actions of these kind. This will also begin to illustrate what is meant by each of these conditions, something I will do in more detail in the coming sections.

⁴³ In chapters 6 and 7 I will discuss in detail how and when we can consider young children as intentional agents. Nothing I say in the current chapter, however, hangs on this, and so if the reader is sceptical about my claims about this example, I encourage them to set *Stacking Rings* aside for the time being

As I stated in §4.1, the characterization I offer of dependent intentional actions is as follows:

Dependent Intentional Action

An event X is a dependent intentional action of an agent A , supported by an agent B , iff

- (D1) X is an event of A 's ϕ -in(g)
- (D2) A is acting with an intention to ϕ
- (D3) B is acting with an intention to help A ϕ
- (D4) B is ψ -ing because their ψ ing contributes to the realization of their intention help A ϕ
- (D5) A guides X such that X occurs in keeping with their direction
- (D6) B is attuned to and responsive to A 's guidance and to A 's actions in realizing their intention to ϕ

Let us consider, then, how this applies to each of the cases I have described. I will at this point, to avoid confusion, dispense with the names that I gave in each particular case. Instead, I will shift to the terminology of 'the supported agent' and 'the supporting agent'.

Dressing

- (1) This is an event of the supported agent's dressing herself
- (2) The supported agent is acting with an intention to dress herself
- (3) The supporting agent is acting with an intention to help the supported agent dress herself
- (4) The supporting agent is doing up the buttons on the blouse because this contributes to the realization of her intention to help the supported agent dress herself
- (5) The supported agent chooses the clothes and begins to dress herself, putting the clothes on in her desired way, and asks for help only when she wants it
- (6) The supporting agent responds to the request for help, doing only what is asked of her in doing up the buttons

Stacking rings

- (1) This is an event of the supported agent's stacking the rings
- (2) The supported agent is acting with an intention to stack the rings
- (3) The supporting agent is acting with an intention to help the supported agent stack the rings

- (4) The supporting agent is laying out the rings in front of the supported agent and prompting him to correct his mistakes because this contributes to the realization of their intention to help the agent stack the rings
- (5) The supported agent identifies the outcome he is aiming at in stacking the rings, arranges them in various ways himself, and the event continues because he is still trying to achieve that configuration
- (6) The supporting agent notices when the supported agent is struggling, and acts in ways that are aimed at helping with particular difficulties and correcting particular mistakes

Driving lesson

- (1) This is an event of the supported agent's reverse parking
- (2) The supported agent is acting with an intention to reverse park
- (3) The supporting agent is acting with an intention to help the supported agent reverse park
- (4) The supporting agent prompts the supported agent to do particular things because this contributes to the realization of her intention to help the supported agent reverse park
- (5) The supported agent is in control of the car at all times
- (6) The supporting agent is reacting to the particular things that supported agent is doing, and is ready to engage the dual-controls if needed — this is true, even if she does not actually engage them

Yoga

- (1) This is an event of the supported agent's holding a yoga pose
- (2) The supported agent is acting with an intention to hold a yoga pose
- (3) The supporting agent is acting with an intention to help the supported agent hold a yoga pose
- (4) The supporting agent is adjusting the supported agent's body as this will enable him to better and more easily hold the yoga pose
- (5) The supported agent could refuse or resist the supporting agent's help and ask for changes in the help being given if it is, for instance, causing discomfort
- (6) The supporting agent is using their yoga and teaching experience to assess what the supported agent is currently doing and identify what he needs to do differently in order to better and more easily hold the pose, and would change the help they are giving if he expressed discomfort

Each of these cases, then, satisfies my characterization of dependent intentional action. This lends support to the claim that these conditions capture dependent intentional actions, that is, the sufficiency of my characterization. These are all examples of events that meet the criteria, and their meeting the criteria shows us that they are dependent intentional actions. Although I had previously identified them as such, it is their meeting these criteria that makes them dependent intentional actions — meeting these criteria shows that we have cases of dependent intentional action. Showing that these meeting these criteria makes these dependent intentional actions points to the sufficiency of this characterization. Of course, since I am introducing a new category, it would be strange for my examples not to satisfy it. However, this category is meant to capture something familiar, a kind of action I maintain is widespread and common, and so it is important to confirm that the characterization I provide actually latches on to those familiar cases.

In these sections, I showed that the characterization I have offered can be met by the examples I have given. It is sufficient to capture the actions that I am describing. I now want to turn to the necessity of these conditions. To illustrate this, I will consider instances in which these conditions are *not* satisfied, showing that these fail to be instances of dependent intentional action. This should both motivate the inclusion of these conditions and serve to further explain the role they are playing.

4.2 The role of guidance and responsiveness

The next chapter is focussed on the claim that condition D1 is the correct way to understand the ownership of a dependent intentional action. Therefore, I will begin these sections by considering condition D2. Although they are not sequential, I want to look at this alongside condition D5. These read as follows:

(D2) *A* is acting with an intention to ϕ

(D5) *A* guides *X* such that it occurs in keeping with their direction

The notion of guidance at work here is one I borrow from Frankfurt (1978), and I will return to this in §5.5. When talking about the distinction between action and other kinds of movement, in contrast to causal approaches to action, he states that ‘[w]hat is not merely pertinent but decisive, indeed, is to consider whether or not the movements as they occur are *under the person's guidance*. It is this that

determines whether he is performing an action' (1978, p158). Given this, what we have in the combination of conditions (D2) and (D5) is part of what is involved in X being A 's intentional ϕ -ing.

I want to also consider condition (D6) here, as it is bound up with condition (D5). Condition D6 is as follows:

(D6) B is attuned to and responsive to A 's guidance and to A 's actions in realizing their intention to ϕ

An agent guides an individual intentional action by acting, it unfolds in keeping with their direction or plan because that is part of what it is for them to act. We could describe the agent of an individual intentional action as responsive to their own guidance, but this would be an awkward and artificial way of speaking. The guidance has the responsiveness built into it when we think of just one agent, because both the guidance and the response to the guidance are just what it is to act individually — it is to do something, and to be in control of what one is doing. The coordination in guidance and response is direct.⁴⁴ In dependent intentional action, the guidance and response is not direct, since the event that is the action involves acting not just by the supported agent, but also the supporting agent. For the supporting agent to be guiding is again for the action to occur in keeping with the supported agent's direction, or in keeping with their plan. The nature of the action is down to the supported agent, not the supporting agent. Of course, the occurrence of these actions requires that the supporting agent act, but the way they act and the shape that the action takes comes from the supported agent. For A to be guiding the action successfully, A must be responsive to their guidance — it is only actually guided by them, if their guidance is why it takes the form that it does. The direct coordination that occurs in individual action must happen indirectly and explicitly in dependent intentional action — B must attend to A 's guidance and actions, and use these to determine how they act.

Some of the guidance will come in the form of clear and concrete instructions – “please pass me that bowl” – but not all of it will – I might gesture towards the bowl, or you might recognize that it is what I need next given what I have done thus far. I cannot give an exhaustive list of the kinds of guidance that we might have over our dependent intentional actions, but the examples I have discussed provide

⁴⁴At least, in the ordinary case — see (Schechter 2018) for an exploration of the indirectness of coordination in split-brain patients.

a variety of forms this guidance might take. In §4.1.1, I described several instances of a supported agent guiding an action:

- The supported agent chooses the clothes and begins to dress herself, putting the clothes on in her desired way, and asks for help only when she wants it
- The supported agent identifies the outcome he is aiming at in stacking the rings, arranges them in various ways himself, and the event continues because he is still trying to achieve that configuration
- The supported agent is in control of the car at all times
- The supported agent could refuse or resist the supporting agent's help and ask for changes in the help being given if it is, for instance, causing discomfort

These are of course specific to the examples that I gave, but we can get from them some more general ideas about the guidance that the supported agent might give:

- The form that the action takes is determined by the supported agent
- The end goal of the action is determined by the supported agent — the intention with which the action occurs is the intention of the supported agent
- The supported agent is in control of that through which the action occurs (eg the car)
- The supported agent could unilaterally bring the action to an end

To get a clearer sense of the necessity of both the supported agent's intention and of *successful* guidance, consider the following case, akin to *Dressing*:

Being dressed

Several years have passed and Andrea is now residing in a long term care facility. Alongside her worsening mobility issues, she has developed dementia, which has progressed significantly. She no longer has any particular interest in her appearance, and expresses no preference between the clothes that she owns. Barbara now works at this facility, and she comes into Andrea's room this morning and persuades her to get dressed, something Andrea would not initiate of her own accord. Barbara selects clothes from the wardrobe, which happen to be the same outfit that Andrea picked out all those years earlier. Barbara dresses Andrea, who moves her limbs as needed,

but does not try to do anything herself, and Barbara does up the difficult buttons with no input from Andrea.

The details of this case clearly differ from those of *Dressing*, and these differences mean that this is not a case of dependent intentional action. The important differences all lie in the input that Andrea, the supported agent of *Dressing*, has in the action that occurs. Andrea does not guide the action. What we have here is not a case of Andrea's dependently dressing with Barbara's help, but of Barbara's dressing Andrea.

Firstly, the intention with which the dressing occurs is not Andrea's. She does not do anything that suggests that she herself intends to get dressed, although she does not resist Barbara's suggestion. If she did resist, but Barbara persisted in dressing Andrea nonetheless, succeeding because of Andrea's increased weakness and mobility issues, we would have a case of her being *acted on against her will*.⁴⁵ This would clearly violate conditions (D5) and (D6), since not only would Andrea not be guiding the action, Barbara would be actively ignoring Andrea's attempted guidance. *Being Dressed* does not fall into this category, since there is no resistance or even reluctance from Andrea — this is something she is amenable to, but which she would not herself have set out to do. This action is occurring neither in accordance with or in defiance of her will — her will plays no role at all in its taking place. I am using 'will' loosely here, I am not suggesting that this term denotes an entity, but rather noting that this action is neither in accordance with any intentions held by Andrea, nor in contradiction to them. It may be that Barbara chooses this particular outfit because she remembers it as a favourite of Andrea's at an earlier time, but this does not bring it closer into alignment with any intentions that Andrea may now hold.

We can imagine another scenario, in which Andrea now holds strong preferences that differ dramatically from those that she had earlier in life — she values comfort over style, and chooses clothes that her younger self would have looked at with disgust. A shifting of preferences is normal across a lifetime, but when this manifests during the onset of dementia, the situation is more complex. We have here the sartorial equivalent of a lifelong vegetarian, whose desire to eat meat in the nursing home causes difficulties in determining which wishes to follow.⁴⁶ Jaworska presents us with a case akin to this, considering Mrs Rogoff who in earlier life 'was an introvert, always carefully guarding the way

⁴⁵ And likely of assault

⁴⁶ Lindemann (2014, Ch.5) considers a case-study in which this happens, which she draws from (Banks and Nøhr 2013).

she presented herself to others' (1999 p105) and whose housekeeper has to determine whether 'Mrs. Rogoff [should] get dressed if her family is coming to visit and she insists on wearing pajamas?' (ibid, p106). In cases like this, what we have is a tension between the expressed preferences of the adult with dementia, and the likely more complex set of views held by their younger self. We have a real difficulty in understanding what it would be to treat them as an agent who can act intentionally, and in determining what it is we should do. Both Jaworska and Lindemann (2010, 2014, Ch. 5) offer careful consideration of these situations, and how it is we can respect the agency and personhood of the adult who his now expressing these new preferences. The point of *Being Dressed* is not to present such a difficulty. We do not have someone whose current preferences and agency come into conflict with their earlier self. Rather, we have someone who does not seem to have preferences either way. Barbara does not go against any of Andrea's wishes if she dresses her, she just simply does not accord with any of Andrea's wishes either. In picking out an outfit that Andrea was once especially fond of, she may be helping to 'hold Andrea in personhood', to borrow Lindemann's expression, sustaining her as who she was, since who she is does not come into tension with it in the way that Mrs Rogoff's desire to wear pyjamas does with her earlier care over her appearance.

The action depicted in *Being dressed* does not come into conflict with Andrea's preferences, then, but it also does not realize any of her intentions. We can, moreover, imagine a case where Barbara's dressing of Andrea is in accordance with her intentions, but does not in any way realize them. Andrea could have decided that this is the outfit that she wants to wear today. Perhaps she is having a more lucid day and is feeling nostalgic for her younger years, perhaps she does not recognize it as an old favourite but can see that it is particularly nice and has been told that she is expecting visitors today and wants to look her best, or perhaps she has simply taken a fancy to it today, liking the feel of the fabric or the detailing. All of these could offer explanations for someone's getting dressed in a particular outfit. If Andrea forms this intention, and Barbara comes into the room and picks out this outfit, without consulting Andrea's preference (perhaps because she rarely has one), and dresses her in it, the dressing does not realize Andrea's intention. Her intention might match with the action that occurs, but it has no role in its occurrence. The action does not occur with the intention that Andrea formed to wear that outfit. Even if the potential supported agent has an intention that the act-type could realize, this is insufficient for the action to be a dependent intentional action. It is not that the supported agent needs to announce the intention in advance but the supporting agent needs to be aware of this intention, and the supported agent needs to be guiding the action as a way of realizing their intention.

It is also important that Andrea guides the action: it takes the shape that it does because that is the shape she gives to it. In *Dressing*, we see her ask for help with the buttons. She might not do this verbally, she might gesture to Barbara. It might even be that they have such a familiar relationship that Barbara recognizes when Andrea has started to give up on doing up the buttons, and offers to help. Perhaps she was primed for this when she saw the blouse that had been selected, knowing that Andrea often struggles with these buttons. These situations are consistent with the action being guided by Andrea, because they are acting in ways directed at helping her to get dressed, and because Barbara only steps in with Andrea's encouragement. If she started on the buttons without any sign-off from Andrea, explicit or implicit, we would find ourselves back in a situation of Andrea's being acted on. Perhaps Barbara wants the process to be over more quickly, or doesn't like watching Andrea struggle — in these situations, it is not Andrea's intention that is the one with which the action occurs. If Andrea decides she wants to give the buttons another try, Barbara needs to be responsive to that, she cannot simply overrule this desire without changing the nature of what is happening. Without this responsiveness to the guidance of A, the guidance is meaningless, and B ceases to support and instead acts on or for A, rather than helping them.

The intention of the supported agent needs to be *why the action is happening*. How much input the supported agent must have for this to be the case is something I will explore further in response to potential objections, but this difference between *Dressing* and *Being Dressed* should hopefully be fairly clear. In the former, Andrea's dressing in this outfit is something that happens because of Andrea's intention to do so, in *Being Dressed* and the variations I have discussed, there is no intention of Andrea's that serves such a role. There must be an intention of Andrea's with which the dressing occurs for it to count as a dependent intentional action. Moreover, the intention needs to be more than the instigating factor, it needs to also be part of the action as it goes on. The supported agent needs to guide the action in service of the realization of this intention, and the supporting agent needs to actively respond to that guidance.

4.3 The necessity of intentional helping

Conditions D3 and D4 are closely related. They state, where A is the supported agent and B is the supporting agent of a dependent intentional action, that

(D3) *B* is acting with an intention to help *A* ϕ

(D4) B is ϕ -ing because their ϕ -ing contributes to the realization of their intention help A ϕ

If we look at the criteria of individual intentional action that I offered in §1.2, we can see a close resemblance. There, I stated:

An event X is an intentional action of an agent A if

- (I1) X is an action of A 's ϕ -ing
- (I2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting
- (I3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

This does not perfectly correspond to conditions (D3) and (D4) of my characterization of dependent intentional action, since it includes a further condition that is absent here — an explicit statement that B 's ϕ -ing contributes to the realization of their intention. However, we can take it that the 'because' of D4 involves the implicit assumption that it does in fact so contribute. If it did not, then the condition would not be satisfied. What these two conditions amount to, then, is that B is engaged in an individual intentional action of the type '*helping- A - ϕ* '.

To motivate this requirement on B 's action, I will begin by thinking about the role that it plays in *Yoga*. In this case, the supporting agent B sets out to teach a yoga class. They are, throughout the duration of the class, engaged in intentionally leading the group. They guide the attendees, instructing them as to which positions will come next and demonstrating them. At times, whilst those positions are being held, they move around the room, helping people in following the practice, and in holding the positions. This is the way in which they help the supported agent A — they adjust his position slightly, supporting him in doing what is needed to complete the class. During this time, they are still engaged in an intentional teaching of the class, but they are also engaged in a helping of A . This helping is intentional, and it is done *to further A 's ends*. Although it also furthers B 's end, in forming part of the teaching of the class, B engages in this action in order to help A in particular, and it is something that is not needed for their realizing their intention to teach the class. This is an act of helping A , rather than of leading the class. A could participate in the class without this help, it is not necessary to B successfully leading the class, but it is needed in ensuring that A can get the most out of this particular element of the class. In this way, then, B is directing their agency at A 's ends. When they are helping A

in the way described in *Yoga*, they are acting with an intention to help A realize their intention to hold that pose, and they are adjusting A's posture in order to help A realize that intention.

For these conditions to not be met is for A's action not to be supported by an agent that intends to help them in this action. To show the necessity of A's action being so supported in characterizing dependent intentional action, I will consider cases in which A's action is aided in some way, but not by the support of an agent that intends to help them with this action. By increasing the amount of help that A receives in a series of examples, I will build up to the requirement that is given by conditions D3 and D4 — that B is acting with an intention to help A ϕ . This will show that without this full requirement being met, we do not have dependent intentional action. I will set aside cases where A engages in entirely unsupported action, since there is no possibility of these being cases of dependent intentional action. First, I will consider cases in which A is supported in a way that that does not involve another's intentional agency. Then I will look at cases in which A's action is enabled by the intentional action of B, but as an unintended and unforeseen consequence of B's action. Then, I will consider cases in which B acts, aware that their action helps A, but in which this help is not B's aim — they are not doing this because of an intention to help A, or because it contributes to the realization of such an intention. Although these examples tend towards dependent intentional action, they will all fail to be such action, in virtue of not meeting conditions D3 and D4. This will illustrate the necessity of these conditions for an event's being a dependent intentional action.

4.3.1 Non-agential support

The first set of cases that I need to consider are those in which the kind of support received is not that offered by an agent. There are various tools that A might use to help his practice — blocks and bands designed to support movements and deepen stretches. When A's action involves these, it is not dependent, any more than it is dependent on the yoga mat he places on the floor, or the floor beneath that mat. These are features of the material world that need to be in place for his successful yoga practice, and the blocks and bands are aides to this practice — they are part of the infrastructure necessary for his action, a term I defined in §4.1 as referring to the objects, conditions and circumstances which enable my action to occur in the way that it does. All sorts of things constitute the infrastructure of our action, some of them more explicitly recognized than others. I cannot bake a cake without an oven, and I cannot easily see things at a distance without my glasses. Although I can function fine without my glasses most of the time, my playing of video-games, for instance, is

worsened by their absence — making out fine details on my television can prove tricky in ways that can slow my reaction time or reduce my precision. They are less central to my action than the console or the controller — I can realize an intention to play *Fallout: New Vegas* without my glasses but not without a console of the right generation to play this game. Nonetheless, the way in which I realize my intention is altered by the presence of my glasses. However, it does not seem right to say that my action is *dependent* on any of these things. If I lack the console I cannot play the game, if my oven is broken I cannot bake a cake, but this doesn't make these dependent actions; it makes them actions that require particular infrastructure if they are to occur. They are act-types that rely on certain aspects of the material world.

I want to pull apart the role played by the console and the role played by my glasses. The former is needed for the action to occur, the latter is an aid to my acting. We often readily accept the idea that things that fall outside of the boundaries of someone's natural body are nonetheless unproblematically part of how they act on and in the world. These are not things that are necessary for a particular action because they form part of the conditions of that action, but rather things that form part of the conditions of the agents acting. This is clear when we think about the aids, supports and prostheses that are often utilized by disabled people, especially, but not exclusively, those with mobility impairments. Francis lists '...eyeglasses, hearing aids, and cochlear implants; 'artificial' limbs; and wheelchairs...' (2009, p205) as among the things that we see as supplementing someone's physical capacities. Moreover, she notes '[n]o one questions whether 'I' am seeing because I wear glasses... [n]o one questions whether 'I' am running if I have a prosthetic foot...' (ibid). In these familiar cases, it is clear to us that these devices help provide the conditions for the enacting of agency — wearing glasses allows someone to gain the information they need to reason in ways that will bring about their desired outcome (or to have a clearer view of their health bar in a video game!) and having a prosthetic foot allows someone to run to the theatre in order to arrive before curtain-up and thus realize their intention to see a play. This is not the same as these things either providing the same kind of support as a supporting agent, or their being part of the conditions of the action itself. Our limbs and eyes are part of how we act on the world, and so too are wheelchairs and glasses. They're part of how the agent acts, they're not part of the particular action in the way that a games console is.⁴⁷

⁴⁷ Although it is of course not rigorous, a Twitter survey (Pennick 2021) of over 400 self-identified wheelchair users found 70.1% of them described themselves as 'walking'. This seems to speak to the idea that these aids provide not a new act-type, but rather a way of engaging in an existing act-type — the wheelchair is how the wheelchair user walks.

As I discussed in §1.3.1, sometimes the means of acting will feature in the content of an intention, when the particular way in which I do something is something I intend. If I want to use up the eggs that might go off, my intention to bake a cake will likely feature these eggs — it might even be that the action is an egg-using-by-cake-baking rather than a cake-baking that uses eggs. But even if I have no particular eggs in mind, my cake-baking will likely involve them. This is not its being dependent on eggs, it is its requiring eggs. I do not need to make explicit that I am going to write with a pen (or a pencil, or a computer), for this is just built into the act-type, any more than I need to make explicit that I'll do it myself. These kinds of objects are not things an agent is dependent on in acting, they're part of what it is to act in that way. To suggest otherwise would turn almost all actions into dependent actions, emptying out this category of any meaning.

We can imagine cases where another agent features in the infrastructure necessary for an intentional action in the way that material objects do, but these cases are not instances of dependent intentional action. If we return to *Yōga*, we can imagine a case in which A treats B as an object to aid in realizing his intention. Much as he can use blocks to support his body, he could use B's body, resting on their limbs in the way he might otherwise rest on an inanimate object. If A does this whilst B is asleep, with no awareness on B's part, then it seems straightforwardly of a kind with using an object in this way. If he does it against B's will, then it seems similar, albeit far less morally ambiguous. If it is done against B's will, then B is treated, at best, as if an object. If A is oblivious to B's (potential) discontent then perhaps he is not doing something deeply cruel, but he is nonetheless failing to recognize B as an agent. If he knows that B does not want to do this, then the cruelty is clear. Either way, B is not playing an agential role in this scenario, and A is not treating them as an intentional agent. B is, of course, still an agent, and that is where the cruelty lies — it is A's failure to respond to and recognize B's agency that makes this so unpleasant. This is far removed from the situation described in *Yōga*, in which B deliberately directs their intentional agency at helping A realize his intention. It is not an action of B's that features here, but B themselves, understood primarily as a material object rather than as a conscious being. What these examples hopefully show is that neither aids to acting nor the objects used in acting provide the kind of support involved in dependent intentional action. These are elements of the infrastructure necessary for the action, not something on which it depends.

4.3.2 Accidental help

In the previous section I suggested that support from material objects is insufficient for dependent intentional action. That the support an agent receives in realizing their intention stems from the actions of another is also insufficient for that action to be dependent. Someone can do something that is helpful in my achieving my ends without even being aware of my existence, let alone the existence of my ends and their furthering of them. The person who holds up traffic by stopping their car to jump out and post a letter holds up the bus. They may do this intentionally, if they see the bus in their rearview mirror, acknowledge that their stopping will impede the progress of the bus, and decide to do it nonetheless. Their doing this could mean that I succeed in catching the bus that I would otherwise miss, and thus the person who holds up traffic intentionally does something that is, in fact, helpful to me. However, they could drive away, thinking that all they had done was inconvenience other people and care not at all about having this kind of impact. Someone acting in a way that is consciously and deliberately selfish could nonetheless be helpful to me by so acting, but since its being helpful is incidental to their realizing their intention, they do not act to help me. The helpfulness is an unintended and unnecessary consequence of what they set out to do, it forms no part of their realizing the intention with which they are acting. As with the role of objects and infrastructure constructed by other people, the unintended consequences of other people's actions are so widespread that to point to these as in some way worthy of special characterization would give a category with no substance. All I am doing is taking advantage of the circumstances created by someone else's actions, even if the connection between my action and theirs is more immediate.

I want to draw a distinction between between doing something that is *helpful* for another agent and *helping* another agent. This is a stipulative point about how I will use each of these terms — the distinction is not alien, but this particular use of language is not necessarily uniformly accepted. Something is helpful with regards to an action if it contributes to the infrastructure surrounding someone's action in a way that aids their acting — it is making the world better suited to the action. To be helpful is to make the world such that someone can take advantage of it in acting. To help someone, however, is to acknowledge the action that someone is aiming to engage in, and to aid with that action in some way, rather than just with the circumstances of the action. One can do something helpful by accident — the person that holds up the bus is helpful, they do something that better suits the world to my ends. One can only help intentionally, just as one can only be kind or considerate or

malicious intentionally. One can say and do things that are beneficial or hurtful accidentally, with no idea that these actions are going to have that effect. Kindness and maliciousness are attitudes towards a person, and being kind and being malicious are terms that pick out when we do things that are beneficial or hurtful intentionally, or at least purposefully. When we are malicious, hurting someone is (part of) the point. In the same way, helping is something one does for a person, it is a term I will use to identify the situations in which something helpful is done purposefully, not merely incidentally. The person that stops the bus does something that is helpful, but they are not helping me. They are not even acting with an awareness that they help anyone. Their action forms part of the infrastructure amongst which I am acting, but they are not helping me. When someone does something that is *merely* helpful (rather than also being a product of their helping) to my acting, I am not acting dependently. Their action's role in the commission of my action is akin to that which might be played by ordinary traffic or a diversion on the route, it is something that leads to the bus being at the stop when I get there, rather than something that supports my catching the bus. Someone's doing something merely helpful whilst acting intentionally does not make them the supporting agent of a dependent intentional action.

4.3.3 Intentionally helping and intending to help

Condition D3 stipulated that the supporting agent, B, is acting with an intention to help A ϕ . The support A receives from B must be agential, and it cannot be merely an accidental consequence of B's intentional action. The reader might object, however, that this does not require that B acts with an intention to help A. B could intentionally help A without acting with an intention to help. In this section, I will argue that, although in most instances, we can distinguish between acting with an intention to ϕ and intentionally ϕ -ing, this distinction does not hold up when ϕ -ing is 'helping A to ϕ '. Intentionally helping a particular agent to ϕ is acting with an intention to help them ϕ , although of course this is not the only intention with which you may be acting in this case. The core claim is that helping someone to ϕ is not something that can contribute to the realization of an intention that is not an intention to help them ϕ . If something is done solely because it furthers an end of your own, it is not helping.

My argument for this will rest on the distinction between being helpful and helping that I laid out in the previous section, but also on two further distinctions drawn by Anscombe (2005, pp141-2). These are the distinction between the general and the particular and between the generic and the specific, as

they apply to the content of someone's ends. Anscombe (ibid) describes an end as particular when 'the end is that something shall hold about a given individual' (ibid, p141), whereas general ends are not about any individual in particular. The use of 'individual' here is not limited to persons – '*This hut* is to be inhabitable' (ibid) is particular, because it is about a given individual, this very hut. Act-types are general, but actions are particular, they are individuals. Anscombe frames the generic and the specific as a matter of form. '[A] man who wants to be wealthy' (ibid, pp142) is aiming at something generic, 'but if he is to achieve this it must take a more specific form' (ibid) — it is not enough to want to be wealthy to bring about that one is wealthy, one has to decide if one wants '[t]he possession of lands, or of a regular income, or of a large sum of money' (ibid) — one has to determine the form of wealth that one is aiming at if one is to actually realize this aim. Specificity clearly comes in degree — wanting to be wealthy is generic, wanting to possess lands is specific, wanting to possess lands in England is more specific yet. Wanting to possess these lands here, pointing at a map or naming them, is not only specific but also particular — it is about some particular entity. Whether an end is general or particular is a question of what it is about, whether an end is generic or specific is a question of the form it takes. Specificity is needed if one is to begin doing something, and particularity is inevitable once one is doing it. Even in bringing about a general end, one has to do it in particular ways. To borrow from Anscombe one last time, 'that there be a copy of the Bible in every hotel room... is a general but specific end' (2005, p142).⁴⁸ Realizing this end, however, will require particular Bibles, particular hotel rooms, and particular people placing the former in the latter.

With these distinctions in mind, consider two cases, both featuring a bus driver. The first case is simply that of the bus driver who is going about their job in what seems likely to be an ordinary way. When the bus driver sets out on their route, they do it knowing that their daily activity will help people get to where they want to go. They may well even like this feature of their job, it may be part of its appeal and something that makes them happy about what they do. However, this helping is general rather than particular. They have a general end, of helping some people. They intend to pick up the people at the stops on their route and they might well have regulars that they often see over the course of their working day and expect to pick them up. However, helping these people in particular, and helping them with what they intend to do, is not part of what is needed to realize the driver's intention. Their intention is to drive their route, to pick up people at stops and to let them out at later stops. Its being realized often involves helping people, but it need not. They still do what they set out to

⁴⁸ Although the Bible is a specific religious text, this intention is not about any particular instances of this text, and so is general.

do if they have a day on which no-one catches their bus. This might be an unusual way of their realizing their intention, or at least a realization of their intention with unusual features (since what they actually do is conceivably minimally different from on other days, since on any given day each stop might be used by no-one), but it realizes it nonetheless. In the ordinary cases where the driver does in fact have passengers, the driver is not trying to get those passengers to where they want to go. Rather, getting them to where they want to go (or some of the way to where they want to go, since bus stops are rarely final destinations) happens in the course of the driver realizing their intention. Even if the driver sees this as something they will do, getting these particular people to these particular destinations is not needed to realize their intention. Getting some people to some places might be a thing they set out to do, but no people in particular and no destinations in particular beyond the stops on the bus route. If one of their regulars was not at the stop, nothing would go wrong in realizing their intention — they would not need to seek out that absent person to be able to realize their intention. If they help people, they are more than merely helpful, they help intentionally. However, they merely intentionally help in the course of realizing their intention, they do not act with an intention to help — they act with an intention to do their job, and helping people is something that ordinarily happens in the course of doing their job. Helping people is for them like wearing down sneakers for Bratman in running, it is something they see as contributing to what they set out to do — and perhaps they think it desirable that it will have this additional effect — but it's not that at which they are aiming.

In contrast to this case, to help clarify the claim I am making, consider the bus driver who helps a particular passenger. This could take a variety of forms. A passenger on an unfamiliar route might ask the driver to let them know when they are at the correct stop. The driver, on arriving at that stop, might call out to that passenger, or even, lean around from their seat to catch the passenger's eye, ensuring that they get off at their destination. One occasionally hears more dramatic varieties of this — a driver who, discovering a lost child on their bus, or a passenger who needs to go to A&E, or someone who has found themselves on the wrong route with an urgent appointment, might diverge from their planned route. Informing the other passengers of their plans, they take the passenger to their destination, or to somewhere far closer to it than the route would otherwise have taken them. When they do this, they are not relating to someone as one of many passengers, but as someone in particular, with a particular end that they are attempting to realize. The helping here is particular and specific, because they are helping someone in particular do something in particular, and the helping is far more specific than merely 'driving the bus.' It is letting the passenger know when they are at their destination so that they can get off, or driving to where it is the passenger wants to go, not merely to

the stops on the route on the assumption that passengers want to get off at those places (else they wouldn't be on the bus).

My claim is that in the kinds of cases I have just described, in which there is a double particularity — both of who is trying to do something, and what they are trying to do — the driver is acting with an intention to help. More than this, they cannot merely intentionally help in this way. As soon as an agent has something with this kind of particularity in sight, they can only help by acting with an intention to help. Helping that lacks this kind of particularity can be something one does either with an intention to help or merely intentionally. The bus driver might have an intention to help people that they do not realize on days when no-one gets their bus. This intention lacks particularity, although its realization will of course be particular. It will be the helping of particular people in particular actions but helping anyone on their bus going anywhere on their route would do to realize this intention. This kind of general helping is something one can do merely intentionally. If people are helped, and the driver foresaw this helping, then, they helped intentionally, because this helping contributed to the realization of another intention — the intention to do their job, for instance. This is not the case with particular helping, however. If one foresees helping a particular person with a particular action, one cannot just do this in the course of doing something else. Here, the difference between being helpful and helping matters. One might foresee that their actions might be helpful to a particular person doing a particular thing, and not seek to avoid that outcome. If one does something that is in fact helpful, and foresaw this, then one is in a slightly different situation to the driver that holds up the bus with no conception of its being helpful. However, one is not intentionally helping this particular agent with their particular action. To help, as I suggested, one has to be doing things in order to help. When helping is general, it does not take as much to do something in order to help, such that one can do it intentionally without acting with an intention to help. It can be done in the course of realizing another intention. Helping with something particular, however, requires directing oneself at bringing that particular thing about. This requires acting with an intention to help, because it is not enough that one does things that might turn out to help someone. One has to actively seek to bring about that which helps.

Once I understand my action as helping some particular person do some particular thing, I am myself invested in its occurrence. I might recognize that my action might bring about circumstances that are helpful for someone to do something, but that is not helping them. I might realize that moving my car might make it easier for my neighbour to park outside their house, but unless I do things to enable

them to do that (knocking on their door to let them know the space will be free, or checking that there is enough room behind my car for theirs), then I am not intentionally helping them, I am merely aware that it might be helpful for them. If, however, I do any of those things, it seems clear that I am not merely intentionally helping them, I am acting with an intention to help them. I am doing things aimed at helping them, rather than just accepting their being helped as a consequence of what I am doing. Helping a particular person with a particular action cannot happen by accident. Something can be accidentally helpful to A's ϕ -ing, but helping A to ϕ requires recognizing what it is that A is up to, and deliberately facilitating this, because they are a particular agent with a particular aim. I understand that aim, and consider what it is I might do to help the realization of that aim, and make sure that I do it. If it is something that I was likely to do anyway, if I do not see this as part of why I am doing it, I am not helping them to ϕ , I am just doing something that is helpful for their ϕ -ing. For me to be helping A to ϕ , I need to be doing things because they help A ϕ , and thus I need to be acting with an intention to help ϕ .

When an agent merely intentionally ψ 's (that is, they intentionally ψ but do not act with an intention to ϕ), they ψ because it contributes to the realization of an intention to ϕ . They do not ψ in order to ϕ , they ψ as part of their ϕ -ing. They condone their own ψ -ing, they realize that they are likely to do it in the course of their ϕ -ing, but they do nothing to guarantee that they ψ . If we think about Bratman's wearing down of his sneaker soles, he is not invested in this occurring, he merely condones its happening in the course of his running the marathon — he does not take steps to avoid it. If he were acting with an intention to wear down his sneaker soles, he would make an effort to ensure that it happen. If this were the guiding intention, the one with which he is acting, he might wear them down manually, taking a pumice stone to them at home. If the marathon running were a way to do this, he might adjust his gait to ensure greater impact on the soles or he might register for a marathon on a hard and rough track. That is, he might alter the character of the running in order to ensure that the soles are worn down. If it is something he does merely intentionally, he is prepared to do it if it contributes to running marathon, but he is not trying to do it. Helping A to ϕ cannot be like this. If I am aware of their intention, and that my action will help its realization, then I have to decide whether I am going to contribute to the realization of *that* intention — am I going to help them to ϕ ? If I do not want them to ϕ , or at least do not want to be part of why they ϕ , I would change how I was acting so as to avoid helping them ϕ . If I want to help them, I would actively ensure I do those elements of my action that would help them to ϕ . If I am indifferent to it, if I do not care either way whether they ϕ , what I do might be helpful, but it is not helping. Once I have identified their intention as something

that I can intentionally help to realize, if I choose to do so, I am not merely condoning the occurrence of the action. I am doing things because they help, I am doing them because they contribute to the realization of their intention. Since the intention they contribute to the realization of is not mine but that of the person I help, I'm not merely helping in the course of realizing an intention — doing it because it so contributes is acting with an intention to help them.

Being the supporting agent of a dependent intentional action requires intentionally helping A to ϕ . This has the kind of double particularity I have discussed. I cannot merely intentionally do it, any intentional helping of this kind, I have argued, is acting with an intention to help. Without such an intention, I am not invested in its coming about, and so I am not helping at all, I am merely doing something that might be helpful. To help is to do something because it is helpful. Dependent intentional action requires acting with an intention to help because it requires the intentional helping of a particular person with a particular action. Without this, an agent is merely helping, or perhaps even merely helpful. The person that is being helped has to feature as a particular agent engaged in a particular action for the requisite kind of helping to happen, and the intention with which the helped person is acting needs to be something of which the supporting agent is aware. Only then can they identify the specific forms their helping must take — they have to know what it is the helped person is doing else they cannot help. Once they know this and they are identifying ways of helping, they are not merely intentionally helping. The intention to help that person with that action is one of the intentions with which they are acting.

To be able to help in just the way that a dependent intentional action involves requires helping that is both specific and particular. My argument in this section is that this kind of helping cannot be something that one merely does. Once the particular action of ϕ -ing — where both the action and the agent are in sight — is one that the agent has grasped, helping with this is something that the agent has to do intentionally. To intentionally help with this particular thing is to direct oneself at helping, because one is recognizing that which the agent is trying to do and how they are trying to do it. One has to actively direct bits of one's action at it, or at least, keep those in place in acting. If I know that moving my car will be helpful, then I might well generate the situation necessary for someone's acting. If I intentionally help someone by moving my car, I need this helping to come off — I need to actually help — and so I need to do things to ensure that they are helped. Only if I do this can I be said to intentionally help them, rather than just to be helpful, and to do this is to act with an intention to help.

It is necessary for dependent intentional action that the supporting agent acts with an intention to help, because this is the only way that one can help a particular agent with a particular action.

4.4 Dependency of actions and of act-types

In this chapter, I have argued that dependent intentional action is a familiar and meaningful category of action, in which one agent directs their agency at the end of another. The supported agent guides the action, whilst the supporting agent, acting with an intention to help them, responds to that guidance. Before I finish this chapter, I want to say something about the act-types involved in dependent intentional action. Some of the actions I have discussed are instances of actions that can also be done individually, such as getting dressed and stacking rings. However, some of them are instances of teaching, and these are among act-types which involve a kind of mutual dependence between the actions of multiple agents. Events of acting in these ways, however, are not necessarily dependent intentional actions — this is a feature of particular actions, rather than of particular act-types.

Consider two cases. First, imagine that B is demonstrating a yoga pose in a class which A is attending. B holds movements at the front of the room as they lead the class through the routine, explaining which things the students might struggle with, and A watches carefully, following B's movements with their own body. Secondly, imagine that B is giving A a lift. B is driving the car, but A is the one determining the destination — B is driving to get A to their destination. In all of these, we have individual actions of individual agents, but they are also things that depend on other agents for their occurrence. The lift is only happening because A wants to get somewhere, the class is only happening because there are students. For an action to be a giving-of-a-lift, there must be someone to whom the lift is given. For an action to be a demonstration, there must be someone to whom it is demonstrated. These are actions that by their nature involve another person. Giving a lift and demonstrating something are both act-types that need to be done to or for someone else. They each have a counterpart, being given a lift and being demonstrated to, which are necessary for them to occur in their ordinary form. We could perhaps have demonstrating without an audience,⁴⁹ but this would be derivative on ordinary demonstration, and would require very particular circumstances to come off rather than just be a failed demonstration. If we think instead about being-given-a-lift and the being-

⁴⁹ For instance, recording a demonstration to be viewed at another time.

demonstrated-to as dependent actions, the same thing is true. These, again, are act-types that need someone else for them to occur.

It is clear that these are act-types that inherently involve a certain kind of dependency. Being given a lift depends on someone else giving you a lift, and giving someone a lift depends on someone being given a lift. Demonstrating something depends on someone being demonstrated to, and being demonstrated to depends on someone demonstrating. That is, these are cases where we have two act-types with the following relationship:

Agent A can ϕ only if agent B ψ s

Agent B can ψ only if agent A ϕ s

Some act-types that fit this characterization are such that there is a symmetry between the names for the act-types (giving a lift, being given a lift), but this is not always the case. To teach someone to ϕ requires someone learning to ϕ . The converse is not always straightforwardly true – you can learn to ϕ without someone else concurrently teaching you to ϕ . There are, however, ways in which anyone learning is nonetheless being taught. It might be that the same person is teaching and learning, as in the case of self-taught skills, or that someone has produced teaching materials from which they are learning. Instances of these act-types require another agent engaging in the corresponding act-type, and so there is a kind of mutual dependency. The question is whether this dependency means that instance of these act-types are always dependent intentional actions.

An initial response here is simply to note that ‘being-given-a-lift’ and ‘being-demonstrated-to’ are not really actions at all, but rather ways of being acted-on. *A fortiori*, then, they are not dependent intentional actions since they are not actions at all. Although this is true, learning is clearly an action, and I still want to reflect a little on the nature of being-given-a-lift and being-demonstrated-to, to clarify how this kind of dependence works. What would it be for all of the cases that have this kind of counterpart-structure to be dependent intentional actions? We would need to identify in each action a supporting agent and a supported agent. Once we attend to these actions, however, it becomes clear that we cannot do this, at least if we understand these actions as a lift-giving and a demonstration. Consider first the demonstration. Demonstrating something to someone does not require letting them guide you in what you demonstrate, and so the demonstrator is not a supporting agent. Similarly, the student is not a supporting agent, because they are not acting in a way that aids the demonstrator in

their demonstrating, except in that they are helping it to be a good demonstration — they are not acting with an intention to help the teacher to demonstrate. The same pattern plays out if we look to the lift-giving. The driver need not be responsive to how the passenger wants to get to where they are going, and they might be giving them a lift because they are heading that way anyway — they are not supporting the passenger. The passenger is necessary for the driver to be giving a lift, but they are not acting with an intention of helping the driver give a lift. The interaction of the act-types does not guarantee that the agents involved are engaged in a dependent intentional action.

There might be instances of these act-types that occur dependently, but this would involve agents supporting the agents involved — they would need to be dependently giving-a-lift or receiving-a-lift, and dependently-demonstrating or following-the-demonstration. The driver's action could be dependent on the passenger, if they were giving the driver directions throughout the lift, or if the driver asked them to “do the gears”. Whether a given teaching or learning is dependent *as a teaching or learning* depends on whether someone is supporting the teaching or learning. For instance, someone with a note-taker in lectures might be dependently learning, in a way that someone whose action is not supported by another is not dependently learning. These same actions might also be dependent instances of other act-types. A receiving-of-a-lift might also be a dependent getting-somewhere, a following-of-a-demonstration might also be a dependent holding-of-a-yoga-pose,⁵⁰ but this will be down to features of the particular actions. The driver who is heading that way anyway is not acting with an intention to help, but the driver who goes out of their way to take someone somewhere they were not otherwise going may well be acting with such an intention. These act-types do not necessitate that instances of them are dependent intentional actions, although they do lend themselves more readily to being instances of dependent intentional action of other act-types. For them to be dependent intentional actions of this act-type would require that an agent involved is doing this particular thing dependently (perhaps teaching through an interpreter, or giving-a-lift with support in driving), rather than simply acting in a way that involves harmony with another action. Dependency of the kind demonstrated in dependent intentional action is a feature of actions, not of act-types.

In this chapter, I have characterized the category of action I call dependent intentional action. In §4.1, I introduced the idea of dependent intentional action, and presented the characterization of this action that I would defend throughout the chapter. In §4.1.1, I gave examples of this kind of action, showing that is a familiar phenomenon, and then in §4.1.2 applied the characterization I am offering

⁵⁰ I will return to the relationship between learning-to- ϕ and dependently ϕ -ing in Chapter 7

to action of this kind. In §4.2, I argued that a dependent intentional action of ϕ -ing requires that the supported agent is acting with an intention to ϕ , and that it occurs with their guidance which is responded to by the supported agent. In §4.3, I argued that dependent intentional action requires that the supporting agent is acting with an intention to help the supporting agent to ϕ . In §4.3.1 I argued that this support must be agential, in §4.3.2 that the help cannot be accidental, and in §4.3.3 that the kind of help being offered cannot happen by an agent acting merely intentionally — they must act with an intention to help. Together, §4.2 and §4.3 defended conditions D2-D6 of my characterization of dependent intentional action. Finally, in §4.4, I argued that the dependency involved in dependent intentional action is a feature of particular actions, rather than act-types.

Chapter 5 – Supported agents as agents

In the previous chapter, I characterized a kind of action that I am calling dependent intentional action. In dependent intentional action, I have suggested, a supported agent realizes an intention with the help of a supporting agent, who directs their agency at the supported agent's end. In that chapter, I set out and motivated five of the six conditions I gave for an event's being a dependent intentional action. In this chapter, I will turn to the remaining condition, which is as follows:

(D1) X is an event of A 's ϕ -ing (where A is the supported agent)

This condition states that dependent intentional actions are actions of the supported agent. This is, I recognize, likely to be the element of my account that causes most disquiet, since these are actions involving multiple agents, and the supporting agent seems more capable with regards to the action than the supported agent. Given this, I will spend this chapter arguing that we can and should understand these actions in this way.

This will be something that I argue for broadly by process of elimination. I will take in turn several alternative explanations of these actions, and show that they fail to accurately describe them. The issue will not be that there are no actions that fit these alternative descriptions, but rather that the actions I have been discussing are not such actions. Considering these alternatives will allow me to map the conceptual space in a way that leaves open only the possibility that we should understand dependent intentional actions as actions of the supported agent. The possibilities are:

- these events are not actions
- these events are the outcome of joint activity (but not joint action) of the supported and supporting agent
- these events are joint actions of the supported and supporting agents
- these events are actions of the supporting agent
- these events are actions of the supported agent

I will consider each of these possibilities in turn, and will ultimately reject all but the last explanation — dependent intentional actions are actions of the supported agent. In §5.1, I will reject the first two possibilities. In §5.2, I will reject the idea that these are joint actions, arguing in §5.2.1 that the supported agent is unlikely to be engaged in a joint action, and in §5.2.2 that the supporting agent is not engaged in a joint action. In §5.3, I will reject the idea that these are straightforward actions of the supporting agent. This will leave me with two possibilities that need to be ruled out — that these are *mere attempts* or *failures to act* by the supported agent combined with actions of the supporting agent, or that there are unacceptable deviant causal chains at work that mean that these are not the actions of the supported agent. I will tackle these in §5.4 and §5.5, ultimately concluding that we can and should understand dependent intentional actions as actions of the supported agent. In §5.4.1 I will set out the idea of an agent acting on another’s behalf; in §5.4.2 I will deny that it is either a mere attempt or a failure to act, and in §5.4.3 I will argue for the permissibility of intending to act with help. In §5.5.1, I will set out the problem of deviant causal chains, showing its potential relevance to dependent intentional action through the notion of heteromesial causal chains. In §5.5.2, I will consider whether purportedly basic dependent intentional actions are in fact basic, in §5.5.3 I will show that the causal chains involved are not deviant, and in §5.5.4 I will argue for the permissibility of heteromesial causal chains when they do not prevent an agent from acting.

5.1 Dependent intentional action is action

In §4.1, I set out the following characterization of dependent intentional action:

Dependent Intentional Action

An event X is a dependent intentional action of an agent A , supported by an agent B , iff

- (D1) X is an event of A ’s ϕ -ing
- (D2) A is acting with an intention to ϕ
- (D3) B is acting with an intention to help A ϕ
- (D4) B is ϕ -ing because their ϕ ing contributes to the realization of their intention help A ϕ
- (D5) A guides X such that X occurs in keeping with their direction
- (D6) B is attuned to and responsive to A ’s guidance and to A ’s actions in realizing their intention to ϕ

I gave several examples of this kind of action in §4.1.1, and I want to introduce one more here, which I will use throughout this chapter to illustrate elements of my argument in defence of condition (D1). Consider a child of, say, six years old. Although they are capable of acting intentionally, there are a great many things they remain unable to do unaided, due to lack of experience and lack of skill. They could, for instance, intend to bake a cake. This intention seems to be just the same intention that an adult could have to bake a cake, in that it is an intention directed at just the same act-type. However, a six-year-old cannot bake a cake alone. Although they can perhaps enthusiastically stir things, they cannot follow a recipe without guidance. They cannot acquire the necessary ingredients or reach the worktops. What realizing the intention involves is radically different from what it involves for an adult. This is partly to do with environmental barriers rather than the child themselves—there is nothing impossible about kitchen surfaces that are the correct height for children. The acquisition of the ingredients might well be thought of as a precursor to cake-baking rather than a component of the cake-baking itself. However, following a recipe and putting a cake into and taking it out of the oven are crucial sub-actions of the cake-baking action, and these are not things that this child can do. These are things that can only be done by a ‘responsible adult’, most likely, a parent or caregiver. I will assume that this is a description of a situation that occurs, even if the reader is inclined to quibble over the exact age of the child involved. The case can then be set out as follows:

Dependent cake baking:

A child intends to bake a cake, and expresses this intention to their parent. The child is unable to bake a cake alone, due to both their lack of skill and their being too young to use the oven unsupervised. With the parent’s help, however, they can and do successfully bake a cake.

The first thing I need to show is that this is an action, rather than an event of another kind. There are two possibilities that are worth considering – bodily movements of a non-agential kind, or joint activity that does not constitute an action.

Consider the event that is dependent cake baking. At one level of description, it can be understood as a series of bodily movements. To get to this description, let us first think about all the things that happen in the making of the cake – reading of a recipe, preheating of an oven, opening and closing of cupboards, weighing, measuring and combining of ingredients, pouring of the mix into the cake tin, opening the oven and placing of the cake tin into the oven, removing of the tin to a cake rack to cool.

All of these involve movements of the bodies of one or both of the agents involved — holding things and moving one’s hand whilst holding them, moving one’s mouth to read the recipe aloud. That we can describe the event in this way does not show that it is not an action, since, as Anscombe (1979) argued, an action need not be intentional under all descriptions to be intentional. In §1.2.1, I made the following claim: *mere bodily movements are those movements that are not up to the agent whose body is moving*. These hand and mouth and arm movements are clearly under the control of the agents involved, the vocal sounds produced by the parent holding the recipe book are under their control. It need not be that they have in focus how they must move their mouth to make the sounds they are making for this to be the case. They do not need to understand exactly how the mechanics of speech work. What matters is that they are reading from the page, and making the noises they make because those are the words on the page, and their doing this is under their control. The hand and arm movements tell a similar story — they are up to the agents. The agents are doing these things, they are not merely happening.

Strictly, all this gives us is that the sub-actions of the cake baking are actions. There is an action of recipe-reading, and action of flour-weighing, and so on. It does not give us that the cake baking is an action. What we do know, though, is not only are the individual sub-actions up to the agents, but they are related. The getting of the flour out the cupboard is not just something that turns out to be useful for the weighing of the flour, it is done in order to enable the flour weighing.⁵¹ All the things the child does are related to one another, and so too are all the things the adult does. This relation is agential, each of them does what they do in service of an end at which they are directing themselves – the child is ϕ -ing, the parent is ψ -ing. The outcome of their ϕ -ing and ψ -ing is that a cake is baked. This does not entail, however, that there is an action of cake-baking. As I noted in §2.1, there is a broad category of joint activity, that is, of ways that the activity of multiple agents can come together. It is possible for the actions of two agents to bring about an outcome at which neither agent is aiming, or for one agent to do something that enables another to later act, without there being an action of the kind in question. For there to not be an action of cake-baking in this situation, one of the following would need to be the case:

- (A) One agent prepares for a cake to be baked, the other finishes baking a cake
- (B) The two agents separately and unrelatedly each do part of baking a cake, such that, taken together, a cake is baked

⁵¹ Arguably, it is part of the flour weighing

In Smith's (2011, p217) case, we imagine that Löwenheim proves a lemma, and Skolem deduces the theorem. There are two ways that we might understand what Skolem does here — we might think that he has done enough to count as proving the theorem, which would leave us in situation (A). This would make this an action still — there is a cake-baking. If it's a cake-baking of the child, then we have no problems here. That is exactly the claim I am defending in this chapter. There might be some resistance to my understanding this as a dependent cake-baking, rather than an ordinary individual cake-baking, but this would be misguided. There are, of course, cake-bakings that happen after someone else has set things out in a way that allows someone else to get started with more ease, but that is not what we have here. This particular cake-baking involves much more dependence on the parent than there would be in that kind of individual cake-baking. On this occasion, this is because of the capacities of the child — they cannot bake a cake without this help — but this does not need to be true for us to see these actions as dependent rather than ordinary individual actions. What matters is not that a cake baking by this agent could not happen in any other way, but that *this* cake-baking could not happen in any other way. This particular cake-baking involves dependency on a supporting agent, and that is what matters. Dependency is a feature of actions. Not only is not a feature of act-types, as I argued in §4.4, it is not a feature of agents. Some agents might be necessarily dependent with regards to particular act-types, but it does not need to be impossible for a given agent to independently ϕ individually for a given instance of their ϕ -ing to be dependent. What makes something a dependent ϕ -ing is the kind of support that is given, and the role that this support plays in this particular ϕ -ing coming off. The specific way in which an agent ϕ s on a given occasion is what makes an action dependent, and so the fact that there could be an ordinary individual action of an agent's ϕ -ing does not rule out some other particular event of their ϕ -ing being dependent. There is no reason to think of dependent intentional actions as ordinary independent individual actions of the supported agent. In §5.3, I will set out why we also cannot understand it as an ordinary individual action of the supporting agent.

The second alternative is that we have an accidental baking of a cake. This is akin to the case in which neither Löwenheim nor Skolem can be said to have proved the theorem themselves. Rather, they each did something that, when combined with that which the other one did, meant that the theorem was proved. It is hard to even conceive of what this would look like in most cases, which indicates the implausibility of this understanding. For something like this to come off, it almost always requires that it is done purposefully. Cakes do not just fall out of things that people do for other reasons. Even if

they did or could, however; this is not what is happening in this case. There is deliberate and careful coordination between the two agents. The parent gets things out the cupboard for the child to use, the child hands the parent the cake tin to place it in the oven. The coordination between their activity is not mere coincidence, it is deliberate on the part of both agents, and is directed at the making of cake. The event that occurs in the kitchen is an action of baking a cake; doubting this would be bizarre. It might be that some dependent intentional actions are of kinds where we can conceive of things of that kind happening through sheer accident and coincidence on other occasions. However, these events are not candidates for dependent intentional action. The kind of deliberate coordination set out in (D3-D6) rules out the possibility that dependent intentional actions are mere coincidences. The baking of a cake is something that both agents set out to have happen, and they do the things that they do in order to bring it about that a cake is baked. These events are clearly actions, not events of some other kind.

5.2 Dependent intentional action is not joint action

The most familiar kind of intentional action involving multiple agents is a joint intentional action. Given this, I need to rule out the possibility that dependent intentional actions are in fact merely ordinary joint actions of the supporting and supported agents acting together. In §2.3, I gave the following as the general form of joint intentional action:

Joint Intentional Action

An event X is a joint intentional action of agents A_1-A_n if

- (J1) X is an event of A_1-A_n 's ϕ -ing
- (J2) A_1-A_n 's ϕ -ing consists of each of A_1-A_n acting intentionally such that A_i is ϕ_i -ing
- (J3) A_i is ϕ_i -ing because their ϕ_i -ing contributes to realizing their intention to ϕ together
- (J4) A_1-A_n are ϕ -ing to realize their intention to ϕ together

To show that dependent intentional actions are not joint actions, I will turn to the content of the intentions involved in each of these kinds of action. In joint intentional action, the agents each intend *to ϕ together* with the other agent(s). For instance, in a joint cake baking, the parent and child would both be acting with an intention to jointly bake a cake.

5.2.1 The supported agent

Consider the dependent intentional cake-baking. As I set things out in §4.1, the intentions involved in a dependent intentional action are an intention *to ϕ* and an intention *to help A to ϕ* . This means that what the child intends is *to bake a cake*. They have the same intention as in an ordinary individual cake-baking. If this were a joint action, the content of the child's intention would need to include a reference to the parent. They would feature not as a means to the child's baking, but as another agent with whom the child will act. If this were a joint action, the child would intend *to bake a cake together* (with the parent). The options for what this involves are that they intend to bake a cake as part of a plural agent (plural intention), they intend to do their part of a joint cake baking (participatory intention), they intend to bake a cake together with their parent, if their parent is willing (conditional intention), or to jointly bake a cake (Anscombean joint intention). The first, second and last of these both involve the child thinking of the cake-baking as a joint action, and this does not mirror the way I have suggested the child thinks about what they are intending. The child intends to bake a cake, and they recognize they need the support of their parent to do this, but that is not the same as intending a joint cake baking. The third, conditional, formulation, might capture the child's intention. The child intends to bake a cake, and in this case they know that they can only realize their intention if the parent is willing to help them. Maybe, we can understand the child as intending *to bake a cake with their parent, if their parent is willing*. If this were their intention, then it leaves open that it is a joint action, assuming we accept a conditional intention model of joint intention.

I am reluctant to so understand what the child intends. In part this is purely stipulative — I have already stated that the intention with which a supported agent engages in a dependent action is an intention to ϕ . However, I want to say some more to shore up this claim — it is not something I stipulated purely for the sake of stipulation. This accurately captures the intention with which the agent acts. Whilst acting dependently, an agent is aware that the way in which they realize their intention will involve the help of another agent. They may not have this in sight in formulating a prior intention, since, as I argued in §4.4, dependency is not a feature of act-types, nor is it necessary that an agent can *only* ϕ dependently for them to ϕ dependently on a particular occasion, as I argued in §5.1. I might intend to do something, expecting to struggle with it alone, before you offer your help. Of course, this could lead to us doing it together, but it could also lead to my doing it dependently. This does not change the content of the intention with which I am acting, and my successful dependent action would realize my intention to act in this way. The unexpected nature of the help is not

necessary for the intention to be an intention to ϕ rather than an intention to ϕ with the supporting agent. In §1.3.1 I discussed the conditions on what an agent can intend, and the relationship between the act-type that an agent intends to engage in and the means by which they will be acting. In particular, I argued that agents do not intend the means of their action. If they do, it is not as a means but as something that is itself being intended. That I have to ϕ in some specific way does not mean that an intention to ϕ is an intention to ϕ in that way, even if I have that way of ϕ -ing in sight. That the supported agent recognizes that they will need to be helped does not mean that what they intend is to ϕ with help. Act-types are generic, and part of intending is identifying the specific way in which that act-type will be realized on a particular occasion. That does not mean that one intends that specific way of realizing the intention, but rather that one intends to do the generic thing and that this specific method has been settled on, at least for the time being. Moreover, even if the supported agent has the more precise intention, it does not entail that this is the intention with which they are acting — an agent can be doing something with multiple intentions. I can be walking to the station both with an intention to enjoy the pleasant weather, and with an intention to catch my train. Similarly, I could be walking to the station with an intention to walk to the station, and with an intention to get home, where the former is a way of realizing the latter. Even if the child is acting with an intention to bake a cake with their parent if their parent is willing, this does not rule out that they are intending this as a way of realizing an intention to bake a cake *simpliciter*.

This is supported by reflecting on what would happen if the supported agent started to ‘go it alone’. Imagine the supported agent, starting to get a hang of the act-type. In this situation, the supporting agent would likely recognize their help is no longer needed. The driving instructor would stop giving instructions, the parent would leave the child to their own devices. In this situation, we would be loathe to think of the supported agent as failing in realizing their intention. They have not been thwarted by the removal of help, since they set out to ϕ , and they successfully ϕ -d. If being supported were part of what they set out to do, then this would make no sense — they would have been thwarted by the supporting agent’s decision to pull back. What this suggests is that, although the supported agent may well have expected the help, what they set out to do did not rely on it. They might even believe the help is ongoing, feeling comfortable in their acting in virtue of the belief that they will be helped if needed. However, what they intend to do cannot involve this help, else the removal of help would prevent them from realizing their intention. This is the situation that we see in joint action — what each agent intends requires the involvement of the other, and if one withdraws from the action, that intention is thwarted. We can see this on all the approaches to the intention involved in joint

action that I detailed in §2.4. If there is a plural intention, then this falls apart when one agent pulls out of the action, since there is no longer a plural subject that so intends. If there is a participatory intention, this is an intention to do one's part of the joint ϕ -ing, which cannot occur without the other agent doing their part. If there is a conditional intention, this can only be realized if the other is willing, which is not the case if the other agent pulls out. If there is an Anscombean intention, then it is an intention to jointly ϕ , and so can only be realized if the other agent is involved. Given that the supported agent's intention is realizable without the participation of the supporting agent, they do not seem to be acting with an intention to act together; that is, with the intention involved in joint action.

5.2.2 The supporting agent

Although I have suggested we should not understand the child as having the conditional intention *to bake a cake with their parent, if their parent is willing*, I recognize that it's possible that this is one of the intentions with which the child is acting. That is, it is possible that the supported agent, as well as intending to ϕ , also intends to ϕ with the supporting agent, if the supporting agent is willing. The way in which they're intending to ϕ with them, though, is likely not the same as the way in which an agent of a joint action intends to ϕ with another agent. It might be that we can capture the child's intention with this form of words, but really what the child intends is not to bake a cake with their parent, but to bake a cake with their parent's *help*. Nonetheless, it is not crucial that we determine whether or not the child has this conditional intention. In considering the intention of the parent, the divergence from the joint action case becomes apparent. If this were a joint action, the parent would need to have the same intention — it is not enough that the child has an intention to engage in a joint-action, the parent must also have it. What the parent intends to do is fundamentally different to what agents of a joint action intend to do — the parent does not intend to bake a cake together with their child. The parent intends to help their child bake a cake. Their intention is explicitly supportive. This does not mean that a parent *cannot* intend to bake a cake together with their child, but that this is a different intention directed at a different action. In the case at hand, the cake-baking occurs only because the child wants it to, and the parent's role is as supporter and aid, not equal partner in the action. The parent's intention contains reference to the intentions, not simply the actions, of the child that they help. It is not just that they intend to dovetail with the actions of the child in bringing about their shared end, but that their intention itself has in its content a reference to realizing the child's intention — the parent intends 'to help my child (realize their intention) to bake a cake'. We can see this even more clearly in the cases where the act-types involved do not seem especially apt for joint action. In

Yoga, the supported agent moves their body into a particular position. We might do yoga together, we might even hold this particular pose together, but this would involve each of us adopting the pose ourselves. We can imagine doing this as part of a joint action, for instance if this pose were part of a piece of group choreography. Saying that A and B held the yoga pose together, however, makes far less sense. The driving lesson case is similar — if the instructor and student described themselves as ‘reverse parking together’ we would think this bizarre. In dependent intentional actions, the supporting agent intends ‘to help the supported agent (realize their intention) to ϕ .’

I have already rejected the possibility that a plural, participatory or Anscombean model could capture the intention of the child, and this is enough to rule out that these approaches cannot capture this event as a joint action. It is, nonetheless, worth noting that they also cannot capture those of the parent. The parent does not intend a joint cake baking, or to do their part in one. Conditional intentions again seem our best bet for this to be a joint action. They include reference to the intention of the other agent – both agents intend to bake a cake together, if the other is willing, that is, if the other has the same conditional intention. What both conditionally intend to do is a joint action. In the case of joint cake-baking, both intend to bake a cake together, given the other has the conditional intention to jointly bake a cake. Each intention is an intention to act together, and it includes in its content the other’s intention to act together as a condition of so intending. The parent’s intention would need to be the intention *to bake a cake with their child, if their child is willing*. The thing that both are willing to do needs to be the same. However, in the dependent cake-baking, the child’s intention is not just the condition of the parent’s willing, but is the thing about which they intend. They do not intend the same thing as the child — to bake a cake — but rather they intend to support the child in doing what they intend to do. We do not have the requisite symmetry between the intention of the parent and of the child for this to be a case of joint action. The parent does not act with the same intention as the child — the intention of the child is primary, it is that towards which the parent’s intention is directed. This is the way in which the parent’s intention is dependent — it depends on the child’s intention for its existence and its content. It is centred on the intention of the child, and is an intention to help bring about the realization of the child’s intention. Given the lack of match between the intention of the child and the parent, this is not a joint action. Instead, this is a dependent intentional action of the child. The action is dependent, because it depends on the support of the parent, and the parent’s intention is dependent, because it depends on the intention of the child. A supported agent intends to ϕ , and a supporting agent intends to help them to ϕ . In joint action, a joint ϕ -ing is aimed at by all the agents, but the supporting agent does not intend to be the agent of a ϕ -ing of any kind.

They intend only to help with a ϕ -ing, and so that ϕ -ing cannot be a joint intentional action. The supporting agent's explicitly supportive intention rules out the possibility that dependent intentional actions are joint actions.

5.3 Dependent intentional action is not action of the supporting agent

Having shown that a dependent intentional action is an action, and is not a joint action of the supported and supporting agent, there remains one alternative to its being an action of the supported agent. This is the possibility that it is instead an action of the supporting agent. In this section, I will give initial reasons for rejecting this possibility in general, before turning to consider two particular reasons for seeing it in this light — understanding what happens as a trying and failing on the part of the supported agent, followed by action of the supporting agent, or understanding what happens as an action of the supporting agent in virtue of a deviant causal chain between the intention with which the supported agent is acting and the action in question.

If we return to the cake-baking, why might we think that this is an individual action of the parent? They are the one that reads the recipe. They take the final steps that turn the cake mix into the cake. The cake is only baked because of them. However, this does not seem enough to think of what happens as the parent baking a cake. Although we might claim that I have baked a cake if I go into the kitchen, find a spring-form tin full of cake mix, and place it in the oven for the requisite time, this would be an unusual thing to say. When we describe someone as baking a cake, we normally mean that they did much more than this. Similarly, in the case where the cake mixture is made by a child, we should not think of placing the mixture into the oven and removing it as sufficient to have baked a cake. The parent has not baked a cake any more than I would have when I stumbled upon cake mix. If we were to praise the quality of the cake-baking, we would praise the child, not the parent. We would not compare it to cakes that the parent had baked alone in determining whether it deserves praise, but to what the child has done before or what we expect of them. If the cake produced was impressive, it would be impressive for the child. We might praise the parent for what they did to help the child bake well, but we would not praise them for baking the cake — it is as supporting agent that they can excel. This suggests that we do not think of this as their action, especially given the argument of Chapter 3 that praise or blame for an action presupposes identity with the agent of that action.

Most importantly the parent simply does not have the right intention for this to be their action. The parent does not intend to bake a cake. It is the child's idea, and the parent acts only in order to help the child. If the child changes their mind and demands to read a book instead, the cake-baking is abandoned. If there is an intentional individual cake-baking, there needs to be an intention to bake a cake that it realizes, or it needs to be something done intentionally in the realizing of another intention. In the case I described, we have a case of the first kind — there is an intention to bake a cake, but it is not one held by the parent. The intention that initiates the action is an intention held by the child. The action itself is guided by the child. This is true even though the parent is the one reading the recipe, because this is being done in response to the child. The parent would not simply plough on through the recipe if the child was stuck on a step, nor would they do bits without explaining them to the child. There are, of course, cake-bakings that look like that, but they are not dependent cake-bakings. The action is happening because of the child and their intention, that intention shapes the action, and if the child lost interest, it would stop. These features of the action are all guaranteed by conditions (D2-D6), and for this to be an action of the supporting agent would clash with this. In §1.2 I gave the following requirements for intentional action:

Intentional Action

An event X is an intentional ϕ -ing of an agent A only if

- (I1) X is an action of A 's ϕ -ing
- (I2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting
- (I3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

Given that I have stipulated that the supporting agent (B) is acting with an intention to help the supported agent (A) to ϕ , if this were B's individual action, B would be ϕ -ing because it contributes to helping A to ϕ . This is not impossible — I might demonstrate how to ϕ to help someone to ϕ — but in the cases in which this happens, the ϕ -ing that B is doing is not the selfsame ϕ -ing as the one with which they are helping A — it is a different token of the same act-type. That an event could be a helping with itself is incoherent.

Despite the above, I recognize there are compelling reasons for denying that the supported agent is the agent of a dependent intentional action. Over the next two sections, I will consider and reject what I take to be the most forceful and plausible of those reasons.

5.4 Dependent intentional action and ways of not ϕ -ing

Recall the distinction I drew in §1.2.2 between actions, mere attempts at action, and failures to act. There, I stated that *mere attempts are attempts to ϕ that do not culminate in the agent's ϕ -ing*, and that *a failure to act is when an agent's not ϕ -ing is in virtue of their not doing something that they need to in order to ϕ* . Given this distinction, someone might object as follows:

‘Sure, the “supported” agent *intends* to act. But they don’t. After they’ve failed, someone else comes along and does it for them’

The objection is that what we have is not a dependent ϕ -ing by an agent, relying on the support of another agent, but instead an agent acting with an intention to ϕ , but not ϕ -ing, followed by an intentional ϕ -ing by another agent. The proposed structure would be:

- (1) *A* intends to ϕ ⁵²
- (2) *A* does not ϕ (either by failing to ϕ or merely attempting to ϕ)
- (3) *B* ϕ s

This casts its net too wide. Things with this structure happen all the time. Sometimes *B*'s ϕ -ing depends on *A*'s failing to ϕ , such as when ϕ -ing is winning a race they are both running. Sometimes *A*'s attempt and *B*'s action are unrelated, they are both attempting the same act-type by coincidence. The objector would have to concede that in the cases I am calling dependent intentional action, *B*'s action is not only causally related to *A*'s attempt (as it is when *B* wins the race) but that *B* acts because of *A*'s failure. They are not merely ϕ -ing. They are ϕ -ing because of *A*'s failure to ϕ . They are ϕ -ing *for A*. That is, they are ϕ -ing because *A* intended to bring about some end that they did not succeed in bringing about, and so they are ϕ -ing for *A*'s sake. This would be doing what I will describe as ‘acting on someone’s behalf.’

⁵² Although this framing seems to imply that they are acting on a prior intention, this is not necessary. We can understand their intention to ϕ as that with which they are attempting.

5.4.1 Acting on someone's behalf

Unlike the structure proposed in the previous section, in which one agent abandons an action, and another carries out an action of the same type, acting on another's behalf has the following structure:

Acting on someone's behalf

- (1) A intends to ϕ
- (2) A does not ϕ (either by failing to act or merely attempting to ϕ)
- (3) B ϕ s
- (4) B s ϕ -ing contributes to the realization of an intention to ψ with which B is acting
- (5) B ϕ s because their ϕ -ing contributes to their ψ -ing

In this context, an intention to ψ is an intention *to bring about A's desired ends*. On this picture, then, B does not help A to do what A intends, but 'helps' A by doing what A intends to themselves do. They act for A . As I have characterized dependent intentional action, it offers a kind of supported independence — it facilitates A 's agency by enabling them to realize their own intention. If these are in fact instances of acting on another's behalf, rather than enhancing someone's agency, they seem to strip them of it.

There is no doubt that *acting on someone's behalf* accurately describes some cases.⁵³ A child might try to get a book from a high shelf, fail, and then a parent get it down for them. I might plan to knit myself a new scarf, not have the time, and later you gift me one. I might attempt to pay for lunch when we are out together, realize I left my card at home, and you pay instead. These are all cases of an agent attempting to ϕ , and someone else ϕ -ing in a way designed to bring about the end state of the first agent's ϕ -ing. The child has the desired book, my neck is warm, our meal is paid for. However, as I have made clear throughout the preceding chapters, we often intend to ϕ not just to bring about a state of affairs, but because we are engaged in some other action. Agents ϕ because they are ϕ -ing — it is a way of realizing another intention. Sometimes that intention will still be realizable if someone

⁵³ There is also a kind of acting on someone else's behalf that goes further than this, potentially requiring no intention on the part of the agent on whose behalf it occurs. Lasting power of attorney, for instance, enables someone to take decisions for another beyond the point at which they are able to do so themselves. This does not stand as a possible explanation of dependent intentional actions, given the absence of the relevant intention on the part of the person on whose behalf the action occurs.

else ϕ s for us. If the child is getting the book because they are reading the book (because they want to read the book), if I am making a scarf because I am going for a walk on a cold day (because I want to go for a walk on a cold day), or if I am paying the bill because I am leaving the restaurant with our debts discharged (because I want to leave the restaurant with our debts discharged), then someone else ϕ -ing will still allow the realization of those overarching intentions. If, however, the child is reaching for the book because they want to be the one to get the book, because they like doing things for themselves, then if the parent gives them the book they cannot realize this intention. If I am knitting a scarf because I want to engage in the process of knitting, because I find this kind of straightforward project especially calming, then you buying one for me does not allow me to realize this intention. If I want to treat you to lunch as a way of thanking you for being a supportive friend, then you getting the bill, even if I pay you back, removes this possibility and spoils the treat. If I am ϕ -ing not just in order to ψ , but in order to ϕ , then you ϕ -ing removes my ability to do that. Nonetheless, sometimes, that I am ϕ -ing ‘wins out’, such that even though I would ideally ϕ myself, your ϕ -ing to enable my ϕ -ing is helpful. We do things on others’ behalf all the time, and often those we act for appreciate it. The question is not whether *acting on someone’s behalf* is something that occurs, but whether this characterizes the cases I have called dependent intentional action.

Why might this be the correct characterization of the cases I have called ‘dependent intentional action’? For some of the cases I have discussed, it is implausible to characterize them as one agent not- ϕ -ing and another agent picking this up and ϕ -ing themselves. Often, the contributions of B (the supporting agent) will be interspersed with those of A (the supported agent). In *driving lesson*, B’s support is ongoing, but so too are A’s contributions to the action. The image of them getting started, stopping, and the supporting agent taking over just does not match the way that the action goes. However, perhaps there is a subset of dependent intentional actions that do not look like this, in which the A’s contribution is temporally bounded such that they play no role in the later parts of the action, and the worry could still be raised about these. If we turn, for instance, to *Dressing*, A does not complete the act of dressing all by herself. She tries to get dressed, picking out the outfit, but she cannot manage all the buttons. Instead, B, as well as helping her manipulate her arms into her sleeves, does up those buttons that A cannot. Perhaps, then, we should understand what happens as A trying and failing to dress herself, and B subsequently dressing A? Getting dressed is not the kind of act-type that we tend to engage in for the process itself, but rather to bring about the end state of being-dressed, and so it is the kind of thing that B could do for A. Of course, she might well be frustrated at not being able to do it herself, and her intention to get dressed would be frustrated, but she might, on

balance, appreciate B doing this for her. If we are to understand this as a kind of failure on A's part, then what kind of failure is this? In §1.2.2, I argued that there are two ways in which an agent that intends to ϕ does not successfully ϕ — those in which the failure is down to the agent, and those in which it is down to the circumstances. If an action does not come about because of me, there is a *failure to act*. If an action does not come about because of circumstances outside of my control, there is a *mere attempt* at ϕ -ing. In both of these cases, the agent is correctly described as failing to realize their intention, but only in the second kind of case can we think of the agent as trying to realize their intention. Failures to act are cases of an agent *giving up* on realizing their intention — perhaps they are easily dissuaded, perhaps they are easily distracted — before they have tried to realize it. This is not to suggest trying is a substantive or distinct thing that agents do, but rather to point to the contexts in which an observer would say ‘well, of course they didn’t do it, they didn’t even try!’ Part of intending to ϕ is having some sense of what one needs to do to ϕ , although of course we need not have all the details worked out. An agent that intends to ϕ , knows that in order to ϕ , they must ψ , and also that ψ -ing is within their grasp, should set about ψ -ing. If they do not, then we are within our rights to think of them as failing to ϕ , rather than merely attempting to ϕ . For the cases I have framed as dependent intentional action to be acting on another’s behalf, the supported agent must be unsuccessful in realizing their intention to ϕ — either because they fail to ϕ or because they merely intend to ϕ . Given this, I will now consider whether either of these possibilities captures what the supported agent does — if they do not, it does not seem that the supporting agent can be acting on behalf of the supported agent.

5.4.2 The supporting agent as not ϕ -ing

Some of the discussion in this and the subsequent section will echo the discussions in §5.2, in which I discussed the intention of the supported agent, but I am turning my focus to a different question here. In this section, I will consider whether dependent intentional actions can be understood as either failures to act or mere attempts on the part of supported agents. If, for instance, *Dressing* is a case of A's not ϕ -ing, then it needs to be of one kind or the other. In all of the examples of dependent intentional action I have given, the supported agent does *try* to ϕ . They set out on ϕ -ing, but they don't get all the way themselves. The buttons are too fiddly, the details of the stack of rings elude the child, the manoeuvre is too intricate for the driver and the yoga stretch is too deep. They all set out to act, but none of them can manage it alone. Are they, then, failures to act? In *Dressing*, A does not do up her buttons, because she cannot do it in virtue of her own abilities. The buttons are small and her joints

ache. What is missing from the attempt is something that would be contributed by A, but her abilities are such that she cannot contribute it. What is missing from the attempt to dress herself is the rest of the buttoning, but she lacks the requisite abilities. She does all she can to realize her intention, because what constitutes all that an agent can do depends on the particular abilities of the agent and the situation in which they are acting. What it is for an agent to do all that they can in relation to some act-type depends on what they can do. A does not fail to act, because she does what she can, but perhaps we should understand this as a *mere attempt*, and *Dressing* as a case of acting on someone's behalf.

If we have a mere attempt, then A must not have done what she set out to do. Whether this is accurate rests on the content of the intention with which A is acting, but also on how she envisages realizing her intention. In the case of *Dressing*, it is clear that A was aware in advance that she may well need help in dressing herself, since she arranged for B to come in the morning to help her with it. She might not know for sure which things she will struggle with on each day, but she knows that she may well struggle with something. What she intends is not to get dressed all by herself, but to get dressed. This is the intention she forms, and this is the intention she realizes. The way she expects to do this is *with Barbara's help*. There is no failed action, because there is nothing she intends to do and subsequently fails to do. In §1.3, I set out Anscombe's A-D series, which gives us, at its end:

the intention *with* which the act in each of its other descriptions was done, and this intention so to speak swallows up all the preceding intentions *with* which the earlier members of the series were done.

([1957] 2000, p46, §26)

This intention is the intention to dress herself. As the A-D series reminds us though, we often form intentions along the way that are swallowed up by this final intention. In *Dressing*, one of these intermediate intentions is an intention to dress herself *with B's help*. This is a more general feature of dependent intentional action. Some cases are like *Dressing*, in which A perhaps always recognizes that they will need help, but this help is nonetheless a means to the end of realizing the final intention, which is to ϕ . In other cases, however, the help will be more incidental. I might set out to bake a cake, and you offer to help me in some way. That I did not foresee your help in advance of my acting does not mean that I cannot realize my intention by being helped. Act-types are general, and actions are particular — there will always be elements of how I ϕ that I did not have in mind when I set out to ϕ ,

and sometimes those details will change as the action progresses. If I intend to go to the library, initially plan to travel by train, but discover that it's delayed, I might get the bus instead. Practical reason enables us to change how we realize our intentions when we encounter hurdles and new ways of doing what we set out to do, and we form intermediate intentions in response to those. One of these intermediate intentions can be an intention to ϕ with help.

5.4.3 Intending to ϕ with help

If I want to suggest that dependent intentional action can involve an intention to ϕ with help on the part of the supported agent, this kind of intention has to be legitimate. I noted in §1.3.1 that we cannot intend the actions of another agent. Is intending to be helped intending the actions of another? If it is, then this intention is illegitimate. Let us reflect again on the two ways in which an agent might form this intention. If the agent does not expect to be helped, then there's no potential problem. In this scenario, A does not intend that B helps them, but rather B offers to help them and A takes them up on this offer. The help is a means to their end that presents itself, not something they intend to bring about. In cases like *Dressing*, however, A sets out to do something they know they can only do with help. Does this require that they intend B's action? Although an agent cannot intend the actions of another agent, they can intend to induce another agent to act. In some cases of dependent intentional action, this might be what happens. This is a less altruistic kind of helping than most of the cases I have discussed, but it does recall some of what I set out in §2.2. Coercion and other kinds of inducing agents to help are things that can be intended. However, in many of the cases I have discussed, it does not seem right to think of A as inducing B to help them, at least with this particular action. Perhaps the financial relationships in *Dressing*, *Yoga* and *Driving Lesson*, induce B to help A, but it seems inaccurate to frame a child as inducing a parent to help them. Even in the cases where A pays B, this happens upstream from this particular act of helping, it is a condition of B's availability to help A. A does not think about how to bring it about that B helps her in intending to act, that B will help her is something that is taken for granted in intending. This is not to say that she feels entitled to the help, but rather that she intends something that is only realizable if B helps, and thinks it is likely that B will do this. Often we intend things that take for granted that certain things will be the case if we are able to realize that intention, but this does not mean that we intend to bring about those conditions. Assuming that B will help is of course not the same as, for instance, assuming that the shop will be open when I intend to buy ingredients, because B is an agent in their own right, and that agency is integral to the support that they will give A. It is more like intending to engage in a joint action, which

also involves an assumption that the other will act. As I argued in §5.2, it is not in fact intending to engage in a joint action, however, because the intentions of the agents involved lack the requisite symmetry. What this comparison reminds us is that it is not unusual to intend to do things that depend on someone else acting. In *Dressing*, A is not intending to make B help her, but rather to act, if B will help her. Since she knows that B likely will help her, this is a reasonable thing for her to intend, but of course she cannot realize this intention if B decides not to. If she refuses to help, A might try to persuade her, but this persuasion is not that in which A set out to engage. What she set out to do is dress herself, and this persuasion, if intentional, is something that she would do in service of realizing that intention. A might even form an intermediate intention to get B to help her, but this intention would not be the one with which A is acting. That the realization of A's intention relies on B doing something does not make it something she cannot intend.

Given that A does not form an illegitimate intention about someone else's action, two things seem true about this case. The agent intends to ϕ , and the dependent intentional action realizes that intention. In realizing that intention, they form an intermediate intention, to ϕ -with-B's-help. This ϕ -ing-with-B's-help is a way of ϕ -ing, and it is something they are doing in order to realize their intention to ϕ . A supported agent realizes their intention in a dependent intentional action, and to think of them as either merely attempting to ϕ or failing to ϕ would be to ignore this. For them to fail, there should be an intention that they do not realize, such that there is something they fail to do. Only if they intended to ϕ -without-help would the dependent intentional action constitute a failure. Although *acting on behalf of another* does accurately characterize some actions, it does not capture all of the cases that I am calling dependent intentional action. These are distinct kinds of case. The supported agent ϕ s, the supporting agent helps them; they do not take over for them.

5.5 Deviant causal chains

There is a remaining possibility that I must rule out if I am to claim that dependent intentional actions are actions of the supported agent. Dependent intentional ϕ -ings are, I have suggested, ϕ -ings by an agent who is acting with an intention to ϕ . However, that an agent intends to ϕ is not enough to guarantee that their ϕ -ing realizes that intention. In §1.3, I introduced the principle I dubbed Realization, which states 'An event X realizes an agent A's intention to ϕ only if X is A's ϕ -ing'. It remains open that an agent ϕ s, whilst intending to ϕ , but this ϕ -ing does not realize their intention

since it is not an intentional ϕ -ing. It is only an intentional ϕ -ing if the intention plays the right role in the ϕ -ing.

It is not enough to say, for instance, that the intention to ϕ leads to the ϕ -ing, or is the reason that $A \phi$ s. It is possible that a deviant causal chain obtains between the intention and the action. Although the intention leads to the action in this situation, it does not bring it about in the right way. The canonical example⁵⁴ of a deviant causal chain is the one given by Davidson as follows:

A climber might want to rid himself of the weight and danger of holding another man on a rope, and he might know that by loosening his hold on the rope he could rid himself of the weight and danger. This belief and want might so unnerve him as to cause him to loosen his hold, and yet it might be the case that he never chose to loosen his hold, nor did he do it intentionally.

(Davidson 2001d, p79)

This is a case in which an agent's intention to do something brings about that they do it, but not in virtue of their intentionally doing it. Clearly, though, this is not the kind of thing I have in mind when I speak of the realization of intentions — this is not the appropriate way for an intention to bring about the end at which it is directed. This is the problem that deviant causal chain cases point to — sometimes an intention can lead to what is intended, but not in the right way. Of course there are multiple possible ways of realizing an intention. Davidson's climber could remove his fingers from the rope in any order, he could do it quickly or slowly. If I intend to bake a cake, then there are all sorts of cakes I could bake which would each mean a different set of steps. These wouldn't be deviant causal chains, because the intention is not merely related to the action, the action realizes the intention. Deviant causal chains are deviant because they are not ways of realizing the intention.

5.5.1 The relevance of deviant causal chains

The relevance of worries about deviant causal chains to dependent intentional action becomes clearer if we turn to an example given by Peacocke. In considering what is sufficient for the non-deviancy of a

⁵⁴ The other central example is Chisholm's nephew, which is as follows:

'Suppose a man believes that if he kills his uncle he will inherit a fortune and suppose he desires to inherit a fortune; this belief and desire may agitate him and cause him to drive in such a way that he accidentally kills his uncle; but it will not follow, as the definition would require, that the man has done anything for the purpose of inheriting a fortune.' (Chisholm 1964, p616)

causal chain, he imagines a case in which ‘the chain from intention to bodily movement pass[es] through the intentions of a second person.’ (1979, p87) He gives the case as follows:

This second person might for instance be a knowledgeable neurophysiologist who decides on a particular occasion to produce in me exactly the motor impulses needed to realize what he knows, from my neurophysiological states, to be my intentions. Is my bodily movement really intentional when my arm moves exactly as I intended it to? It is not plausible to say that it is so without qualification.

(Peacocke 1979, p87)

Of this, Peacocke says ‘[i]t is not clear whether there is such a person as *the* originator of the bodily movement in our example, but if there is, it is certainly not the person whose brain the neurophysiologist is inspecting.’ (ibid, p88) He is adamant that, if we are to develop an account of intentional action that does not include those events caused by deviant causal chains from intention, it must include the requirement ‘that the chain from intention to bodily movement should not run through the intentions of another person.’ (ibid)

Pears considers something similar as follows:

...a physiologist always observes the impulse in a patient's motor nerve, and always reproduces it artificially beyond the point of severance. In such a case the agent's movement of his right index finger would not be his intentional basic action, in spite of the fact that its causation started from the essential initiating event, and that it was reliable.

One reason that might be given for this verdict is that the intermediate stages are not the normal ones for a human agent. But this cannot be quite the right reason. For we would allow that it was a basic intentional action if the gap in his motor nerve were bridged by a prosthetic device. Such a device would be an action-aid in the sense in which people now have hearing-aids. But the physiologist and his apparatus could not be regarded as the patient's actions-aids, because the apparatus was not regularly attached to his body and its successful operation required the physiologist's agency.

(Pears 1975, p67)

Pears and Peacocke have developed a picture on which some of the cases that I would describe as dependent intentional actions are in fact not cases of intentional action, since the intention does not cause the event in the requisite way. Dependent intentional actions involve one agent's intentions being realized in part through the intentional action of another agent, and the idea that causal chains of this kind are deviant is one that is more widely accepted. Bishop uses the term *heteromesial* to refer to causal chains that run 'from intention to matching behavior through the intentional actions of a *second* agent' (1989, p125), although a causal chain's being heteromesial does not require that it solely passes through the intentional actions of another agent.

In §1.3.1, I characterized a basic action as something an agent can do 'just like that', such as moving a limb, without needing to perform any further actions. The focus on bodily movement in the above passages indicates that the picture presented by Pears and Peacocke is one on which *all heteromesial causal chains to basic action are deviant*. If the realization of intention to engage in basic action goes via the agency of another, then it is not in fact a realization of that intention. A causal chain connecting an intention to a bodily movement that goes via another agent, such as the supporting agent of a dependent intentional action, is a deviant causal chain. The neurophysiologist's action is akin to the anxious sweat on the climber's palms. It is caused by the intention but not in a way that means that the intending agent brought about the intended outcome through their intentional action.

Although this is framed in the terms of causal approach to intentional action, the worry motivating it is present for all models. We do not want to include in our account of intentional action cases in which it seems wrong to think of the agent as intentionally bringing about the intended outcome. The Anscombean model is such that it does not run into this risk as often, but it is not immune to the threat, it could not see the climber's partner being dropped as an intentional action. Although the intention to drop their partner is what led him to drop them, it is not *why* he dropped them. If you asked him why he dropped his partner, he would tell you that he did not, or at least, that he did not do it on purpose. This kind of deviant causal chain is already debarred from being an intentional action on the Anscombean approach. The examples with the neurophysiologist, however, are trickier. Of the person whose movement is caused in this kind of way, Pears notes that 'he is, and feels that he is, the originator of an intentional action, even though he may know that, in a sense, the experimenter is its originator' (1975, p66). Given this, he may well have an answer to Anscombe's 'why?', since he will understand himself to be acting, will be attempting to act, and will be indirectly bringing about his intended outcome. If we cannot allow intentions to engage in basic action to be realized through the

intentions of another agent, then we seem at risk of many dependent intentional actions in fact not being actions at all.

The worry, then, is for cases in which the dependent intentional action is something that we would think of as a basic action. Although I have not identified any of them as basic actions, since some of these central cases consist of simple bodily movements, it seems worth considering this possibility. These cases are also far from the only examples of dependent intentional action that might be considered basic action. Consider, for instance, a case of dependent nail-painting⁵⁵ — perhaps I ask you to open the bottle of nail varnish, since this particular movement causes me wrist pain, and you support my arm in order to steady my hand as I paint the nails on my right hand, since I am less dextrous with my left hand. It might be difficult to think of these kinds of actions as things that the agent can do ‘just like that’, since in many cases it is exactly this ability that they lack. These actions are dependent because the agent cannot realize their intention ‘just like that’, they need help. Nonetheless, if they are bodily movements, they seem to be cases of potential basic action, and so the problem arises. Moreover, in at least some of these cases, the help is with something that someone can do unaided, just perhaps less well. Thus, it seems right to consider some of these as potential instances of deviant causal chains.

5.5.2 Basic actions

There are two strands to my rejection of this analysis of dependent intentional action cases. Firstly, I want to resist the idea that these are really cases of basic action. I will then set out reason to think that the heteromesial chains involved in dependent intentional action are not deviant, even if we think these are in fact basic actions and are thus apt candidates for deviant causation. To take the first of these, my claim is this: the description under which this event is deviantly caused is not the description under which this event is a dependent intentional action. At one point, Peacocke suggests:

For instance, if I suffer from akinesia, and I get another person to move my arm for me, and this movement is a signalling to a third person outside the window, it may not be wrong to say that the movement of my arm (*not* my moving it - no such thing occurred) is intentional under a description relating to signalling.

(Peacocke 1979, p88, fn23)

⁵⁵ Thanks owed to Rosie Parry for suggesting this act-type

This acknowledges the importance of the description under which an agent intends to whether we can understand them as intentionally acting. He suggests that if we accept that this kind of case is non-deviant, then we should restrict the prohibition on heteromesy in realizing an intention 'to descriptions that are 'basic' in the agent's plan.' (1979, p88, fn23). What is it for an action to be basic in the agent's plan? Judging by the waving example, it seems that what matters is how the agent describes that which they are doing. If the agent describes their action as a movement of their body, then it is basic in their plan, if they otherwise describe it then it is not. In *Dressing*, *Driving Lesson* and *Stacking Rings*, even though the agent obviously intends to move their body, and they do not need some downstream effect to come off to succeed, they do not conceive of the action as a movement of their body. They conceive of it as a dressing or a parallel parking or a stacking of rings. It is exceedingly rare for an agent to hold in mind all of the bodily movements that are needed for them to act in the intended way. Even in *lôga*, the agent does not solely conceive of this as a bodily movement. They conceive of it as a particular yoga pose, as part of completing the routine, as improving their strength or flexibility. Quite often, then, there is reason to resist the applicability of the threat of deviant causal chains to dependent intentional actions.

5.5.3 Deviant causal chains and guidance

Even if the reader is not convinced by the claim that these actions are not basic in the relevant sense, there is a further line of resistance to the challenge that heteromesial chains to basic action are necessarily deviant. The challenge that deviant causal chains provide to an account of action is a familiar one in philosophy, that of distinguishing between the good cases of a certain phenomena and the bad ones. Deviant causal chains are discussed not by opponents of a causal approach to action, but by proponents of it — they are used to refine approaches so as to ensure that any account of action does not cast its net too wide. There are a variety of solutions to the general problem of deviant causal chains.⁵⁶ Harman (1976) points to the difference between the good and bad cases in the following way, discussing a case much like Chisholm's uncle, in which Mabel accidentally runs over Ted as she backs out of her driveway on her way to his house to run him over:

The point is that, even though she intends to kill him by running over him and does kill him in that way, she does not do what she intends. One kills someone

⁵⁶ See (Mayr 2011) and (Piñeros Glasscock and Tenenbaum 2023, §2.4) for discussion of the landscape of proposed approaches.

intentionally if one intends to kill him, does kill him, and thereby does what one intends; but in this case Mabel does not do what she intends. Her intention is not simply "to kill Ted" or "to kill Ted by running over him." It also includes a plan specifying how that intention will lead her to do what she intends to do. She does not in the example do what she intends, because what she does differs significantly from the plan that is part of her intention.

(Harman 1976, p444)

In this case, it is clear that the event of running Ted over does not realize Mabel's intention, for she did not simply intend to run him over, but to do so in a particular way, having located him first. The solution that Harman gives rests on the idea that 'a positive intention to do something is the intention that that very intention will lead in a more or less explicitly specified way to one's doing the thing in question' (ibid, p445). Although he acknowledges that there are various ways in which such a plan might change over time, he is nonetheless committed to the idea that the intention to ϕ involves a plan for how one will ϕ , and that deviant causal chains are cases in which the ϕ -ing does not happen in accordance with those plans. I have reiterated that I do not understand intentions as involving commitments to a particular way of ϕ -ing — they involve commitments to identifying a way of ϕ -ing, but that is not the same as the plan being necessarily built in to all intentions. There might be intentions to run over Ted that *are* realized by hitting him as one backs out of the drive. Despite this, the idea that the event has to come off in the way that the agent intends is clearly key to the solution.

If we return to the characterization of dependent intentional action that I gave in §4.4.1, my strategy for dealing with the challenge of causal deviance is already apparent:

(D5) A guides X such that X occurs in keeping with their direction

In motivating this condition in §4.2, I noted that the notion of guidance I was using was one drawn from Frankfurt (1978). He rejects the causal approach to action, but this idea is still valuable in setting out how cases of causal deviance differ from those of genuine action. Setiya (2003; 2007) draws out the details of how a causal approach might use this notion, but since I am not committed to such an approach, and I have suggested that the worry raised under the title of 'deviant causal chains' is one that the Anscombean need also address, I will not do this here. Instead, I want to spell out what Frankfurt means by guidance, and why the presence of this in dependent intentional actions is sufficient to resist the accusation of deviance. Guidance, as Frankfurt presents it, is the difference

between a mere bodily movement and a bodily action, and it is, rather than something that comes before action (such as an appropriate cause of plan), something that happens whilst the agent is acting. Frankfurt notes that:

during the time a person is performing an action he is necessarily in touch with the movements of his body in a certain way, whereas he is necessarily not in touch with them in that way when movements of his body are occurring without his making them.

(Frankfurt 1978, p158)

Guidance is the way in which an agent is in touch with their action — they guide the movements that their body makes. This is how Frankfurt distinguishes intentional actions from other movements — actions are intentional movements, and intentional movements are those that are guided by an agent. He captures this as follows:

Behavior is purposive when its course is subject to adjustments which compensate for the effects of forces which would otherwise interfere with the course of the behavior, and when the occurrence of these adjustments is not explainable by what explains the state of affairs that elicits them. The behavior is in that case under the guidance of an independent causal mechanism, whose readiness to bring about compensatory adjustments tends to ensure that the behavior is accomplished. The activity of such a mechanism is normally not, of course, guided by us. Rather it is, when we are performing an action, our guidance of our behavior.

(Frankfurt 1978, p160)

This is a description of the kind of guidance that an individual agent has of their own action, but it is also extendable to dependent intentional actions. The supported agent need not have a detailed plan in advance of how they will realize their intention, and the dependency might be unexpected, but as long as they stay in control of the way in which it occurs, they guide it. They adjust to what the supporting agent does, and they suggest adjustments to the way in which the supporting agent helps them. Those suggestions are, as I state in (D6), something to which the supporting agent is responsive. Dependent intentional actions occur in the way that the supported agent intends them to, and with their guidance throughout, even if this guidance is less direct than the guidance we have over the movements of our bodies alone. Frankfurt imagines a driver whose car coasts downhill at a speed with

which the driver is happy, and who thus never intervenes to alter the speed. ‘This would not show that the movement of the automobile did not occur under his guidance. What counts is that he was prepared to intervene if necessary, and that he was in a position to do so more or less effectively’ (1978, p160). The supported agent is able to intervene in how the action comes off, including bringing it to an end, and this means that the action is occurring under their guidance in the way that makes it their intentional action.

5.5.4 Kinds of heteromesy

The guidance of the supported agent is, I believe, enough to conclude that dependent intentional actions are not in fact instances of deviant causal chains. The supported agent has the right kind of relationship to the event, that of ongoing guidance, to make it their intentional action. It might be objected that the issue was never that these are deviant causal chains in general, but that they are heteromesial chains. The issue is that the intentions of another agent figure in the event. For this particular worry, I want to turn to Bishop (1989), from whom I borrowed the term ‘heteromesy, and who argues that not all cases of heteromesy are cases of causal deviance. He notes that:

Sometimes the second agent’s involvement in the causal chain *preempts or blocks the agent’s exercise of direct control* over his or her bodily movements. Then the second agent is no mere cog in the mechanisms that realize the first agent’s direct control. Rather, the second agent is part of a system that provides the first agent with, at best, only *indirect* control over the movements of his or her own body.

(Bishop 1989, p159)

The problem arises in cases of ‘*preemptive heteromesy*’ (ibid), cases in which the second agent’s input impedes the control that the first agent would otherwise have. Pre-empting involves cutting in on what the first agent is doing, doing it instead of them or for them. In what I will call *supportive heteromesy*, the second agent does not preempt the first agent’s ability to realize their intention, but facilitates it. Bishop’s example is of ‘a person fitted with a successful prosthetic neural replacement, which one day breaks down but is briefly repaired by having a second agent intentionally hold the broken wires together until they can be resoldered’ (ibid). Here, the heteromesy allows the agent to realize their intention, rather than impeding this.

His concern, like Peacocke's, is with cases where the intervention of another agent happens at the neurophysiological level, and the use of examples of this kind is, I think, part of the reason we are so inclined to see heteromesial chains as inimical to something's being an intentional action. He resolves the problem through a notion of 'continuous regulation' that he draws from Thalberg (1984), suggesting that '[w]hat matters is whether the 'agent *remains in direct control of* his or her bodily movements, though only by dint of another's action' (1989, p159). This notion of direct control is necessary, perhaps, for the kind of cases that Peacocke and Bishop discuss, but not for dependent intentional actions. Here, guidance is sufficient. Consider the case of dependent nail-painting. I am the one in control of the brush, you steady my hand but which nail is painted next, where on the nail I begin, and whether I give any nails an additional coat of varnish all fall on me. Clearly, I am guiding my behaviour. My guidance does not need to go via the actions of another in the way that it does if the other agent is bound up with the neurophysiological mechanisms that are involved in the transition from intention to action. In both Peacocke and Bishop's examples, the internal feedback mechanisms are facilitated by the second agent, whereas in my case they are involved in my acting on that internal feedback. The same is true in *Yoga*, as, for instance, A's pain or discomfort will impact how he furthers the stretch, and B's input is only to enable A to do this. B might indicate particular sensations that they expect A to be feeling, to identify those to be avoided or encouraged, but recognizing and responding to those sensations will be down to B. In dependent intentional actions, the heteromesial chains we find are not deviant, because the action occurs under the guidance of the supported agent. It is the action of the supported agent.

In this chapter, I have defended (D1), the claim that dependent intentional actions are actions of the supported agent. I have done this by considering the landscape of possible alternative views. By rejecting these in turn, whilst considering the relationship that the supported agent bears to a dependent intentional action, I have shown that the best understanding of these events is as actions of the supported agent. In §5.1, I argued that these events are in fact actions, rather than some other kind of event. In §5.2, I argued that they are not joint actions, by considering first the intention of the supported agent in §5.2.1 and then of the supporting agent in §5.2.2. I then turned to the possibility that these are actions of the supporting agent, rejecting this in §5.3 by considering how the agents involved would understand this action, and its failure to meet the conditions I gave on individual intentional action in §1.2. I then turned to a particular version of the argument that this is the action of the supporting agent in §5.4, considering whether we should understand them as acting on behalf of the supported agent. In §5.4.1 I argued that so understanding the action misses what it is the

supported agent intends to do. In §5.4.2 I considered why we cannot understand this as either a mere attempt by the supported agent or as a failure to act. In §5.4.3, I discussed the legitimacy of intending to act with help. In §5.5, I turned to the challenge that the causal chain connecting the intention of the supported agent to the action is deviant, such that the action does not realize their intention. In §5.5.1, I set out the worry that this notion presents, explaining the idea of heteromesial causal chains in particular. In §5.5.2, I argued that those dependent intentional actions that appear to be basic actions are not, since they are not so understood by their agents. In §5.5.3 I drew out the notion of guidance that is present in the characterization of dependent intentional action, and finally in §5.5.4 I suggested that a causal chain going via the intentions of another agent is insufficient for its being deviant. Having rejected the alternative explanations in this way means that understanding dependent intentional actions as the actions of the supported agent is legitimate, and is preferable to alternative explanations we might offer. A dependent intentional action is the action of the supported agent.

Chapter 6 – Barriers to intentional agency in early childhood

In this chapter and the next, I am interested in the kind of agency that young children have. I am setting aside dependent intentional action for now, to return to it in chapter 7. In this chapter, I will be considering two questions – can young children engage in individual intentional action, and can young children engage in joint intentional action? I will not presume that these questions have a hard and fast answer as applies to all act-types — it may be that a young child can individually intentionally ϕ , but not individually intentionally ψ , and yet can jointly intentionally ψ . In particular, I will be interested in what I will call *novel* act-types — those that the child does not know how to perform — since these are the ones it seems most likely that a young child could not intentionally perform. My conclusion will be that there are no insurmountable barriers to individual and joint intentional action in early childhood. This is not the same as arguing that young children do in fact engage in either individual or joint intentional action, but rather that there is nothing that prohibits them from so doing. It will, nonetheless, be left open at the end of this chapter that, rather than acting intentionally, young children engage in a more minimal kind of purposive action. In the following chapter, I will set out the way that dependency facilitates intentional action in early childhood.

I will begin in §6.1 by considering what is distinctive about young children as agents. I will then consider a potential barrier to a young child's intentionally ϕ -ing — that they do not know how to ϕ . In §6.2 I will introduce the idea that one might need to know how to ϕ in order to act with an intention to ϕ , and in §6.2.1 I will discuss how this concern applies to young children in particular. In §6.3 I respond to this concern. I will give several reasons to reject the possibility that this lack of knowledge prevents young child from intentionally acting, arguing that young children have the understanding of at least some novel act-types that is requisite for intending to engage in them. I will discuss in §6.3.1 how it is that we determine that an agent is acting intentionally, suggesting that such an understanding applies to at least some action of young children. In §6.3.2 I will consider the role that membership in a community of language users plays in the ability of ordinary adult agents to act with an intention to ϕ , and turn in §6.3.3 to the knowledge that we have of ϕ -ing in intending to learn to ϕ . These considerations will lead me to the conclusion in §6.3.4 that an agent does not need to

know how to ϕ in order to act with an intention to ϕ or to ϕ intentionally, but rather needs to *know what it is to ϕ* , which is a more minimal kind of knowledge. I will defend the claim that the following kind of knowledge — knowing what it is to ϕ — is sufficient for intentionally ϕ -ing:

Knowing what it is to ϕ is having sufficient understanding of ϕ -ing to know that one is ϕ -ing when one is ϕ -ing, and to identify which kinds of ϕ -ing might be ways of ϕ -ing.

This knowledge includes an understanding of the kind of act-type that ϕ -ing is, such that an agent that knows what it is to ϕ could identify some elements of ϕ -ing in advance, but it does not require being able to form a detailed plan for ϕ -ing, or being able to engage in all the steps of ϕ -ing without guidance.

I will then turn to the question of joint intentional action in §6.4. In §6.4.1 I will show that in order to engage in joint action, one needs to conceive of the other as an agent. I will set out in §6.4.2 why we might doubt a young child to be capable of this, considering experimental evidence, before turning to evidence in §6.4.3 that complicates this account. I will then, in §6.5, consider alternative explanations of the evidence, setting out alternative approaches to the ability of young children to understand the agency of others in §6.5.1 and §6.5.2. I will conclude in §6.6 that we should recognize young children as capable of engaging in joint action. Regardless of whether the joint actions that young children can engage in are intentional, it will be clear that they can recognize others as agents. Young children are purposive agents, who have sufficient understanding of novel act-types to intend to engage in them, and who can recognize others as agents and engage in joint action with them. I will end with a note of caution, that we should not necessarily taken this to be evidence of intentional action rather than of purposive action on a non-intentional kind. Nonetheless, it is evidence that these are not insurmountable barriers to young children acting intentionally, either individually or jointly.

6.1 Early childhood

To understand the nature of young children's agency requires first understanding who I mean by young children, and what their agency looks like. There isn't clear agreement about the term 'early childhood'⁵⁷ but I will be generally focussed on children aged 0-5, and often especially focussed on

⁵⁷ Unicef (2022) says 0-8 years whereas the Center on the Developing Child at Harvard University sometimes prioritizes ages 0-3 (e.g. 2023) and at other times considers the broader range of 0-5 years (e.g. 2007)

children aged 2-5. These ages aren't given as a firm rule, not least because the developmental trajectories of individual children can be radically different. They are given more as an indication to the reader of the kind of agents I am discussing. Children of this age are rapidly changing and acquiring new abilities, but they are not yet fully fledged rational adult agents. This is illustrated in the ways that we treat young children. For instance, in the UK, the age of criminal responsibility is 10,⁵⁸ as it is 'conclusively presumed' (*Children and Young Persons Act 1933* (23 & 24 Geo. 5. c. 12), section 50) that children below this age cannot be guilty of any criminal act. This stems from a presumption these children are not capable of 'distinguishing between right and wrong' [2009] UKHL 20 (p6). Being incapable of distinguishing between right and wrong is, in part, being incapable of fully appreciating the nature of various act types and their consequences in the way that we often associate with intentional action. Although these laws are not making a direct claim about intentional action, it does indicate that we treat children of the ages I am considering as having more limited agency than those above that age.

Often a young child will want to do something they have never done before. In early childhood, we are continually doing things for the first time. This is a straightforward feature of this being early in one's life — there is just not much opportunity to have done things before — but also a consequence of the developmental trajectory of the child. MacMurray contrasts human infants with those of other animal species, who, although dependent, are able to engage in some of the things needed for their survival from birth:

The child has to start from scratch and has to learn everything. All his skills are acquired... [and] the child's acquiring of skills is a cumulative process. Simple skills are used in acquiring more complex skills and the process goes on indefinitely.

(MacMurray 1961, p55)

As we learn to do more things, we can, with our new-found skills, attempt new act-types. In the first twelve months of life, children develop the capacity for visually-guided reaching (Bruner 1973). That is, they come to be able to act on a desire to grasp an object, using their visual system to enable them to successfully reach for and grasp it. Once the child has come to be able to do this, they then develop the ability to do things that depend on this skill. For instance, between 8 and 12 months of age,

⁵⁸ This is codified in section 50 of the *Children and Young Persons Act 1933* (23 & 24 Geo. 5. c. 12) as amended by section 16 of the *Children and Young Persons Act 1963* (c. 37)

‘detour-reaching’ emerges. This ‘requires reaching away from the goal object at the outset of the reach’ (Diamond 2006, p71) so as to avoid obstacles in their way.⁵⁹ Only in virtue of first being adept at straightforward visually-guided reaching, involving no obstacles, can the child come to engage in this more complex kind of reaching. What this speaks to is the way that over the child’s developmental trajectory, new action-possibilities open up to them. It is only once they have a handle on how to do the things that are prerequisites for and building-blocks of an act-type that it becomes a possibility for them. The speed at which they are learning new ways of acting means that there are continually new things for them to do. A young child is constantly doing things for the first time.

To think a little more about what young children can and cannot do, let us consider the anecdotes in the following tweets:

not saying my toddler knows how to use a pub but he did just wake up from his nap and head towards the bar shouting “i get.....drink!!”

(Miller 2022b, on Iggy, age 35 months)

if betty’s hat falls off when she’s on my back i stop a friendly looking person and ask them to put it back on her because i can’t reach. this system has now broken down because she’s learned that if she takes her hat off and throws it on the ground she gets to make a new friend

(Miller 2022a, on Betty, age 13 months)

In both of these tweets, we have a description of a child that clearly wants to do something, and acts to bring about that which they want. What they are doing, however, is radically different to the way that a rational adult agent would go about doing those things. The toddler heading off towards the bar does not stop to check if anyone else wants anything or take care to ensure he has a method of paying. He simply knows that one gets drinks in a pub and that one does it from the bar. His younger sister, dropping her hat on the floor, deploys a method of making a new friend that is not available to you or I. If I were to drop my hat, it’s possible that someone else will see me do it and give it back to

⁵⁹ These examples take place in infancy, but this does not rule out their relevance to my argument. They are early cases of a phenomenon that persists, but the smaller number of act-types available to the infant makes the way that one grows out of others more clearcut. The features of early childhood that make dependent intentional action harder to argue for than in adulthood are present in infants, such that attending to them does not stack the deck in my favour. There are just additional barriers to the action of this younger cohort, such that turning to them cannot make things easier for me!

me, but there's no guarantee of this occurring, because Betty's mother is not there to ask someone. Even if it is returned to me, I am not an adorable small child, and so our interaction will not play out in the way that hers does. That these young children are ϕ -ing in a way that I would not ϕ is not itself a problem. There are lots of ways of realizing an intention to ϕ , as I have repeatedly noted. Betty's strategy generally works, and Iggy's likely will not, but this is also not a problem. I can act intentionally and unsuccessfully. The question is whether either of them are acting intentionally.

To consider whether young children are able to engage in intentional action, I need to have in sight what intentional action requires. In §1.2, I gave the following conditions for some event's being an individual intentional action:

Intentional Action

An event X is an intentional ϕ -ing of an agent A only if

- (I1) X is an action of A 's ϕ -ing
- (I2) This ϕ -ing contributes to the realization of an intention to ψ with which A is acting
- (I3) A is ϕ -ing because their ϕ -ing contributes to their ψ -ing

These are necessary, but potentially insufficient, conditions on some event's being an intentional action of an agent A , and so for a young child to be able to engage in such an action, there need to be actions of theirs that satisfy these conditions. The first condition, that the event is an action of the child, is not a problem — young children clearly act. The second and third conditions, however, provide some requirements that an agent must meet in order for their ϕ -ing to be an intentional action. From (I2), it is clear that the agent needs to be able to act with an intention (even if it is not an intention to ϕ). (I3) places a condition on how the agent relates to their ϕ -ing. The agent needs not only to be acting with an intention to ψ , but to be ϕ -ing because this contributes to their intention to ψ . This requires some degree of understanding of ψ -ing and ϕ -ing, and this is what I will explore in the following sections.

Although merely intentionally ϕ -ing might seem less complex than acting with an intention to ϕ , it requires that an agent understand that ϕ -ing is something that they will do in the course of realizing their intention to ψ . To reduce the demand on the agent, I will generally consider only intentional ϕ -ings that occur when an agent acts with an intention to ϕ . This only requires the agent to understand

one act-type in the requisite way in order to intentionally engage in it, rather than understanding two and having a grasp on the relationship between them. This means that what I need to determine is: *are young children capable of ϕ -ing with an intention to ϕ ?*

6.2 Know-how

In the following sections, I will consider a potential barrier to a young child's ϕ -ing with an intention to ϕ , when ϕ is a novel act-type. As I stated in the introduction of this chapter, a novel-act type is one that a child does not know how to perform. In the following sections, I will first set out why we might think know-how is necessary for acting with an intention to ϕ . I will then set out three reasons for denying the necessity of know-how for acting with an intention to ϕ , arguing instead that an agent needs to *know what it is to ϕ* . A precise analysis of this concept is beyond the scope of this project, but it describes something fairly intuitive, and I will draw out two key elements:

- knowing what it is to ϕ is sufficient for knowing that one is ϕ -ing when one is ϕ -ing
- knowing what it is to ϕ entails knowing enough about ϕ -ing to identify which kinds of ϕ -ing might be ways of ϕ -ing

'Know-how' is a substantive philosophical area in its own right, and this thesis is not a project in epistemology. This does not prohibit me, however, from using the term 'know-how.' Analyses of know-how do not disagree as to whether this is a kind of knowledge, but about the form this knowledge takes. Even the 'intellectualist' who Ryle tells us 'assumes that knowledge-how must be reducible to knowledge-that' (1945, p8) acknowledges the existence of knowledge-how as something we can recognize and discuss. Given this, I will take for granted that we can describe an agent as knowing how to do various things, and not knowing how to do others. I will use this term in a way that is perhaps more expansive than some uses of it, but not without precedent, at least in ordinary speech. I know how to speak and walk and read, which I can do without thinking at all about how to do them. I know how to knit, which I can do generally by feel and memory, only occasionally visually inspecting my work to check it is going to plan. I know how to bake a cake, although I might need to check through a recipe before and during baking. I do not need to have a cake recipe memorized to know how to make a cake, I need to know how to understand an instruction to cream together butter and sugar, to know the importance of precise measurements, to know how to salvage wet ingredients that have split, and to know how to tell when a cake is done and when it needs longer in the oven. I,

arguably, do not know how to speak German — I can understand a significant amount of German, beyond a rudimentary level, but my grasp of the language is insufficient to describe myself as knowing how to speak German, rather than, for instance, ‘being able to get by’ in the language. To borrow an example from Lewis (1976), I also do not know how to speak Finnish, but my lack of knowledge here is different. In German, I might trip over genders and adjectival endings, or not know a particular word. I could not even confidently identify whether a given sentence is in Finnish rather than one of several other languages I do not know how to speak.

Knowing-how to ϕ is not necessary for ϕ -ing. This is not a novel thought; we can find it in Aristotle:

it is possible to produce a grammatical result by chance, or by following someone else's instructions. To be grammarians, then, we must both produce a grammatical result and produce it grammatically – that is to say, produce it in accord with the grammatical knowledge in us.’

(Aristotle, 1999, Book II, Chapter 4, §§1-2, 1105a)

That is, I can do something that is apt for know-how without knowing how to do it. I could do it by sheer accident. What might seem similarly uncontroversial is the claim that I cannot act with an intention to ϕ if I do not know how to ϕ . How could I possibly intend to speak to you in Finnish, if I do not know how to do that? To be acting with an intention to ϕ is to be directing one’s agency at realizing that intention, thinking about how I will realize that intention and what I need to do to do those intermediate steps. If we return to the conditions I gave on intentional action, often what I am doing is intentionally ψ -ing because ψ -ing contributes to realizing the intention to ϕ with which I am acting. In order to be acting with an intention to ϕ , I need to be committed to my ϕ -ing (rather than merely interested in it, for instance). If ϕ -ing is done by ψ -ing, rather than ‘just like that’ then I also need to be able to ψ in order to ϕ . What this perhaps suggests, then, is that I cannot act with an intention to ϕ if I do not know how to ϕ .

6.2.1 Early childhood and know-how

If know-how is necessary for intending, then by definition, one cannot intend to ϕ , when ϕ -ing is a novel act-type. Given, as I discussed in §6.1, young children engage in a huge number of novel act-

types, there is a significant barrier to young children acting with intentions. The act-types that are not novel, especially in very early childhood, are those that we may be inclined to doubt have the status of intentional action at all. Are the movements of our limbs intentional? Are cries when hungry intentional? Whether you conceive of these actions as manifestations of intentional agency will partly rest on how you understand the role that practical reason plays in agency,⁶⁰ but regardless, we might be loathe to think of these as things as done with an intention. Even if there are some actions that are intentional, if intending required know-how then the frequency with which a young child does something they do not know how to do entails that there would be reason to think that little, if any, of the young child's action is intentional.

To illustrate the extent to which what the young child does involves their engaging in act-types that they do not know-how to engage in, I want to consider the developmental trajectory of visually-guided reaching, which I mentioned in §6.1. Visually guided reaching is something we do constantly — when I answer my phone, or I pick up my fork to eat, or reach for my cup of tea, I am engaging in visually guided reaching. However, this is something that the young child has to learn how to do, over many attempts. At first, on identifying an object for which they want to reach, there is ‘...prolonged looking and, very shortly after, there is action of the mouth and tongue and jaws — the area to which a captured object will be transported once effective, visually guided reaching develops’ and after this comes ‘...antigravitational activity of arms and shoulders, clenching of fists in a "grab" pattern, movement of arms, ballistic flinging of clenched fists’ (Bruner 1973, p2). What the infant wants to do is identify the location of the object, reach for it, grab it, and bring it to their mouth with their clenched fist, and then likely gnaw (or gum) on it. At this early stage, although they are engaging in these constituent acts, they do not occur in the right order to bring this about. ‘In time, and probably *by virtue of sheer practice of the act...* the act is successfully executed. An object is captured and brought to the mouth’ (ibid, p3 emphasis mine). Only by continually trying and failing to grasp the desired object does the infant come to be able to do this successfully — although they know from the start what they want to do, they do not initially know how to do it. Even in their first successful attempt, which we would describe as visually-guided reaching (rather than attempted or unsuccessful visually-guided reaching) it is not clear they know how to do it.

Once the act is successfully executed and repeated with success, that is, constituents are put stably into proper serial order, there often appears a sharp alteration in the

⁶⁰ See §1.4

structure of the act used for achieving an intended outcome. For example, shortly after the first successful taking of an object, the fist, rather than being closed prematurely which often happens before successful capture, now remains open at maximum extension until the object being sought is touched.

(Bruner 1973, p3)

Only once they are already able to doing this thing do they begin to refine it, altering their behaviour in ways that we would think of as part of their knowing how to do it. In the early attempts, they do not know how to engage in visually-guided reaching. It is clear, nonetheless, that in these early attempts, they want to do it. Of course we can know how to do something and still get better at it, such as the way that over time the child's hand 'begins to close gradually to the shape of the object as it approaches rather than after it gets there' (ibid, p4). Unlike this change, which is arguably a refinement for the sake of efficiency, knowing that you should not close your fist before you grasp the object you are reaching for seems more like a basic component of reaching for things than it does polishing a skill. It is part of what it is to know how to reach for an object. This is an extremely basic act-type, on which a huge number of other act-types rely. If this cannot be done with an intention to do it without knowing how to do it, then it seems there are massive barriers to understanding young children as acting with intentions. Of course, this skill is learnt at a very young age, but similar patterns will play out across early childhood with increasingly complex skills, and the same problem will persist. Rödl describes just such a case as follows:

Consider a child, three years old, let us say, learning to carve a stick. The child wants to carve his stick... He knows it involves a knife and somehow applying the knife to the wood. But that is about it. And yet he wants to do it... he wants to do what she, the parent is doing or does.

(Rödl 2016, p94)

This is a case in which the young child wants to ϕ , but does not know how to ϕ , and cases of this kind are common. Given that they do not know how to ϕ , it seems worth saying a little about how it is they can want to ϕ . Firstly, ϕ -ing could be an act-type that features in the child's desire *de dicto* — they know it by name or description, having heard another child on the playground talk about doing gymnastics, or visiting the zoo, and know that whatever that thing is, that is something they want to do. That they cannot fill in the details will not dissuade them, they are sure it is something fun and exciting,

something they are certain that they want to engage in, even if they are wrong, and the act-type in question is something deeply tedious. They may well persist in doing it when they discover this, in a kind of stubbornness designed to show that they really did want to do it all along. Alternatively, ϕ -ing might feature in the child's desire *de re* — they want to do that thing there, that that person is doing, without knowing what it is by name. 'That thing that Daddy is doing', 'the thing that they were doing in that TV show' or simply 'that', accompanied by them pointing at someone engaging in it. They may well express a desire to join in with some ongoing activity, without knowing what that activity is. In this second kind of case, the child can come to know the name of the act-type, and they seem to have more of a sense of what it involves than in the first case.⁶¹ Nonetheless, it seems possible that they do not know what it is to ϕ — if we told them that they were more than welcome to ϕ , they would not know where to start. They want to ϕ , but do not know how it is that one ϕ s — they do not know what ϕ -ing is.

6.3 Intending to ϕ and knowing how to ϕ

Although, as I have suggested, there is *prima facie* reason to think that an agent must know how to ϕ in order to act with an intention to ϕ , over the following sections I will argue that this is not the case. We need to know something about ϕ -ing, but it is not necessary to know-how to ϕ . I will first give three arguments against the idea that an agent must know how to ϕ in order to act with an intention to ϕ . Having done this, I will set out what an agent must know about an act-type in order to act with an intention to engage in it. This is what I call *knowing what it is to ϕ* . I will then give reason to understand young children as having that knowledge.

Consider the following passage from Anscombe:

Since a single action can have many different descriptions, e.g. "sawing a plank", "sawing oak", "sawing one of Smith's planks", "making a squeaky noise with the saw", "making a great deal of sawdust" and so on and so on, it is important to notice that a man may know that he is doing a thing under one description, and not under another... For this reason, the statement that a man knows he is doing X does not imply the statement that, concerning anything which is also his doing X, he knows

⁶¹ There is not necessarily a sharp distinction — the child might both know that they want to do *that* and also that *that* is called knitting, but even in this hybrid case they do not know what knitting is in the way that I do.

that he is doing that thing. So to say that a man knows he is doing X is to give a description of what he is doing *under which* he knows it.

(Anscombe [1957] 2000, pp11-12, §6)

Agents know what they are doing, and the intentions with which they are acting, under a description. That description is often a name of the act-type. Understanding our actions in this way does not require knowing everything that so acting involves. That the young child does not know that in wanting to ϕ , they want to engage in an act-type that necessarily involves their ϕ -ing, does not necessarily prohibit that them from intending to ϕ . This may seem like a slightly cheap way out of this problem, since what is at stake is whether one can really intend to ϕ without this kind of knowledge of ϕ -ing, but it indicates the direction in which I will be proceeding. What I need to determine is whether the use of ‘ ϕ ’ by agents that do not know how to ϕ picks out ϕ -ing in the kind of way needed for acting with an intention to ϕ . In the following sections, I will give three reasons for thinking that knowing-how to ϕ is not necessary for an agent’s intention to ϕ to be a genuine intention – ‘ ϕ ’ can often pick out ϕ -ing even if the agent does not know how to ϕ .

6.3.1 Identifying intentional action

The best guide to what an agent is doing what they take themselves to be doing, although observing them often tells us without needing to ask. Anscombe reminds us that attributions of actions ordinarily work by simply attending to other agents. She notes that, if you were asked what someone was doing ‘...in a very large number of cases, your selection from the immense variety of true statements about him which you might make would coincide with what he could say he was doing’ ([1957] 2000, p8, §4), noting that ‘[she is] sitting in a chair writing, and anyone grown to the age of reason in the same world would know this as soon as he saw [her]...’ (ibid). We determine what someone is doing by attending to them, both to their actions and their claims. We can often work out what they are doing by looking at them, but the conclusions we draw through observation can be overruled by the agent — this is why what matters in that first Anscombe quote is that our attribution would coincide with the agent’s own claim. Some actions are such that we could only understand them as being done in virtue of the agent setting out to do them — the observer that sees Anscombe writing will not think that this is being done accidentally or unintentionally. Even if she would much rather be doing something else, there is *some* pro-attitude held towards writing by her, else it would not be occurring. It can’t happen by

accident or unintentionally (even if coerced). Given this, what reason have we for denying that the child who states that they want to carve a stick, that obtains the materials they think it involves, that closely watches their parent carving a stick, is acting with an intention to carve a stick? That they do not know how to do this does not seem enough, especially since they will often vocally express their desire to do this in just the way that you or I might express an intention.

I do not, however, want to suggest that there is some particular expression of intention that is necessary for acting with an intention. This is true for ordinary adult agents, but especially for young children whose linguistic capacities often restrict their ability to express what they are doing. The CDC's child developmental milestones (2022) suggest that at 30 months a child should be saying about 50 words, and that child of this age '[s]ays two or more words together, with one action word, like "Doggie run"', with their communicative skills gradually increasing, such that at 4 years a child '[s]ays sentences with four or more words'. These are language/communication milestones, not cognitive ones, and so they do not obviously mirror the trajectory of understanding of the things that these words pick out. A lack of expression of intention does not indicate a lack of intention. Piaget suggests that '...verbal judgement lags behind effective judgement: the idea of autonomy appears in the child about a year later than cooperative behaviour and the practical consciousness of autonomy' (1983, pp113-14). Taking this seriously means that we should not look to what the child can tell us about what they can do, but what the child can do.

We ordinarily know what someone is doing simply by observing them. This is what it is to be an intentional agent situated amongst intentional agents — what they are doing is legible to us from attending to their behaviour. We recognize them as intentional agents, and we see not a series of movements, but the action of an agent. We may not be able to know exactly why it is that someone is doing something, nor which of the many reasons for which they could be doing it is operant for them. When we look at someone, we see that they are pumping the water, and we understand this as something they are doing intentionally. We do not need to know the intention with which they are acting to know this, just to recognize it as something done with an intention. Some actions are things only doable by intentional agents — like writing a thesis — whereas some others are doable by intentional and non-intentional agents alike — like eating and drinking. Some of the actions of the young child are legible to us as intentional, even if we think that some of their actions are not, in the same way that the actions of the adult agent are so legible. This is the case even when the child does not express an intention in advance of or during action — sometimes we can tell by looking.

It might be argued that the situation described above does not apply to the action of small children. Of course we read the actions of rational adult agents as intentional, because we know them to be capable of intentional action. Similarly, we are not going to understand the behaviour of a robot vacuum cleaner as intentional, because it is not capable of intentional action. However, this is to misunderstand the role that recognizing intentional actions plays in our understanding something to be being capable of intentional action. Part of this, of course, comes from preconceptions about the kind of agent that it is. I could take the same route around my living room as a robot vacuum cleaner, and an observer would almost certainly understand my action as intentional, but not that of the vacuum cleaner. A large part of this is, of course, that they would recognize me as an adult human, and the vacuum as a tool that had been designed and programmed for a particular purpose. What it can do and does is very limited, and so questions about the nature of its agency simply do not arise. However, often it is less straightforward than this, and we find ourselves looking at what an agent is doing to determine the kind of agency they have. Consider, for instance, the way one might understand their pet. I know that my cats have desires and beliefs of some kind. They sit near the kitchen door when they are angling for an early dinner, because they want food and believe (correctly) that this is where their food comes from. That this is what they were doing is clear from the enthusiasm with which they eat if this does coincide with their meal times. Sometimes, however, they seem to do things that belie more complex mental states — getting up on the surfaces that they know they are not supposed to be on (the dining table, the record player), since we will have to get up to chase them off, and if we do that we might be more likely to follow them to the kitchen to feed them. Of course, this is in the main just the fondness of a pet owner, believing in the intelligence of the animals of which they are especially fond. If, however, the animal exhibited repeated behaviours, best understood as involving more complex mental states than we think the animal capable of, we would start to question the assessment of their abilities. This is exactly what happens in situations where animals do appear to exhibit unexpected capacities. Koko the gorilla, for instance, was understood by some as using sign language to lie (Brooks 2018) and as combining words to create new vocabulary (Silver 2016), but by others as doing no more than ‘flailing around producing signs at random’ (Pullum 2018). What is contested here is whether she did in fact do the things some saw her as doing, because if so, this entails that she had more complex capacities.

Consider the anecdote in the first of the tweets which I presented in §6.1. The expression “[I] get..... drink!!” is not something that the child in question would describe as an intention, but it seems to

express one. Whether or not the child has the conceptual framework to understand this as an intention is not what matters, what matters is that it is the intention with which they are acting (or, at least, attempting to act). Acting with an intention does not require using the language of ‘intention’ and its cognates in its expression, nor does it require expressing it at all. It might not even require language, as the second tweet’s anecdote about the toddler’s 13-month old sibling implies. To understand her behaviour as anything other than an action directed at meeting new people requires reading “against the grain”, looking for an alternative explanation of her action to the one that makes most sense of what she is doing and that comes most naturally. Whether an agent is capable of intentional action is something we in part determine based on whether they do things that we recognize as intentional. We do not decide if they can act intentionally and then decide if their behaviour is so legible. It is therefore reasonable to ask whether we can recognize young children as engaging in intentional action, as a way of determining if they are so capable. This is not to beg the question about children’s capacity for intentional action, but rather to engage in the standard process for identifying the capacity for forming intentions. The most coherent understanding of the anecdotes detailed in the tweets above is that they are instances of goal-directed action, and quite possibly intentional action. This supports the idea that we should look to what young children do, rather than what they know-how to do, to determine whether they are able to act intentionally.

6.3.2 Semantic externalism

The reason for insisting that intending to ϕ requires knowing how to ϕ is that an intention to ϕ is taken to require a thorough understanding of what ϕ -ing involves. Of course, an agent need not be certain which of the ways of ϕ -ing available to them is the one they will adopt,⁶² but nonetheless they need to know how to ϕ . Without this knowledge, it might be thought that an agent cannot settle on a method of ϕ -ing, and go about ϕ -ing, which they are able to do if they intend to ϕ .

⁶² In the background of this recurrent discussion of the relationship between means of acting and intention is the distinction that Travis (2013) draws between the ‘general’ and the ‘historical.’ Expressions such as ‘I intend to ϕ ’ are general in ways that particular actions are not — there are many ways of ϕ -ing, but any instance of ϕ -ing cannot contain that generality. An instance of ϕ -ing is a historical particular, and is limited in virtue of this. Many other, different, events could also have satisfied my intention to ϕ . See (Wallage 2020) for a discussion of this distinction – although Travis, and thus Wallage, focus on perception, the distinction is also illuminating in thinking about action. Something similar comes out in Davidson, when he notes that ‘Any one of an indefinitely large number of actions would satisfy the want and can be considered equally eligible as its object.’ (2001a, p6) as well, of course, in the sections of Anscombe (2005, pp142) I discussed in §4.3.3

If we require this kind of understanding of what an act-type involves for an agent to be able to form intentions involving that act-type, we will have to draw implausible conclusions. Often, we intend to do something with a more minimal understanding of what the act-type involves. Consider Burge's arthritis sufferer, who, failing to understand that this condition involves inflammation of the joints, relays to the doctor their concern that they now have arthritis in their thigh (1979, p77). Burge argues that, given the patient accepts the doctor's clarification, and recognizes that their thigh-pain must have an alternative source, their use of the word 'arthritis' did in fact refer to arthritis. Concurring with this is to acknowledge that the meaning of one's language (including in thought) stems not only from the things that we associate with it, but also from its meaning in the community of language-users of which one is a part. In particular, Burge argues that this is true of these words when they have *oblique-occurrences* — that is, uses that are 'not freely exchangeable with all extensionally equivalent expressions' (1979, p76). This is exactly the kind of occurrence we have in mind when we think about desires or intentions to act. Imagine the patient, who has not yet learnt that arthritis cannot occur in their thigh, reminding themselves that they intend to pick up their prescription because they intend to lessen the pain that their arthritis is causing in their knee. Here, it seems clear that they do intend to do this, and yet they don't know what it is to do it, because they do not know what arthritis is. What matters for their being able to have intentions about their arthritis, following Burge, is not what they individually know about exactly what the term 'arthritis' refers to, for instance, but that their use of these words is part of a community of language users, and that they are being used to refer to the things to which they do in fact refer.

The consideration of Burge (and semantic externalism more generally) shows us that the cost of concluding that we can only intend to do things if we have a complete understanding of what they involve is too high. Someone can intend to treat their arthritis symptoms, even if all they know is that this will involve speaking to a doctor, and perhaps trialling several medications. They can intend this, even though this makes the act-type that features in their intention 'treat arthritis', and this includes a term whose meaning they do not understand, and relies on mechanisms they do not know. What matters is not whether I fully understand what ϕ -ing involves, but that I have sufficient grasp on it, as part of a community that fills out what these terms mean. There is a description under which they can intend to treat their symptoms, if they know what kind of things they need to do to. Similarly, there is a description under which the child understands what they want to do, even if the details of what that involves need to be filled out by the presence of other language users in the community. Maybe they know it involves getting on a train first (the zoo), or that they will need to wear special clothes and no

shoes (gymnastics). They might not know how to do it, but they know enough of what it is to do it to get started with it. This does not mean that they do in fact intend to do it, but it cannot rule out their so intending, without also ruling out that adults intend to engage in act-types like treating their arthritis symptoms. The child has the same kind of knowledge of ϕ -ing that any of us has — limited, in part constituted by external features, and just about enough for our intentions to grip on to ϕ -ing. To treat the child as unable to intend to ϕ has unpalatable consequences for all intentions towards act-types that we do not understand in fine-grained detail.

6.3.3 Learning to ϕ

There may well be cases in which we want to insist that the child's lack of knowledge of ϕ -ing is significant enough that we cannot attribute to them the intention to ϕ . They really do not know anything about ϕ -ing, so they really cannot intend to ϕ . They cannot intend to reflect on their professional development, for instance, because they cannot hold any pro-attitudes towards this act-type. There are too many concepts involved that they cannot begin to get a handle on. There are also cases where it seems like the act-type is one that is in principle attainable for them, but they nonetheless lack knowledge of it, such as the young child that wants to tie their shoelaces or tell the time, but does not yet know how. Unlike in situations where they know a bit of what is involved in ϕ -ing, but are still working out the details, these can be situations where they do not know how to get started with ϕ -ing at all. First, I will suggest that intentionally learning to ϕ is often sufficient to be acting with an intention to ϕ . I will then also suggest that, for some values of ϕ , acting with an intention to learn to ϕ involves intentionally ϕ -ing, because these are act-types that one learns by doing. These relationships between learning and doing indicate that know-how is unnecessary for intentionally ϕ -ing. Here, I am diverging a little from focus on ϕ -ing with an intention to ϕ . However, in §6.1, I motivated this focus on the grounds that the other kinds of intentional ϕ -ing involved acting with an intention to ϕ , and thus put a further demand on the agent. When the ϕ -ing in question is learning-to- ϕ , this worry does not arise, and so the shift to intentionally ϕ -ing seems permissible. If the reader is unhappy with this, the first of my arguments, that intending to learn to ϕ is sufficient for acting with an intention to ϕ , clearly targets the child's ability to act with an intention to ϕ .

As I have noted, early childhood is a time of constant and rapid development — young children are attempting all sorts of actions for the first time, and also gradually learning how to successfully carry them out. Although the young child learns to do things at a greater frequency than the rest of us, they

are not the only agents that can learn new act-types. The following is a common form that this might take:

I intend to ϕ , but I do not know how to

I intend to learn how to ϕ , which will enable me to ϕ

Thus: I intend to learn to ϕ because I intend to ϕ

I can act with an intention to ϕ when I do not know how to ϕ , because I can intentionally learn to ϕ . What I need to know about ϕ -ing in order to act with an intention to ϕ , is what I need to do in order to get going with ϕ -ing. Learning to ϕ is sometimes part of how one gets going with ϕ -ing, and a part of how one realizes an intention to ϕ . I need not know what learning to ϕ will teach me (if I knew this, I would not need to learn it) to know that it will enable me to ϕ , and thus to realize my intention to ϕ . Much as I can intentionally crack an egg because I am acting with an intention to bake a cake, I can intentionally learn to ϕ because I am acting with an intention to ϕ .

Moreover, we often learn new act-types by doing them before we know how to do them.

Virtues... we acquire just as we acquire crafts, by having first activated them. For we learn a craft by producing the same product that we must produce when we have learned it; we become builders, for instance, by building, and we become harpists by playing the harp.

(Aristotle 1999, Book II, Chapter 2, §3, 1103a)

Aristotle is clear that we can ϕ without knowing how to ϕ , at least when ϕ -ing is either a craft (ie a skilled act-type) or a virtuous action. This is what is involved in learning to ϕ — we begin by ϕ -ing, without knowing what it is to ϕ . We might want to describe these as something other than acts of ϕ -ing, perhaps as (mere) attempts-to- ϕ or failed- ϕ -ings, at least initially. As we progress in our learning, we will be tending towards the successful ϕ -ing, even though the first time we do it correctly we might not know how it is that we did this. We saw this in the infant's visually-guided reaching, with them still working out how to achieve their ends after they had already done so. Cases like this abound outside of infancy and early childhood. Before I could knit, it was something I wanted to be able to do, and I learnt in the only way that one can, by knitting. Almost every first time knitter begins with some DK yarn and 4mm needles (because these are the most readily available supplies) and embarks on knitting

what can loosely be described as a scarf. The early rows are full of uneven tension and dropped stitches, sometimes they bring the yarn the wrong way around the needle to make an ill-formed stitch. As their rectangle grows, however, the rows get better and better, and by the end of this, they are able to knit, albeit at a beginner level.⁶³ Cases like this appear across the board — we might decide to try a new sport, to learn to drive a car, to experiment with cooking a new kind of dish. All the time, we intentionally do things without knowing how to do them, and whilst acting with an intention to do them.

Consider, for instance, someone arriving at their first driving lesson. They may have little understanding of what driving involves. Nonetheless, if they have expressed a desire to learn to drive a car, have seen cars being driven and known that people are in control of them, and explained their desire in virtue of, for instance, wanting to be able to get about more easily, or travel long distances without needing to get on public transport, we would say that this is something they want to do. If they had not only done this, but had obtained their provisional licence, booked themselves an intensive course of lessons, and registered for their theory test, we would say that they intend to learn to drive a car. And yet, they do not know how to drive a car. They will learn how to do it by doing it. For many act-types, learning to ϕ requires ϕ -ing, and thus in intending to learn to ϕ , I intend to ϕ (it's just that I expect to initially ϕ badly). This might seem a kind of merely intentionally ϕ -ing, if it were not for the relationship between ϕ -ing and learning-to- ϕ . These are not unrelated act-types. Learning to ϕ is rarely something one does for the sake of it, but rather in order to be able to ϕ , once one has learnt. Although we might, in some situations, terminate our explanation of what we are doing at 'learning to ϕ ', we in fact have a variation on the structure I set out earlier. When I learn to ϕ by ϕ -ing, because I am acting with an intention to ϕ (once I have learnt how), then I am ϕ -ing because I am acting with an intention to ϕ . Not only do young children know enough about the act-types they learn to, in principle, be able to intend to engage in them, there are clear cases of young children learning to do things that are of act-types that we may well consider necessarily intentional. Reading a book is clearly something that can only be done intentionally — perhaps my eyes glance across a sign and I take in what it says unintentionally, but a book is something that one opens only to read. The young child learns to do this by doing it, initially in a stumbling way, but by doing it nonetheless. In learning to read, and other act-types that are by their nature intentional, it seems especially hard to deny that the young child acts intentionally. We should think of the child that does not know how to ϕ , and who is

⁶³ The transition from this to skilled action is discussed in depth by Piñeros Glasscock (2021) who dubs it 'Novicehood-to-Mastery'

learning to ϕ , as having sufficient understanding of ϕ -ing to act with an intention to ϕ . This does not, I should reiterate, entail that they do in fact intend to ϕ , rather than want to ϕ , for instance, but it remains possible that they intend. In the next section, I will set out the kind of knowledge that we need to have of act-types in order to intend to engage in them.

6.3.4 Knowing what it is to ϕ

Over the last few sections, I have argued that acting with an intention to ϕ does not require that an agent knows how to ϕ . An agent must, however, have some kind of knowledge of ϕ -ing in order to intend to ϕ . This is a more minimal kind of knowledge which I will call ‘knowing what it is to ϕ .’

I can have never travelled to the British Library before, but know what it is to do that. I can understand that it is of a kind with travelling to other libraries, and to other places in the surrounding area, can know what it is to plan a route, to take a train, and what I might need to do to ensure my travel is successful. Perhaps I have even ‘travelled’ the route from the station on Google Streetview, seeing what I would pass, where I would need to turn. That is not to say that I know *exactly what it would be like* to go there, because I have never done it, but that I know *what it is* to go there. I can do something for the first time and also know what it is to do that before doing it. I know what it is to travel to the British Library, even if I have never done it. To know what it is to ϕ is to know enough to act with an intention to ϕ . It consists of two key elements:

- knowing what it is to ϕ is sufficient for knowing that one is ϕ -ing when one is ϕ -ing
- knowing what it is to ϕ entails knowing enough about ϕ -ing to identify which kinds of ϕ -ing might be ways of ϕ -ing

The first of these is the more straightforward of the pair. Knowing what it is to ϕ is having sufficient understanding of the kind of thing that ϕ -ing is to recognize oneself as ϕ -ing. It is not practical knowledge, which is knowledge of what one is doing, but it is the knowledge that enables one to have practical knowledge. This does not mean that I need to ϕ to know what it is to ϕ , but rather that if I were ϕ -ing, I would be able to recognize it. The second of these is closely related to this, since an agent’s ability to recognize themselves as ϕ -ing relies on their knowing that the things that they are doing are ways of ϕ -ing. If I thought that going to the British Library involved cracking eggs in my kitchen, I would not know what it is to go to the British Library, because I would be incorrect about

the kind of act-type it is. If I started going the wrong way, because I do not know exactly where the British Library is, I clearly don't know how to go to the British Library, but I do know what it is to go to a library, I'm just doing it wrong. If I get to the library, I have done this intentionally and with an intention to do it. What is involved in knowing the kinds of ϕ -ing that are potential ways of ϕ -ing is dependent on the context and on the abilities of the agent in question. If ϕ -ing is a skilled action that is unknown to the agent, it might be enough that they know how to find a way of learning to ϕ . If it is travelling to a new destination, it may well involve planning routes and considering methods of transport. It does not require that an agent knows what they will do in ϕ -ing, just that they know enough to "get going". One way an agent might get going with ϕ -ing is by asking someone to show them how to ϕ . If we think back to Rödl's carving child, this is what they do. They want to do *that thing*, and they know that they can do that by turning to a parent for help. In Chapter 7, I will explore this in more detail, looking at the way that dependency supports the intentional nature of the young child's action, but at this point I just want to note that it seems coherent that this is a kind of knowing what it is to ϕ . I can be acting with an intention to get to the library if I plan on asking a friend about the route, and it does not seem to prohibit the young child from intending to ϕ that they are going to ask an adult how to do it.

6.4 Joint intentional action

Thus far, I have argued that lack of know-how is not a barrier to individual intentional action, and thus young children can in principle intentionally engage in novel-act types. I will now turn to consider whether the agency of young children prohibits them from engaging in joint intentional actions. In particular, I will discuss whether young children are capable of understanding and relating to other agents in the way needed to jointly intentionally act with them. I will set out why this kind of understanding of another agent is necessary for joint intentional action, before turning to empirical research. I will first consider research that suggests that young children are not capable of understanding others in the requisite way, before turning to alternative interpretations of this research. I will then set out several ways of modelling the understanding that young children have of other agents, such that they could, in principle, engage in joint intentional action. I will also consider the implications of this research for questions about individual intentional agency in early childhood. Most of the empirical work into intentions and intentional agency in early childhood is of the kind I will consider in the following sections, investigating how young children describe and understand the action of others, rather than the kinds of actions they engage in individually. This is, I take it, in part

due to the difficulty of finding ways of assessing whether a particular action is intentional, especially when its agent has limited spoken language.

Although I will be considering the ability to engage in joint intentional action, this is valuable in assessing their capacity for individual intentional action. The relationship between the capacity for individual and joint intentional action is unclear, but these capacities are either developed together, or one after the other. If the ability to engage in joint intentional action requires first being able to engage in individual intentional action, then young children engaging in joint intentional action is evidence that they must already be able to individually intentionally act. If the two develop at the same time, then evidence of one is evidence of the other, and the child that can jointly intentionally act can also individually intentionally act. The final possibility is the most complicated — it may be that the ability to engage in joint intentional action arrives before the ability to individually intentionally act. If this is the case, then evidence of joint intentional action is not evidence of individual intentional action. However, if a child can recognize another as an intentional agent in the way that joint intentional action involves, this does open the door to a different possibility — that they can engage in dependent intentional action. If they can recognize a parent or caregiver, for instance, as helping them, then they can, in principle, dependently intentionally act. I will explore dependent intentional action in childhood in Chapter 7. First, I want to set out the kind of understanding that one must have of other agents in order to jointly act with them, and consider whether young children can have this kind of understanding.

6.4.1 Understanding others as agents

In §2.3, I set out a general picture of joint intentional action, arguing that it is compatible with the dominant accounts. The picture I gave is as follows

Joint Intentional Action

An event X is a joint intentional action of agents A_1 - A_n if

- (J1) X is an event of A_1 - A_n 's ϕ -ing
- (J2) A_1 - A_n 's ϕ -ing consists of each of A_1 - A_n acting intentionally such that A_i is ϕ_i -ing
- (J3) A_i is ϕ_i -ing because their ϕ_i -ing contributes to realizing their intention to ϕ together
- (J4) A_1 - A_n are ϕ -ing to realize their intention to ϕ together

On this picture, joint intentional action involves agents acting with an intention to ϕ together. Acting with an intention to ϕ together with other agents involves understanding what one is doing as involving those other agents. This is part of the difference between joint intentional action and other kinds of joint activity. If you and I are both acting with an intention to bake a cake together, realizing that intention involves all of us directing ourselves at the baking of a cake. This does not require that we all do the same things, but it does require that we are acting in service of our shared goal of baking a cake. More than this, though, it requires understanding the other agents involved as acting with such an intention. I can only intend to bake a cake with you if I understand you as at least able to intend to bake a cake with me. If I am acting with an intention to bake a cake together, then I need to understand you as acting with this intention. If you are not acting with this intention, then we cannot bake a cake together, but your actually holding this intention is not a requirement for my acting with it. I do, however, need to believe that you are acting with such an intention (even if I am mistaken). This follows from the previous sections of this chapter, in which I argued that acting with an intention to ϕ requires having sufficient understanding of ϕ -ing to identify ways of ϕ -ing. If I do not believe that you are acting with an intention to ϕ together, then there are no ways of realizing an intention to ϕ together available to me, since any kind of ϕ -ing together involves you.

This claim might be too demanding — it might be that I could intend to coerce you to ϕ together with me, as a way of realizing an intention that I currently hold to ϕ together. However, even if I step back from the requirement that I understand you as acting with an intention to ϕ together, I need to understand you as capable of acting with an intention. I can only engage in intentional action with you if you are capable of intentional action, and so I can only intend to ϕ together with you, given that this will require you to ϕ together with me, if I understand you as an agent capable of acting intentionally. Our realizing that intention will require you to direct yourself at the realization of this intention in the same kind of way that I do. If I do not think you have the requisite kind of agency to do this, then I cannot engage in joint intentional action with you. Doing this, then, requires me to understand others as intentional agents. Over the following sections, I will consider whether young children can understand others in this way.

6.4.2 Young children's understanding of other agents

Before considering what this research shows, I want to say a little about a background understanding of agency that is present in much work in empirical cognitive science and developmental psychology. The 'desire-belief theory of intention' (Bratman 1987) understands intentions as arising out of the conjunction of a belief and a desire. It is captured in the following claim from Davidson:

R is a primary reason why an agent performed the action A under the description d only if R consists of a pro attitude of the agent towards actions with a certain property, and a belief of the agent that A, under the description d, has that property.

(Davidson 2001a, p5)

Remembering that for Davidson, the primary reason for an action is an intention, this is a model of intentions that takes forming both beliefs and desires as necessary for forming intentions. This does not mean that beliefs and desires are sufficient for intentions (since it might be possible for an agent to form both of these without possessing intentional agency), but it does take that they are necessary. We can see this understanding in the way that empirical research into the mental states of young children is carried out. In particular, we see the idea that understanding of both beliefs and desires is necessary for understanding intentions. Wellman and Wooley (1990) propose that children at the lower end of the age-range with which I am interested do not understand action in terms of beliefs and desires. Instead, older 2-year olds understand actions only in terms of 'simple desires'. They found that even some 3-year olds would only offer desire-explanations of actions, even when prompted to give belief-explanations (ibid, p273). For the belief-desire theorist, this would seem to entail that these children are unable to understand others as intentional agents.

In the previous paragraph, I indicated one of the kinds of experimental data that is used to identify the development of the ability to understand others as intentional — *the kinds of explanations agents give of others' actions*. If a child does not seem able to explain another's actions in terms of both beliefs and desires, given the commitment to belief-desire psychology, then there is no reason to think them able to understand another as an intentional agent. Alongside this, there is a focus on false belief tests as markers of Theory of Mind. Theory of Mind is the ability to understand others as having mental

states distinct from one's own. This does not necessarily require that an agent with Theory of Mind explicitly grasps such a theory, rather than that such a theory provides an explanation of their abilities. Adjudicating between approaches to theory of mind is a project of its own, my purpose in referring to it here is just to identify the difficulty that false belief failure is supposed to point to, a difficulty with recognizing the mental states of others. In the classic false belief test, a child is asked what some particular agent will do next in a given scenario:

Suppose Maxi's mother transferred the chocolate Maxi put into the green kitchen cupboard to the blue cupboard while he is out playing. Maxi, feeling peckish, returns to the house. *Where will he look for the chocolate?*

(Perner and Roessler 2010, p200).

Passing this test and stating that Maxi will look in the green cupboard is taken as evidence that the child has Theory of Mind, since this answer relies on attributing to Maxi an action that is motivated by a false belief. Children do not start reliably passing these tests until 4-5 years of age⁶⁴, which is taken to show that before this age, they do not understand the 'explanatory role of beliefs' (ibid) in action. Belief-desire psychology suggests that understanding an agent as intentional requires understanding them as acting based on what they believe will enable them to satisfy their desires. If failure to pass false belief tests indicates a failure to attribute false beliefs to other agents in situations in which they would hold them, then the role, nature, and formation of beliefs is not properly understood. These considerations seem to point to young children, at least those below the age of 4, being unable to fully understand others as intentional agents.

However, not all research supports this picture. If instead of looking at how young children explain others' actions or the mental states we attribute to them, we look at what they do, they seem to understand others as intentional. There are several studies suggesting that by the age of 2, young children appear able to identify and respond to the intentions of others. Meltzoff (1995) found that 18-month olds seemed to understand what adults were intending and failing to do, carrying out the intended action in these instances with the same frequency as when they see the adult succeed at this act-type. This seems to show understanding of others as acting intentionally. Carpenter et al (1998) considered an even younger cohort, of 14-18 month-olds,⁶⁵ and found that when adults performed

⁶⁴ See (Wellman, Cross, and Watson 2001) for a meta-analysis of 178 studies

⁶⁵ This seems to be the lower age limit for this kind of behaviour, with Bellagamba and Tomasello (1999) finding that 12-month olds did not replicate intended, but unsuccessful, actions with any real frequency.

two actions, one vocally marked as intentional (“There!”) and one vocally marked as accidental (“Whoops!”), the children replicated the intended action with twice the frequency of the accidental action. This again speaks to an understanding of actions as intended by agents. Somewhere else this comes out clearly is in experimental work looking at the development of altruism. ‘Altruism’ is not understood in this literature as the kind of substantive attitude philosophers might refer to by this term, but rather as a thinner notion that’s fairly interchangeable with ‘helping’. It might be that this helping is altruistic, that is it might be being done selflessly, or at the expense of the altruistic agent, but this is not something that is measured. What it is, however, is helping with something that the altruistic agent is not themselves doing, and in that sense it is selfless — it is not self-directed or immediately self-benefitting.⁶⁶ Warneken and Tomasello (2007) found 14-month olds to altruistically help others in achieving simple goals. In tasks where the experimenter appeared to accidentally drop an item, 14-month olds ‘reliably handed out-of-reach objects’ (ibid, p278) to the experimenter. They note that ‘the majority of children helped and did so across several trials...’ (ibid) suggesting that this helping is not merely a one-off occurrence. In most (71%) of these cases, helping happened in the first 10 seconds after the object was dropped, during which the experimenter was merely looking at the object, although sometimes it did not occur until the experimenter looked at the child (20%) or later referring to the object (but not the child) (9%) (ibid). No matter which stage the helping occurred at, it speaks to an understanding of what the experimenter wanted to do — the difference is in terms of the spontaneity of the child’s helping, not obviously in terms of their understanding of the adult as intending to act. Moreover, we are still considering children much younger than my cohort, and children even a few months older than this were found to help with more complex tasks. When Warneken and Tomasello carried out these and other tests with 18-month olds (2006), they found that these children helped with book-stacking, and opening cupboards that experimenters were unable to, as well as retrieving desired objects from a box by a method unavailable to the experimenter, but not in control conditions in which the experimenter did not appear to be accidentally struggling with a task.

⁶⁶ It seems worth commenting on the way that young children are often settled into the experimental setting, which is via a reciprocal play condition. This seems to ‘prime’ the children for the kind of altruistic behaviour that is being tested for (see (Cortes Barragan and Dweck 2014) for a discussion of this effect). Although this kind of play might bring altruistic behaviour ‘to the fore’ for the young child, it can’t generate a tendency to engage in it. What it might do is reduce our reasons for thinking of this as altruistic, since they have recently engaged in a kind of helping behaviour that benefitted them. For my purposes though, this is not a problem — it does not undermine the claim that the child understands what the adult is intending to do, nor the claim that they understand helping, even if their helping is less ‘altruistic’ than we might otherwise think.

This empirical evidence is obviously not enough alone to show that young children can, in fact, understand others as intentional agents, but it does challenge the authority of the research I cited earlier in this section. As I noted in §6.3.1, we cannot assume that the limits of a child's communicative capacities map onto the limits of their cognitive capacities. Given this, this second strand of research, which looks at what children can do rather than what they can tell us, also seems a better way of judging the understanding that young children have of others, even if they do not verbally communicate that understanding.

6.4.3 Intention, desire and belief

Moreover, all of this has operated on an assumption of the order in which children come to acquire concepts. Since these kinds of approaches think of being intentional as something that is added on to desire-belief psychology, until the child has an understanding of desires and beliefs, they cannot have an understanding of intentions (and if they do not have desires and beliefs, they cannot have intentions). As Tollefsen notes, '[i]f intention is essentially tied to the notion of belief, then it would seem that an understanding of intention presupposes an understanding of belief' (Tollefsen 2005, p88 fn.20). This is reflected in the tendency to look to passing the false-belief test as the first marker of theory of mind. However, it's not clear that the assumption that we must understand desire and belief before we can understand intention is correct.

Lillard and Flavell (1992) compared the success of 3-year olds in a 'false-desire' scenario (one in which another agent wanted something to be the case that the child knew not to be the case) and an otherwise identical false-belief scenario. The results 'imply that children can understand that the contents of desire states can be different from reality earlier than children understand this about belief states' (ibid, p629). Wellman and Wooley (1990) argue that, at least in the context of understanding others actions, what precedes a desire-belief psychology is a desire-psychology, not a belief one. Both of these papers disrupt the idea that understanding of belief is more basic than understanding of desire. Of course, desire and intention are not the same. However, there is also evidence that intention might come before belief. Gopnik and Slaughter (1991) found that 3-year olds had more success in reporting their past desires and intentions than in reporting their past beliefs, suggesting that they had a better handle on conative states than cognitive ones. Moses found that 3-year olds 'understanding of unfulfilled intentions was excellent, and significantly better than their understanding of false beliefs' (1993, p1). What is particularly striking is that these children did better in reporting others false-beliefs

in the context of intention than on straightforward false-belief tests, with Moses remarking that ‘many 3 year-olds appeared to understand that, when an actor intends to achieve an outcome but fails to do so, then the actor mistakenly thought he or she would achieve that outcome.’ (ibid, p21) It is possible that this speaks to the child having a stronger grasp on belief than the false-belief test suggests, but it is also possible that it indicates that the development of the understanding of belief is in part something that grows out of the understanding of intention. Either way, it seems to count against the idea that a fully-fledged understanding of belief must come before any understanding of intention. If this is the case, then this disrupts the idea that we should wait for success in the false belief-test before we take someone to understand others as capable of intentional action.

In this section, I have considered whether young children are able to understand others as intentional agents, which I have argued is necessary for their engaging in joint intentional action. Having considered empirical evidence that speaks to lacking this understanding, alongside evidence that complicates this claim, I will turn in the next section to what I take to be the key consideration here. Does young children’s failure of false-belief tests entail that they cannot understand others as intentional agents? Their ability to so understand others is crucial to their being able to act jointly with others, and thus to the overarching question of the second part of this chapter — can young children engage in joint intentional action?

6.5 False-belief tests and theory of mind

In the previous section, I indicated that there is some evidence that challenges the idea that failing false-belief tests indicates an absence of theory of mind. I will now expand on alternative approaches to false-belief tests and to theory of mind in early childhood. I will begin by considering what young children’s failure of false-belief tests actually indicates, detailing two ways that young children might understand others actions which are compatible with their seeing others as intentional agents even whilst failing false belief tests. Then I will consider an account that posits a more minimal theory of mind, one that does not rest on someone representing the representational states of others, and which is potentially attributable to young children. These accounts, I suggest, give us reason to think young children are able to understand others as intentional agents, despite failing false-belief tests.

As Perner and Roessler note, if you accept belief-desire psychology, you assume that ‘to see people as acting intentionally, you have to think of them as acting on the basis of what they believe:

understanding intentional action requires a grasp of the explanatory role of beliefs.’ (2010, pp199-200) The research I discussed in the beginning of §6.4.2 shows that children of 2-3 cannot offer such explanations, but the research I discussed in the latter part of the same section suggested that children of this age can understand others as acting intentionally. In response to this tension, Perner and Roessler turn to the false belief tests, offering an alternative explanation of what these show.⁶⁷ Recall their example of a false-belief test, which I gave in that section

Suppose Maxi's mother transferred the chocolate Maxi put into the green kitchen cupboard to the blue cupboard while he is out playing. Maxi, feeling peckish, returns to the house. *Where will he look for the chocolate?*

(Perner and Roessler 2010, p200).

To pass this test, a child must state that Maxi will look for chocolate in the green cupboard, since this is taken to be evidence that the child understands that Maxi will falsely believe that the chocolate is still there. Perner and Roessler are interested in a particular kind of failure of this test. Three-year olds do badly at this task, but in a non-random way. ‘They reliably predict that Maxi will look in the blue cupboard. They don't suggest that he will look under the kitchen table or in the playground or in the loft.’ (ibid) We ordinarily interpret this as the child giving an answer in terms of what it makes sense for Maxi to do, given that he wants the chocolate and this is where it is located. Thus, as Perner and Roessler point out, on this interpretation ‘children do think of Maxi as an intentional agent. They assume that Maxi acts in a purposive manner; that he will do what, given his purpose, he has good reason to do (ibid, p201). Therein, they contend, lies evidence that children who fail false-belief tests in the standard way *are* capable of understanding intentional action. Young children’s response to the test is most naturally understood as their being agents with a handle on what is is to act intentionally. This, combined with their apparent failure to understand the role that beliefs are taken to play in intentional action, creates an apparent tension for belief-desire psychology. Over the next sections I will consider several alternative accounts of the understanding young children have of the action of others, first those suggested by Perner and Roessler (2010) and then one given by Butterfil and Apperly (2013). These approaches all present young children as able to understand others as acting in goal-directed ways, even perhaps acting intentionally, without needing to attribute false-beliefs to others.

⁶⁷ They are not alone in reconsidering whether apparent failure in false belief tests really indicates an absence of understanding of others as holding false beliefs, with Onishi and Baillargeon (2005) suggesting that children as young as 15 months can pass non-verbal false-belief tests.

6.5.1 Teleological and hybrid accounts

Perner and Roessler present two, related, accounts of the understanding that young children have of others. Young children's tendency to fail false belief tests by voicing what the other agent should do, not what they will do is, as I have noted, something they take as evidence that young children do understand others as intentional agents. The difficulty stems from taking this understanding to rely on belief-desire psychology, which would require young children to understand beliefs as playing an explanatory role. This would require them to attribute to Maxi a true belief that he cannot hold, which would then indicate that they do not understand how we acquire beliefs, and thus lack adequate understanding of beliefs.

Perner and Roessler argue, then, that young children's understanding of intentional action must rest on something other than a belief-desire model of intention. Their proposals rest on a distinction between *objective* and *subjective* reasons, which they characterize by drawing on Williams's (1981) would-be gin and tonic drinker, who mistakenly believes that the glass of petrol before him contains gin. On Perner and Roessler's terminology, the agent in question has a subjective reason to add tonic to the glass and drink it, but they do not have an objective reason to do this.

The distinction turns on whether or not a reason statement is to be understood as relativized to the agent's current perspective. To say that you have a subjective reason to drink the stuff is to say that, from your perspective, it looks as if you have an (objective) reason to do so.

(Perner and Roessler 2010, p204)

The agent, they suggest, has a (fully) objective reason not to drink this liquid, since it is petrol and thus non-potable. He would have a (partially) objective reason to drink it if it were gin, since he desires a gin and tonic— he has an objective reason, relative to his desires and instrumental beliefs.

Perner and Roessler offer two potential accounts of children's understanding of others as intentional that rest on the distinction between objective and subjective reasons. The *hybrid account* suggests that

young children find actions intelligible in terms of partially objective reasons: reasons that are relativized to the agent's desires but not to her instrumental beliefs... children

might take it that given Maxi's desire to eat some chocolate, he has a good reason to look in the blue cupboard (for doing so will as a matter of fact enable him to satisfy his desire.)

(Perner and Roessler 2010, p204)

On this picture, young children 'conceive of Maxi's reason as a combination of a desire and an objective instrumental fact' (ibid). To understand another as acting for a reason is to grasp what it is an agent wants to do and what in fact would enable them to do this, without needing to understand the agent as holding beliefs about that fact. That is, understanding them involves combining 'a mental state defining the objective of the action, and an objective fact, determining the means to achieve the objective.' (ibid, p205).

Alongside their hybrid account, Perner and Roessler discuss a *teleological account*, which

maintains that children find actions intelligible in terms of fully objective reasons, relativised neither to the agent's instrumental beliefs nor to her pro-attitudes. They conceive of Maxi's reason in terms such as these: 'Maxi needs his chocolate. (Or: It is important, or desirable, that Maxi obtain his chocolate.) The way to get it is to look in the blue cupboard. So he should look in the blue cupboard.'

(Perner and Roessler 2010, p205)

What is distinctive about this kind of model is it does not involve the attribution of mental states to Maxi. The explanation it posits the child offering is not in terms of Maxi's subjective beliefs or desires. It depicts intention not as 'a state that provides the agent with a subjective reason' (ibid, p213). Rather, that 'an agent has this intention follows from the fact that his action is open to a reason-giving, teleological explanation' (ibid).⁶⁸ This, of course, leaves children with an understanding of the actions of others that will fail at times, both in contexts like Maxi's where the agent is acting on a false belief, and in ones where their values differ from those of the person whose action they are attempting to understand. If they do not view the obtaining of chocolate as desirable, then, they will not understand someone's action as being directed at obtaining chocolate.

Perner and Roessler consider research that speaks in favour of these two accounts, although they consider that the evidence that might be rallied in favour of the hybrid account does not in fact speak

⁶⁸ This echoes my discussion in §6.3.1

to young children attributing mental states to others. Although the hybrid account removes the need to understand others as having subjective reasons or false beliefs, it still assumes the attribution of desires to others. The evidence that is taken to speak to this ability, they suggest, might really only support the teleological approach. For instance, they consider Repacholi's and Gopnik's (1997) finding⁶⁹ that most 18-month-olds in their study seemed to recognize an experimenter's preference for a foodstuff (generally broccoli) that they themselves did not like and offered it in response to a request for food instead of the one that they themselves preferred. Perner and Roessler suggested that this does not speak to the child attributing a subjective desire to the experimenter. Instead, children may well understand these kinds of preferences in terms of '*objective* purposes' (Perner and Roessler 2010, p212) rather than subjective enjoyment. Drawing on Perner et al (2005, p239) which proposes that 'the subjective nature can be captured within objectively good person-food combinations (broccoli in her mouth is good) and objectively bad ones (broccoli in my mouth is bad)', they suggest that for the child 'the import of A's and B's differential preferences is that it is objectively desirable that A, but not B, should obtain some broccoli.' (2010, p212)

As well as this offering us a way of accounting for children as understanding intentions, this approach comes from a place of taking the appearance of children's behaviour seriously. The starting point is that children are acting as if they understand intentions — why else would they give the answers that they do in the false belief task? Given that the most coherent and natural explanation is that they understand intentional actions, this a guiding approach that we should endorse. In the next section, I will consider another approach to young children's understanding of others as acting intentionally.

6.5.2 Minimal theory of mind

Part of the cognitive hurdle that is posed by understanding others as acting on beliefs and desires is that this involves representing representational states and attitudes. Beliefs and desires are representational, in that they represent the world being some way, and understanding others beliefs and desires as representations of this kind involve representing these representations. This lay in the background of Perner and Roessler's approach, and in this section I want to turn to an account that focusses on this issue, considering how theory of mind might be possible without representing representational states.

⁶⁹ They also consider Bartsch's and Wellman's (1995) finding that '2-year-old freely verbalise subjective preferences' (Perner and Roessler 2010, p211) through a similar lens.

Butterfill and Apperly (2013) construct their minimal theory of mind to explain ‘what enables those with limited cognitive resources or little conceptual sophistication, such as infants, chimpanzees, scrub-jays and human adults under load, to track others’ perceptions, knowledge states and beliefs.’ (ibid, p1). Although their focus is wider than mine, considering also non-human animals, they share my interest in understanding the action of young children, and also argue that we can attribute to them more understanding than we might inclined to otherwise. In Apperly and Butterfill (2009), they draw on the contrasting evidence that 2-3 year olds fail ‘critical tests of belief reasoning’ (ibid, p953) and yet 1-2 year olds appear to pass false-belief tests. The evidence for this, again, comes from observing the behaviour of these children, and recognizing that this is the most coherent explanation of their behaviour — in this case, by monitoring looking-times in various scenarios, such as in Buttelman et al (2015). That the child looks longest in some cases is understood as the child’s expectations being violated because they were expecting the experimenter to act on a false belief they understood them to have.

Theory of mind allows agents to track things like the beliefs and desires of other agents. Tracking these states enables an agent to understand and predict the behaviour of those whose states they track. If I can track Maxi’s belief about the location of the chocolate, then I can correctly predict that he will look for it in the wrong place. In standard theory of mind, this tracking is representational. On this picture, then, being able to track others’ beliefs, and thus attribute false beliefs to them, requires representing others’ propositional attitudes. Even if, as I have been assuming, intentions are act-directed rather than having propositional content, this still involves their having a kind of representative content. Standard theory of mind requires the representation of representations, which is cognitively demanding, and provides a significant barrier to understanding others as capable of intentional action. Instead of this, Butterfill and Apperly propose a list of principles, none of which require the representation of representation, and which build on the previous principles in ‘such a way that it would be coherent to suppose that [an agent] has the abilities codified by the first n principles only’ (2013, p8), but do not assume that this series of principles represents an evolutionary trajectory. They build these principles out of notions that do not involve representations of representations:

- *goals* – ‘for an outcome, g , to be the goal of some bodily movements is for these bodily movements to occur in order to bring about g ; that is, g is the function of this collection [of movements]. Here ‘function’ should be understood teleologically.’ (ibid, p9)

- *an agent's field* – this is a set of objects, and can be approximated as those objects that an agent can perceive at that time
- *encountering* – this is a relation between an agent and an object, and an agent is encountering an object if it is in their field. It is not a representation of the object that is encountered, it is the relation between the agent and the object
- *registration* – this is a relation between an agent, an object, and a location, that obtains if the agent most recently encountered the object in that location

From these, they produce the following principles

- (1) '[B]odily movements form units which are directed to goals.' (Butterfill and Apperly 2013, p10). Goals are not intentions, but outcomes at which bodily movements are directed, represented as a function of bodily movements
- (2) '[I]f an outcome involves a particular object and the agent has not encountered that object, then that outcome cannot be a goal of her actions.' (ibid, p11)
- (3) An agent registers an object at a location iff they most recently encountered it at that location, and correct registration is a condition for successful action. A registration is correct if the object is in this location
- (4) '[W]hen an agent performs a goal-directed action with a goal that specifies a particular object, the agent will act as if the object were in the location she registers it in.' (ibid, p16)

These principles take it that those with minimal theory of mind have representational states, but they do not require an agent to represent another's representations. It is a kind of theory of mind that an agent could possess without being able to represent others representational states. If an agent can do what is captured in these principles, they could understand another as engaging in goal-directed action. Their suggestion is that this more minimal tracking system constitutes a kind of theory of mind possessed by, amongst others, young children, and continues to exist once we develop a secondary, more sophisticated, tracking system.⁷⁰

There are two central things to be taken from the teleological and hybrid approaches and minimal theory of mind. Firstly, that we should take children's apparent understanding of intention as evidence of their understanding of intention. If the best and most natural explanation we can offer of what

⁷⁰ Perner and Roessler (2010, p218) similarly suggest that mature understanding of the actions of others may have its roots in the teleological understanding they attribute to children.

they are doing is that they understand others as having capacities that we do not ordinarily think they could attribute to others based on our theory, we should consider that they are doing exactly this and that the theory is wrong. Secondly, they show that there are feasible explanations of how children might be able to understand others as having the kind of mental states needed for intentional action, without needing to attribute to children the kind of cognitive sophistication involved in understanding others as having representational mental states like beliefs and intentions. The teleological approach and the minimal theory provide a solution to the belief-desire theorist's struggle with the idea of young children as understanding intention. If we adopt Apperly's and Butterfill's idea of two-systems, we can see this early system as prior to and distinct from full-blooded understanding of intentions.

6.6 Joint action in early childhood

Over the last few sections, I have considered whether young children can understand others in the way necessary for joint intentional action. I have argued that there is at least good reason to think that young children are able to understand others in the requisite way. I now want to consider this more directly. Rather than looking for evidence of the kind of understanding necessary for joint intentional action, I will look for evidence of joint intentional action itself.

There is evidence that infants begin to engage in patterns of joint attention from a very early age — Scaife and Bruner (1975) report it in infants as young as 2 months, and Butterworth and Cochran (1980) confirm that it has begun to develop by 6 months of age. Joint attention, that is, attending to that which another agent is looking at with them, involves recognizing the other's looking as purposive. Attending to what another attends to relies on recognizing them as so attending rather than merely accidentally gazing in that direction. Joint attention is generally understood to play a role in the acquisition of language (eg Bakeman and Adamson (1986), Tomasello and Farrar (1986). Southgate, both individually (2020) and in co-authored work (Grosse Wiesmann and Southgate 2021), argues that infants are alter-centric, contra a more conventional understanding of them as first only recognizing the existence of their own perspective.⁷¹ Given this, if the development of the ability to act intentionally rests on being able to understand others as being at least purposive, there is reason to think that this develops very early.

⁷¹Richmond (2000), for instance, discusses this understanding

There is also evidence of early engagement in joint activity, including joint action. Warneken et al (2006) build on existing research into young children's involvement in various kinds of joint activity. For instance, Eckerman (1993), observed dyads of same-age toddlers with a particular interest in 'acts thematically coordinated to the peer's actions that allowed the peer to continue in his or her action while expanding the action to include both children.' (ibid, pp334-5) Imitative acts, in particular, manifested frequently in the behaviour of the toddlers, occurring 'at the rate of two or three times per minute' (ibid, p335) amongst toddlers of 28 months and older. As the dyad grew older, this pattern emerged, at a median age of between '20 and 24 months' (ibid, p336). This shows young children engaging in social action with other agents. Engaging in this kind of imitative behaviour indicates a kind of acting with another that involves a sense of what it is the other is setting out to do, and how one might do that. The examples that Eckerman gives include 'hiding but at a different location, throwing a beanbag instead of a ball... marching around in a circle on the box before jumping off, or jumping off and falling in a nonliteral manner instead of simply jumping off' (ibid, p335). These are things that involve understanding what it is the other is doing and taking on this action in their own behaviour. This is an understanding acquired in the context of being with another, and is directed at a purpose that they have access to only in virtue of acting alongside the other. It might seem that this is a kind of parallel action, rather than cooperative, but Warneken et al incorporate a second kind of activity into their consideration, namely cooperative problem solving:

Cooperative problem-solving tasks are by definition novel for children, and they have a clear goal that can only be achieved through coordinated action (nothing useful or interesting can be accomplished individually). The coordinated action also needs to take place basically simultaneously and the roles are typically different; therefore, the "imitative pattern" used by younger children is not readily available as a strategy

(Warneken, Chen, and Tomasello 2006, p641)

Brownell and Carriger (1990) showed that '24- and 30-month-olds were able to coordinate behavior with one another quickly and effectively' (ibid, p1164) when presented with cooperative problem-solving tasks, but generally this has been studied in older children. Brownell and Carriger, like Eckerman, are interested in peer behaviour. In that context, the capacities and understanding of the toddlers provide the ceiling on what it is that the dyad can achieve.

Warneken et al (2006), however, paired 18-24 month old children with adults to engage in both social and problem social activities. They are also interested in the phylogeny of these skills, and so carried out the same experiment with chimpanzees for comparative purposes. On their methodology, ‘the adult stops playing his role at key moments’ (ibid, p641) as a way to ‘assess shared intentionality.’ (ibid) As they explain, ‘[i]f subjects had formed a shared goal with the other and the partner is not doing his part of the joint activity, they should attempt to reengage him’ (ibid), and so looking for the child’s attempts to re-engage the adult allowed them to determine ‘if the children had indeed formed with the adult a joint goal... and if she understood the different roles involved.’ (ibid) During this age range, the capacity for this kind of joint activity is clearly developing, as ‘[i]n three of four tasks, children at 24 months coordinated their actions with that of the partner more skillfully than the 18-month olds’ (ibid, p650). They conclude:

Our interpretation is that the children in our study were attempting to reengage the adult toward their joint goal. One could claim that when the child is paired with a more skillful adult partner, she might allow the adult to fully regulate the activity while the child executes her action without regard for the partner. Having the adult partner interrupt his activity should help to decide whether this is the case or not.

(Warneken, Chen, and Tomasello 2006, p652)

The children’s attempts to re-engage their partners ‘can be taken as evidence that the children comprehended their own and the partner’s actions as interconnected parts of a joint activity toward a joint goal (joint intentions)’ (ibid, p653). They were not ‘only interested in retrieving the object in problem-solving tasks’ (ibid), since if this were the case, they would try to do this alone, but instead their ‘results indicate that they understood the tasks as involving two roles and were motivated to cooperate with the other and repair breakdowns when they occurred.’ (ibid). Moreover, they concluded that ‘children seem to be motivated not just by the goal but by the cooperation itself’ (ibid, p659) and clearly had a strong sense of the nature of the activity as something with a particular structure, since ‘after only one brief viewing of two adults engaging in the task, children seemed to form a conception of how the game “ought” to be played’ (ibid, p660). Their assessment is as follows:

Children seem to be understanding the social, even normative, structure of the game as defined by the joint goal of the participants and their joint intentions for reaching that goal (even if the goal was simply playing a game together). Moreover, children

were so engaged socially that sometimes they even turned the tasks aimed at retrieving an object into a game; after they retrieved the object, they were not interested in it, but immediately put it back in the apparatus to start the game again.

(Warneken, Chen, and Tomasello 2006, p660)

The children were actively engaged in joint action, seeking to bring about new patterns of joint action, as well as to continue the ongoing ones. Warneken et al state that '[o]n the most generous interpretation, they have learned to form with others a joint goal, which, when it is breached, they attempt to reinstate; they understand that to accomplish the task the two participants must form such a joint goal and joint intentions to effect it' (ibid). The pattern of attempted re-engagement means that their behaviour cannot be explained away as the children complying with the goal of the adult, because this gives no reason that the child would try to reinitiate the action after the adult appeared to give up on the goal.

If, then, as Warneken et al suggest, 18-24 month olds are capable of engaging in this kind of joint action, then at least when they are acting with another, they can act intentionally. They are acting in ways directed at goals, with another acting alongside them. Goal-directed and intentional actions are not one and the same, of course. My cats act with the goal of getting food, but they do not act intentionally. The adult actors in the experiments, however, are clearly engaged in intentional action. They are not only acting intentionally during the experiment, but also acting on a prior intention given its role in the overarching study. Even if we are wary about thinking that children 'understand that to accomplish the task the two participants must form such a joint goal and joint intentions to effect it,' (2006, p660), it seems less controversial to say that they do form such goals and intentions. They recognize the action is not completed when the adult disengages, because they know that there is something at which the action is aimed at that has not yet been achieved. The attempts to re-engage the adult in the activity are different to the continued pawing of my cats at the kitchen door, in that the cats seem to be doing no more than trying to satiate a need or desire, rather than having a goal, understood by them as a goal (or something akin to this), as that at which the action is directed. On this understanding, then, the child is capable of joint intentional action, when they act alongside an adult agent.

I have argued that we should not take young children to lack requisite capacities for intentional action. I considered two central potential barriers. In §6.3, I argued that although children lack know-how of

many act-types, this does not prohibit them from acting with an intention to engage in those act-types. I argued that, rather than know-how, an agent needs a more minimal understanding of ϕ -ing to act with an intention to ϕ . I argued that young children's knowledge of act-types is in keeping with the kind of knowledge that adult agents often have of act-types that they intentionally engage in. I called this knowledge 'knowing what it is to ϕ '. I then turned to joint intentional action, considering the kind of understanding of other agents that is necessary for acting jointly with them — I must understand another as an intentional agent to act with an intention to ϕ together with them. I considered why we might doubt young children able to so understand other agents, before presenting empirical evidence that seems to confound this. I then presented several approaches that allow children to understand young children as having the necessary understanding of others. This leaves it possible that young children are able to engage in both individual and joint intentional action.

I want to sound a note of caution at this point. That empirical research frames something as evidence of intention, or evidence of understanding of intention, does not guarantee that this term is being used as we would use it philosophically. I alluded to this kind of difference in terminology in §6.4.2, when discussing understandings of 'altruism'. In philosophy, 'intention' has a particular meaning, picking out something that features in a kind of agency that involves, for instance thinking about how one will realize the intention, and being committed to this in the face of difficulties and situations that demand changes in plans. Some of the uses of this term in empirical work can mean something more akin to 'goal' or even 'desire'. However, I did not argue that I had shown that young children do in fact engage in intentional action. I argued for the absence of barriers to their engaging in intentional action. It might be that this evidence shows they understand other agents as purposive in a more minimal way, rather than intentional. However, this does not entail that they *cannot* understand other agents as intentional. That they understand them as purposive is, moreover, an important step to understanding them as intentional. In this chapter, I have shown the absence of certain barriers to young children acting intentionally, both individually and jointly, but I do not claim to have shown that they do so act. In the next chapter, I return to dependent action, arguing that dependency facilitates intentional action, even if young children cannot act intentionally without support.

Chapter 7 – Dependent intentional action in early childhood

In the previous chapter, I considered the nature of agency in early childhood. I concluded that there are no insurmountable barriers to young children engaging in individual or joint intentional action. However, I did not argue that they do, in fact, engage in such action. In this chapter, I will argue that young children can engage in dependent intentional action. I will begin in §7.1 by considering an argument that suggests that a supporting agent can enable a child to intend to ϕ , rather than merely desire to ϕ . I will then turn in §7.2 to one of the two notions that I will focus on for the majority of this chapter, that of scaffolding. I will set out the way that this notion is characterized in developmental research, and discuss its pedagogical role. I will set out the second key notion, executive function, in §7.3, discussing its development and its role in intentional action in §7.3.1. I will look at the role of scaffolding in the development of executive function in §7.4, and in §7.4.1 show that alongside helping the child learn executive function, scaffolding can provide exogenous executive function for the child's actions. Finally, in §7.5, I will argue that young children can engage in dependent intentional action with the supporting agent scaffolding their action.

7.1 Supported intention

In the previous chapter, I discussed potential barriers to young children engaging in intentional action. The possibility of these barriers rested on a picture of intentional agency as something that arrives at a certain point in the developmental trajectory of the child. The empirical work I considered relied on a belief-desire picture of psychology and of intention, which tallies with the kind of additive approach I discussed in §1.4.2. Questions about the point at which a child develops the ability to act intentionally can also arise on a transformative approach. Bakhurst (2011) and McDowell (1996) understand education, and thus child development, as a process of becoming autonomous. On this picture, human beings 'are born mere animals, and they are transformed into thinkers and intentional agents in the course of of coming to maturity' (McDowell 1996, p125). In early childhood, then,

young children are not able to engage in intentional action, if we understand education as these approaches do.

Rödl (2016), offers a transformative account of practical reason that nonetheless rejects the view that autonomy is something that can be ‘impressed upon the child by a parent or teacher’ (ibid, p87). He claims that this model of education ‘disrespects the depth to which reason defines the human being’ (ibid) and that the development of autonomy is the child coming into their nature. We should not understand the child as transforming from their first, animal, nature to their second, personal, nature, and only on the second acquiring practical reason. What education does is realize an already present nature, in keeping with the origins of the word in the Latin ‘educere’ or *lead out*. Rödl argues that rationality is part of the nature of persons, claiming:

[A] person is a subject not only of sensory consciousness, but of rational consciousness. A person is a subject of concepts. Thus its self-movement is of a different character, which we indicate by calling it intentional action. Acting intentionally, a person is not only conscious of the individual it desires, but is conscious of itself as acting in pursuit of it; he not only takes the apple, but does so with knowledge that this is what he is doing. His action is as such an application of the concept of his action.

(Rödl 2016, p91)

Although ‘person’ is often used to pick out those I have been calling ‘ordinary adult agents,’ he is clear that he takes this to apply from the beginning of human life, suggesting that ‘[t]he legs of a human newborn are rational legs, personal legs, we may say, even as he moves them randomly during his first weeks of life.’ (ibid).

We can distinguish between latently intentional actions and manifestly intentional actions — that is, those that are intentional in virtue of the nature of the agent and those that are intentional in virtue of the intention with which the agent acts. On this picture, then, young children can engage in latently intentional action just in virtue of the kind of being they are. I did not include this model in Chapter 6, since it sidesteps rather than engages with the worries I was addressing at that point. However, in presenting this view, Rödl also gives us a picture of a young child engaging in manifestly intentional action. He builds out the case I quoted in §6.2.1 as follows:

Consider a child, three years old, let us say, learning to carve a stick. The child wants to carve his stick; however, his intention is indeterminate, as he has no distinct idea of what it is to carve. He knows it involves a knife and somehow applying the knife to the wood. But that is about it. And yet he wants to do it. His intention has a determinate content and thus is capable of governing action as it contains a reference to his parent's consciousness: wanting to carve, he wants to do what she, the parent is doing or does. This consciousness of the general, of the action form, carving, is a shared consciousness; and it is comprehended to be shared by child and parent. And therefore, the parent's conception of how to carve governs the child's action not from the outside, but rather as giving determinate content to the child's own consciousness, which is a shared consciousness. The shared consciousness of the general... the form of which the child supplies from himself, allows a content to figure in the child's will that he cannot supply from himself. In this transaction, or interaction, therefore, the child is not moved by something outside and extraneous, but by himself, by his own rational consciousness. In due course, the child's dependence on the parent for content will diminish. But all throughout the process, it is the child's own will that rules, joined together as it is with the will of the parent in a shared consciousness of the action form.

(Rödl 2016, p91)

Rödl is proposing a picture on which the support offered by the parent includes making it the case that the child intends. The particularity of their relationship, and its role in drawing out the autonomy of the child through education, makes intentional what might have otherwise been merely a desire to ϕ . The child sets out to do *that thing*, and the parent's support furnishes the child with a determinate intention to carve a stick. This is not the child being a member of a community of language users, but the child identifying the particular act of carving a stick, and the parent being present in a way that enables this to be something the child can do intentionally. Without the parent, the child does not know enough to get going with carving the stick, but with the parent there, prepared to support and help the child, they can do this. When they carve, with the support and guidance of their parent, they act intentionally. They do so in a way that is dependent on the parent, since this is the origin of the child's determinate intention. In virtue of acting with the support and engagement of an adult, the child engages in intentional action.

I detail Rödl's account not to argue for the merits of this particular picture of intentional agency or intentional action, but because it points to the ideas that I will be discussing in this chapter. What Rödl presents is a kind of dependently intentional action. In Chapter 4, I presented cases of dependent intentional actions in which an agent realized an intention in virtue of help from a supporting agent. In those cases, the intention with which the supported agent acts stands free of the support — they might intend to realize it with help, as I discussed in §5.4.3, but this does not mean that the help factors into the formation of the intention. The content of the supporting agent's intention is dependent on the intention of the supported agent — the supported agent intends to ϕ and so the supporting agent intends to help them ϕ .

But in the carving case, as well the carving being dependent in the ordinary way, since the adult facilitates and supports the carving, the intention with which the child acts is also dependent. The supported agent's intention to ϕ depends on the supporting agent to give it content. That the child is acting intentionally is something that is possible only dependently, it seems. In the following sections, I will explore the way that support in early childhood can facilitate intentional action that could not otherwise occur, not just because of the limited knowledge of the child, but also because there are aspects of acting intentionally that the child cannot perform unaided.

7.2 Scaffolding

That children can do more when they are supported by another agent is not a novel thought. In the early 20th century, Vygotsky described the 'Zone of Proximal Development' (ZPD) in the following way:

‘The ZPD of the child is the distance between the level of his actual development, determined with the help of independently solved tasks, and the level of possible development, defined with the help of tasks solved by the child under the guidance of adults or in cooperation with more intelligent peers.’

(Vygotsky [1935] 2011, p204)

Reflecting on the ZPD draws our attention to the way that support can extend what a child is able to do. The level of 'actual development' is that which the child can do unaided, but the ZPD indicates

that their capacity outstrips this once they are supported. The notion of the ZPD appears in a few places in Vygotsky's work, where it is in part about understanding the best dynamics for learning, and is in part for understanding what it is the child will soon be able to do alone. He stated that 'what is indicative of the child's intellectual development is not only what he can do himself, but probably more so what he can do with the help of others' (ibid, p203). In part, then, his point is that we should understand the child's development in terms not of individual action, but dependent action. It seems consistent with this that we understand the activity that lies within the ZPD – the things that the child can do supported – as things that the child can do, or at least as things that we should take into consideration in assessing the capacities of the child. Even if in thinking of these supported actions as those of the developing child, I am exceeding what Vygostky intended with this notion, the idea of the ZPD nonetheless points to the way in which support extends that of which the child is capable.

If the ZPD indicates that the child can do more when supported, the notion of scaffolding explains how. Although the latter concept was not developed in explicit engagement with the notion of the ZPD, that the two are complementary is something that has been recognized by those involved in work on scaffolding (Wood and Wood 1996, p5). The idea of scaffolding stems from the work of Bruner (for instance, Bruner 1973), who co-authored the 1976 paper 'The Role of Tutoring in Problem Solving', which gives empirical support to the theory. In this study, Wood, Bruner and Ross observed children aged 3, 4, and 5, as they undertook an unfamiliar task with the support of a 'tutor', an adult experimenter who knew how to complete the task, and who provided 'scaffolding' to the children's attempts.

This scaffolding consists essentially of the adult "controlling" those elements of the task that are initially beyond the learner's capacity, thus permitting him to concentrate upon and complete only those elements that are within his range of competence. The task thus proceeds to a successful conclusion.

(Wood, Bruner, and Ross 1976, p90)

This should sound familiar from my discussions of dependent action. The support provided in scaffolding is not that different from the support provided in ordinary dependent action, but in the context of development — that is, in the scaffolding that goes with the action of young children — there are some important differences.

The task in which the children were supported was the assembly of a wooden-block puzzle, consisting of twenty one blocks that could be combined to form a six tier pyramid. The first five tiers of the pyramid consisted of four equal sized blocks, which had to first be combined into pairs, which are then joined together, with a single piece forming the top layer. Since the pieces in each layer were smaller than those in the last, they could only be placed in one layer, but it was designed so that the pegs that connected the pieces could fit into the holes in any other piece. The children were first presented ‘with the 21 blocks of various shapes and sizes spread out in a jumble’ (ibid, p92) and invited to play with them, before being guided through assembling the pyramid.

In discussing how we learn new skills, Wood et al state that

comprehension of the solution must precede production. That is to say, the learner must be able to *recognize* a solution to a particular class of problems before he is himself able to produce the steps leading to it without assistance.

(Wood, Bruner, and Ross 1976, p90).

If an adult, unfamiliar with this particular object, were asked to assemble it, they would have the requisite skills to work out what the solution is and thus how to achieve it. They might not yet know how to solve the problem, but they know how to solve problems like it, they have likely encountered many puzzles. They may do it dependently, with the support of another agent, but this support is not necessary. If you leave them for long enough, they would eventually work out how to assemble it. If they didn’t, moreover, they would know how to find out. Although we often learn from other agents, we need not — we might get better at yoga faster if we take a class with a teacher that helps move our bodies into the requisite positions, but we can use other kinds of teaching resource to get there in the end. What is distinctive to the child is that they might well not know what it is they need to work out how to do — they recognize that they need to put the pieces together, but they do not know, for instance, how to identify which pieces should go where. The younger children, in particular, needed guidance in this — ‘[n]one of the 3-yr-olds could put four blocks together correctly’ (ibid, p94) — but although they struggled with correctly assembling pieces, ‘they were just as adept at recognizing an appropriate one when they encountered it’ (ibid). The young children could identify their desired outcome, but unlike adult agents, they struggled to identify how to achieve that outcome. Recall Röd’s would-be wood-carver — they do not know what it is they need to do, beyond that it is ‘carving wood.’

They may well know the kind of end-state they are aiming at, but they do not know what it is they need to do to achieve that.

Scaffolding, as described by Wood et al, is what the tutor provides to enable the child to learn what it is they need to do. They characterize it in the following way:

well executed scaffolding begins by luring the child into actions that produce recognizable-for-him solutions. Once that is achieved, the tutor can interpret discrepancies to the child. Finally, the tutor stands in a confirmatory role until the tutee is checked out to fly on his own.'

(Wood, Bruner, and Ross 1976, p97)

They go on to detail six key features (1976, p99):

- '1. Recruitment. The tutor's first and obvious task is to enlist the problem solver's interest in and adherence to the requirements of the task.'
- '2. Reduction in degrees of freedom. This involves simplifying the task by reducing the number of constituent acts required to reach solution.'
- '3. Direction maintenance. Learners lag and regress to other aims, given limits in their interests and capacities. The tutor has the role of keeping them in pursuit of a particular objective.'
- '4. Marking critical features. A tutor by a variety of means marks or accentuates certain features of the task that are relevant. His marking provides information about the discrepancy between what the child has produced and what he would recognize as a correct production.'
- '5. Frustration control. There should be some such maxim as "Problem solving should be less dangerous or stressful with a tutor than without".'
- '6. Demonstration. Demonstrating or "modelling" solutions to a task, when closely observed, involves considerably more than simply performing in the presence of the tutee. It often involves an "idealization" of the act to be performed and it may involve completion or even explication of a solution already partially executed by the tutee himself.'

Hammond et al (2012, p275) set these out in contrast with non-scaffolding actions, which helps indicate the ways in which scaffolding is distinctive, by contrasting scaffolding support (on the left) with non-scaffolding behaviours (on the right):

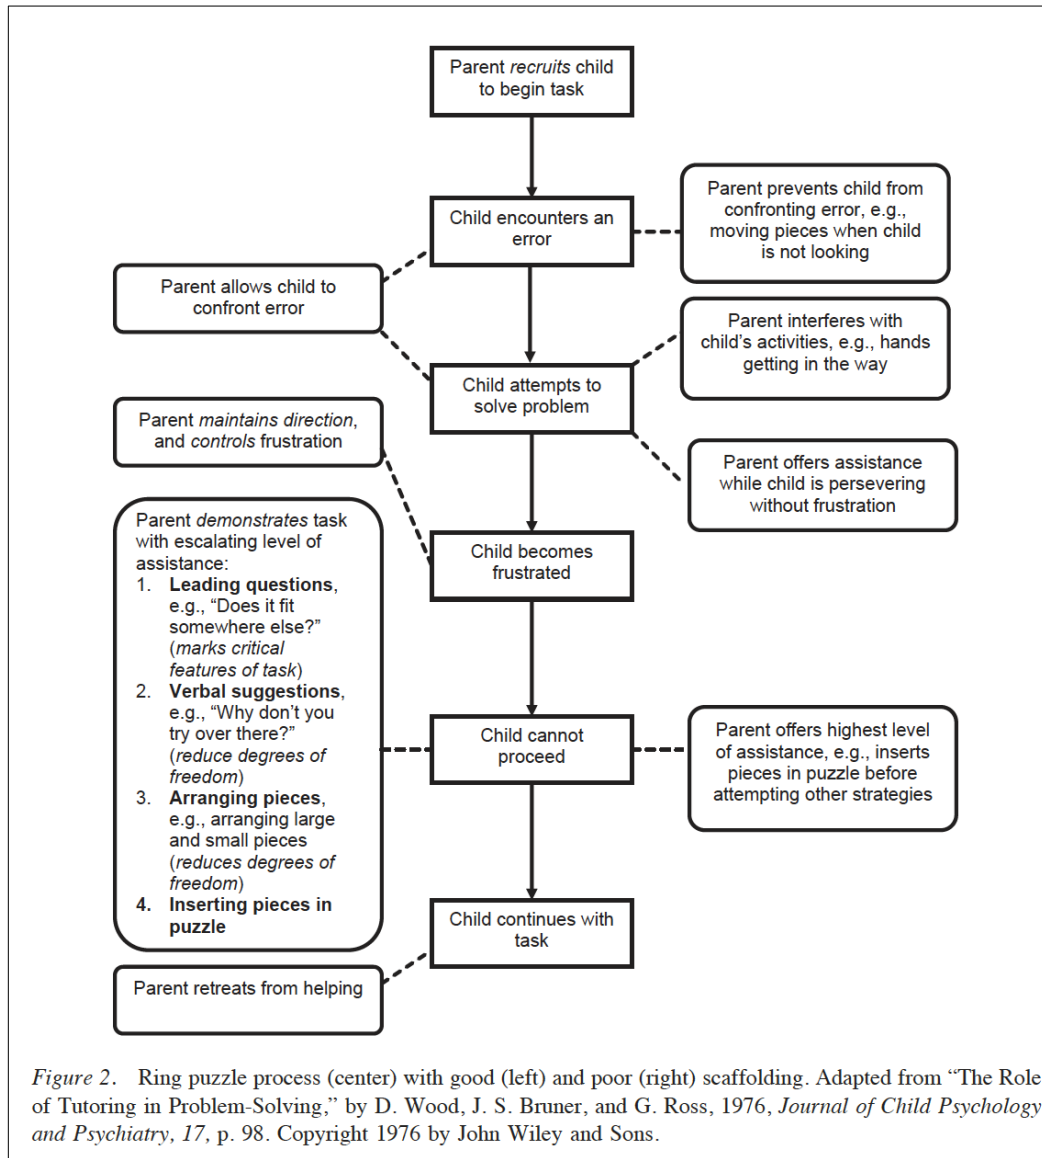


Figure 1 (Hammond et al. 2012, p275)

Reproduced with permission from American Psychological Association. No further reproduction or distribution is permitted.

For instance, they contrast the scaffolding action of allowing a child to confront an error they are making with the non-scaffolding action of preventing them from confronting the error by moving the pieces when the child is not looking. Sometimes, scaffolding is doing nothing at the right moment to

allow the child to work through the problem on their own. Similarly, interfering, or offering assistance when the child is attempting to solve the problem without frustration, are non-scaffolding actions, whereas scaffolding support is tailored to encouraging the child to persist on their own, and making problem solving easier only when they are struggling and unable to continue.

Each of these 6 features are described as elements of scaffolding as an explicitly pedagogical activity — they are things that are done as part of teaching the child how to carry out the act-type, in what Wood et al call the ‘tutorial process’. This need not be as rigidly defined and bounded as the kind of guidance that happens in an experiment designed to gauge how children learn. As they note,

Whether he is learning the procedures that constitute the skills of attending, communicating, manipulating objects, locomoting, or, indeed, a more effective problem solving procedure itself, there are usually others in attendance who help him on his way. Tutorial interactions are, in short, a crucial feature of infancy and childhood

(Wood, Bruner, and Ross 1976, p89)

Scaffolding is introduced in the literature as something that serves an explicitly pedagogical role — Wood et al identify its role in ‘tutorial interactions’, which serve to teach young children skills that they need for problem-solving. Although the tutor is supporting a child in solving a particular problem, they are not only teaching them this. They are teaching them how to solve problems. This is, as Wood et al note, commonplace in early childhood. In the next sections, I want to discuss the ways in which scaffolding is not solely an element of teaching. Not only does it enable a child to learn to act individually, it enables them to act whilst supported. I will do this by turning to the notion of executive function, which is, I will suggest, an integral part of much intentional action. Setting out scaffolding in this way, and its relationship to executive function, will enable me, in the final section of this chapter, to argue that scaffolded actions are dependent intentional actions, such that we can understand the scaffolded child as acting intentionally.

7.3 Executive function

Intentional action, as I have discussed, requires an agent not only to identify an end, but to work to realize that end, considering how they might go about doing what they set out to do, drawing on their

existing knowledge, and responding to hurdles that they come up against. Given this, executive function (EF) underpins much of our intentional action. The term is described variously as referring to:

- ‘cognitive processes that are required for the conscious, top-down control of action, thought, and emotions, and that are associated with neural systems involving the prefrontal cortex’ (Müller and Kerns 2015, p571)
- ‘higher mental processes that allow for flexible and complex goal-directed behavior... especially in novel and ambiguous situations...’ (Hammond et al. 2012, p271)
- ‘higher order, self-regulatory, cognitive processes that aid in the monitoring and control of thought and action’ (Carlson 2005, p595)

Although there is disagreement about how exactly to define the term amongst these authors, and others, there is commonality between these descriptions. *Executive function* refers to processes that are necessary for and integral to the completion of intentional action, by enabling the agent to ‘direct’ their own thoughts and action.

What are these processes? Some notable examples are ‘planning,... error correction and detection, and resistance to interference (Carlson 2005, p595) ‘anticipation, goal selection,... initiation of activity, self-regulation,.. and utilization of feedback (Anderson 2002, p71). There are three primary processes that are often seen as the ‘core’ of executive function,⁷² which are:

1. *Inhibition*,⁷³ that is, the ability to ignore distraction and stay focused, and to resist making one response and instead make another
2. *Working memory*, that is, the ability to hold information in mind and manipulate it
3. *Cognitive flexibility*, that is, the ability to flexibly switch perspectives, focus of attention, or response mappings’

(Diamond 2006, p70)

Whether or not we are inclined to think of those listed earlier as distinct from these three, they are bound up with them. Anticipation and planning clearly involve working memory, for instance, and self-regulation and error correction rely on inhibition (or inhibitory control (Marcovitch and Zelazo

⁷²For instance, by (Best and Miller 2010), (Center on the Developing Child and National Forum on Early Childhood Policy and Programs 2011)

⁷³Some models of executive function take it to amount to inhibition and/or working memory, with its other features stemming from these (Müller and Kerns 2015, pp575-585 for details of this discussion)

2009)). Cognitive flexibility (or attention-switching (Bibok, Carpendale, and Müller 2009) or shifting (Best and Miller 2010)) runs through all of these further processes — it is what allows us to think about each step in a plan, to consider alternatives and select goals, to understand and respond to feedback. Cognitive flexibility enables us to do this even when these also rely on inhibition and/or working memory. These three processes are also the most consistently researched,⁷⁴ and for these reasons, focussing on them initially seems the best way to track the development of executive function. Attending to three processes also encourages us to think of EF as something other than a singular whole, but this isn't to commit to an approach that suggests that we can separate out the component parts and study them in isolation (a factor analysis approach). The elements of EF cannot easily be separated out from one another, nor from the other processes they interact with. We can see the way it is bound up with other processes in the rife problem of task impurity in developing experimental procedures, since a task involving EF will also involve other cognitive processes, impeding our ability to measure EF alone.⁷⁵ I want to think of EF as something that arises out of the interaction of its component processes. These processes could (and do) develop at different rates, and EF would be directly impaired by impairments in one process, but also indirectly impaired as the impairment in one process impairs others. For instance, an impairment in working memory, would impact cognitive flexibility and so overall EF would be impaired both by the impairment to working memory and the way in which that further impairs cognitive flexibility.

Executive function is something that arises out of the combination of the processes (or 'cognitive abilities' (Diamond 2006, p70)) that make it up, and so they are all needed — having working memory without inhibition or cognitive-flexibility is to lack executive function, at least as I will be using the term. That is not to say that we must have them all to the same degree, or that one cannot be stronger than the others, but that executive function, understood as that which allows and enables control of actions, requires their interaction. Inhibition and attention-switching are quite readily understood as abilities, ones that we could be better or worse at. We might be inclined to think of memory as more like a resource than an ability, but it is generally understood to be 'comprised of a storage system and control system' (Müller and Kerns 2015, p577), and the control system is clearly an ability. Working

⁷⁴ This does not mean that this is the only model of executive function. Anderson (2002, p73) considers executive function to have four domains, each of which consists of several processes. Executive function is then what arises out of the interaction of these discrete domains: cognitive flexibility (which he takes to include working memory), goal setting, information processing, and attentional control (of which inhibition is a part). He claims that '[a]ttentional control processes greatly influence the functioning of the other executive domains, while information processing, cognitive flexibility, and goal setting domains are inter-related and inter-dependent.' (ibid)

⁷⁵ See Müller and Kerns (2015, pp577-579) for discussion of the difficulties in developing tests for EF.

memory is not merely the holding of information in the short-term, but the operating on it that is needed to utilize it in action — I need not only to remember which stitch I used last in my knitting, but to be able to use that information to determine which to use next. All three, then, are things that we can do, in some sense, as well as being things that help us do much else. Executive function is what we have when we can do those things, and when those abilities can interact in the ways that allow us to act intentionally. It is not a single ability, but we can understand it under the umbrella term because of the close relation between the three. Executive function, so understood, facilitates and supports integral parts of intentional action. I will explore this role in the next section, as I discuss the development of executive function.

7.3.1 The development of executive function and its role in intentional action

Executive function starts to develop in our first year. Bruner (1973) details the way that the infant's capacity for visually guided reaching develops in the first 12 months of life. He describes this as something that stems from an intention in the child, but as his conception of intention⁷⁶ differs significantly from that we would ordinarily use the term, I will describe this as goal-directed. An object elicits a goal to grasp the object, and the infant makes various attempts at grasping the object. 'In time, and probably by virtue of sheer practice of the act in the presence of the releasing stimulus, which permits the coordination of internal and peripheral feedback, the act is successfully executed' (ibid, p4). Perseveration is continuing to repeat an action after the stimulus for that action has been removed or altered, and the tendency to perseverate is used as a measure of executive function, particularly inhibition and working-memory, in childhood. Often, tests of executive function involve a change in the rules that a child is being asked to follow (Somsen 2007, for instance), and the ability to adapt to the new rule rather than perseverate with following the original one indicates developing executive function. Adapting to changes in circumstances in order to solve problems is part of how we act on intentions — we often have to alter the strategies we are using in the face of changing circumstances, and the ability to do this is part of intentional agency. Detour-reaching, which I discussed in §6.1, requires that infants can hold in mind the location in which an object has been hidden for longer, and that they control their behaviour so that they do not perseverate in looking in a location that is no longer correct (Diamond 2006, p71). This 'reaching requires holding a goal in mind, planning, and

⁷⁶ 'Intention as used here involves an internal discharge in the nervous system whereby an act about to occur is not only produced in the effectors by the usual motor volley but is also signaled to related sensory and coordination systems...' (Bruner 1973, p2)

inhibiting the strong tendency to reach straight for the goal' (ibid). As the infant gets better at tasks of this kind, knowing how and where to move their limbs to grasp the object, correcting their movements based on unsuccessful attempts, and responding to obstacles in the environment, they are relying on their working memory, being cognitively flexible in responding to feedback, and using inhibition to cease their initial, unsuccessful, courses of action. In the first year of life, then, the child begins to develop executive function.

This development continues into adolescence and even early adulthood (Best and Miller 2010), and executive function can vary between adults.⁷⁷ My interest is primarily in the early stages of this development. This is not quite the shift from Incompetence-to-Novicehood (Piñeros Glasscock 2021, p2) - I am not trying to make sense of how EF processes could come 'from nowhere' - but it is also not the shift from Novicehood-to-Mastery (ibid) as EF has a 'protracted development and late age of maturation' (Best, Miller, and Jones 2009, p189), continuing to develop through adolescence and into adulthood. Some aspects of inhibitory control do not reach adult levels until after the age of 15 (Huizinga, Dolan, and van der Molen 2006, p2029), but they also begin to appear by age 4 (Best and Miller 2010, p1646). This is part of a broader trend for a significant development in EF processes between the ages of 2 and 5. Carlson (2005) assessed EF in children at ages 2, 3, 4 and 5-6. The tests that are appropriate at 2 years of age are not all appropriate at 6 years of age, and vice versa, but amongst the tests that were administered to multiple age groups, 65% showed significant age-related improvement (ibid, p612). Moreover, although there was not a significant difference between the performance of young and older 3 year olds (36-41 months and 42-47 months respectively), there was 'significant improvement from young and older 3 to young 4, and again from young 4 to older 4 and 5+' (ibid, p607). This suggest that between 3 and 5, there is rapid improvement in EF skills. This period, when the child has already developed the very basics of EF, but long before they are refining them to adult levels, is the one in which I am interested.

When we think about young children, then, we are considering a group that is still developing abilities that are necessary for intentional action — executive function allows us to act on our intentions, to solve the problems that we might face in so doing, and to direct ourselves at our ends. In early childhood, we have not yet worked out how to act in novel situations, to work out how a new act-type might occur, or how to apply a familiar act-type to an unfamiliar context. This is what executive function facilitates. Once we have the rudiments of the core abilities, we can begin to act in goal-

⁷⁷ For instance, executive function deficits are a feature of ADHD in many, but not all, cases (Willcutt et al. 2005)

directed ways successfully, but since their development goes on into adulthood we are not equipped in the same way that a fully-fledged rational agent is. Nonetheless, what is clear is that young children have the beginnings of the requisite skills to act intentionally. EF enables agents to *do* various things — to resist particular responses, to develop plans of action, to consider various options, for instance. Given the way in which the abilities involved in EF are integral to responding to novel circumstances, it seems right to at least tentatively think of some of them as skills (although, again, not to think of EF as a skill.) Given this, I want to think back to Aristotle’s picture of learning crafts (which are a kind of skill), which I cited in §6.3.3. He remarks that ‘we learn a craft by producing the same product that we must produce when we have learned it; we become builders, for instance, by building, and we become harpists by playing the harp.’ (1999, Book II, Chapter 2, §3, 1103a32). We learn skills by engaging in the skill. Small argues that Aristotle is not puzzled by the idea that we can only learn to do something by doing it, because ‘the source of the learner’s action is the teacher’s skill’ (2014, p99). This echoes the thought that Rödl presented, that the child’s intention is rooted in the parent’s consciousness, and it’s one I will spell out in the following sections. I will look at both the way that caregivers support the acquisition of the skills involved in EF, and the way that this support can be part of a dependent intentional action of a child who as yet lacks these abilities when acting individually.

7.3.2 Teaching executive function

Research suggests that scaffolding by parents and caregivers is part of how children learn abilities like inhibition and attention switching. If we reflect on the six key features of scaffolding, we can see that they relate to elements of executive function. Many of them are clearly connected with *inhibition*. *Recruitment*, *direction maintenance*, and *frustration control* are part of ‘the ability to ignore distraction and remain focussed’, and *marking critical features* and *reduction in degrees of freedom* will help the child ‘to resist making one response and instead make another’. *Working memory* and *cognitive flexibility* will be guided by *demonstration*, which encourages the child to respond to the tutor’s perspective, remember what is demonstrated and adapt it to the next elements of the task. *Marking critical features* will also involve the switching of perspective and focus involved in recognizing what it is that makes these features critical, using *cognitive flexibility*, and deploys *working memory* in using what has been shown in new elements of the task. Scaffolding, then, relates to specific elements of executive function, and I will now turn to the evidence that it helps children to learn executive function.

This begins in infancy. Bernier et al (2010) investigated the impact of three aspects of parental behaviour on executive function (in particular, on the development of inhibition, working memory, and cognitive flexibility). They documented maternal sensitivity and mind-mindedness behaviours when the children were 12 months old, and autonomy support at 15 months old, and assessed EF at 18 and 26 months. *Maternal sensitivity* ‘consists of appropriate and consistent responses to infants’ signals,’ and *mind-mindedness* is ‘the parent’s tendency to use mental terms while talking to the child.’ (ibid, p328). They describe *autonomy support* as consisting ‘mainly of scaffolding, respecting the child’s rhythm, and ensuring that he or she plays an active role in successful completion of the task.’ I am inclined to think of all these features of autonomy support as elements of scaffolding. Scaffolding involves respecting the child’s rhythm and ensuring that the child plays an active role in the completion of the task — without these features, the child is not in control and the parent is not scaffolding the child’s action. This is reflected in Hammond et al’s (2012, p275) elaboration of the process of scaffolding, which, as I discussed in §7.1, emphasises the importance of only intervening when the child is struggling, and only in ways that support their solving the problem themselves. Bernier et al found that all three of these parental behaviours predicted improvements in some domain of EF. Although they concluded that mind-mindedness was the best explanation of improvements between 18 and 26 months, their central finding was that ‘autonomy support is the aspect of parenting that was most robustly related to age-specific indices of child EF.’ (Bernier, Carlson, and Whipple 2010, p335) They went so far as to say their ‘results would appear to suggest that autonomy support not only relates more consistently to child EF than maternal sensitivity and mind-mindedness but also that it is responsible for most links found between child EF and either sensitivity or mind-mindedness’ (ibid). That is, not only did they conclude that scaffolding in infancy contributed to the development of EF in infancy, but that scaffolding also played a significant role in the contributions of other parental behaviours to EF at this age.

Not only does empirical research suggest that scaffolding is influential on the development of EF in infancy, but it also indicates that it plays a role in the pre-school years. Hammond et al (2012) carried out a longitudinal study, assessing EF at ages 2, 3 and 4, and parental assistance on a problem-solving task at ages 2 and 3. They framed scaffolding as ‘emerg[ing] from the interaction of the parent and the child’ (ibid p274) which echoes my suggestion that the additional features of Bernier et al’s autonomy support are also part of scaffolding. They go on to say that ‘its structure is determined by the interaction between parent and child.’ (ibid). They found that scaffolding at age 3 significantly predicted EF at age 4, indicating a relationship between them. They also found an indirect

relationship between scaffolding at age 2 and executive function at age 4, via verbal ability. (ibid, p278) This suggests that scaffolding is part of how children learn executive function, and the apparent relationship with verbal ability also lends support to thinking of executive function not as a single discrete skill but something that is bound up with a variety of other cognitive skills (as well as itself arising out of the interaction of various cognitive skills.)

Bibok et al (2009) consider which kinds of parental support encourage the development of EF in more detail. Much like Hammond et al, they consider not only whether parental support is being offered in solving a problem, but what kind of support is being offered. Instead of simply characterizing the support offered as scaffolding/not scaffolding, however, Bibok et al distinguished between three kinds of puzzle-solving activities that the child could be engaged in, and two kind of utterance the parent could be making. They observed children aged 20-29 months carrying out the ring-puzzle task (also used by Hammond et al) which involves assembling concentric rings, which each are split into several identical pieces (although the pieces of each ring are different from those of the other rings). They classified children's activities based on whether they were an incorrect placement of a piece in a space on the board (curvature), a backwards placement (backwards) or a correct placement (correct) (ibid, p26). They classified task-relevant parental utterances as either *directive* or *elaborative*. Directive utterances 'command, direct, or state the *future* course of action the child should take next' and elaborative utterances by the parent that 'either elaborate on or evaluate the child's *presently* occurring course of action'. (ibid, pp26-27). The novelty of the puzzle meant that parents would likely 'be unable to rely on foreknowledge of their children's difficulties' (ibid, p27) meaning that they would struggle to provide 'anticipatory scaffolding' (ibid). Scaffolding had to be done in response to the child's activities and difficulties, rather than in expectation of them, as the parents were not in a position to predict them in the way that they would be with a familiar task. Given this, elaborative utterances were more likely to be scaffolding than directive ones, since these are future-looking, and anticipatory scaffolding was unlikely. Moreover, directive utterances are by their nature less likely to be scaffolding, since they are 'defined as utterances giving children less opportunity for choice, instead emphasizing expected activities or behaviors.' (ibid, p20) Their research showed that elaborative utterances were predictive of attention-switching EF (ie cognitive flexibility) in the children, whereas directive utterances were not (ibid, p28). This means that it showed scaffolding to be predictive of attention-switching EF in a way that non-scaffolding behaviours were not.

Bibok et al's analysis found that it was not only the content of the utterances that mattered, but also the timing – part of why elaborative utterances are valuable is that they are contingently relevant to the child's activities at the time of utterance. This timing alone is not enough, as directive utterances that were contingently relevant were not found to be significant in EF (Bibok, Carpendale, and Müller 2009, p28) — it is the delivering of content that is relevant to what is being done rather than what will be done in the future, at the very time that is done, that provides the kind of support that aids the development of EF. By making elaborative utterances at the time of the child's acting, the parent provides guidance that the child can relate to their activities, and thus 'construct the meaning of the utterances' (ibid, p30). The parental utterances are not simply things said to the child about their activities, but things that the child understands through and in relation to their activities. I am using 'activities' broadly here, to include the child's cognitive activities – it is not merely that the parent helps the child do something different, but also that they support the child's working out how to tackle the problem. The distinction between physical and cognitive activities in the context of this kind of problem-solving action is fairly arbitrary, but the thought is that not only does the scaffolding support the current attempt by guiding the child in another direction, but it also helps them consider how to go about this kind of problem-solving. It impacts not just their current cognitive activities but their capacity for various cognitive activities and which ones they might initiate at which moments. They are being taught not just how to solve this problem, but how to solve problems.

What the research suggests, then, is that scaffolding facilitates the learning of executive function. I now want to suggest that this is achieved in part by providing executive function. The scaffolding agent supports the child in successfully acting, in part by playing the role that executive function might otherwise play. This is reminiscent of a suggestion made by Clark and Chalmers that a thinker's 'mental states [could] be partly constituted by the states of other thinkers' (1998, p17) but the claim I am making is a little different. It is not that the child's mental states are partially constituted by the caregiver's, but rather that they outsource some of the abilities involved in intentional action to the caregiver. In the next section, I will consider how this kind of support not only teaches executive function, but how it provides executive function during teaching.

7.3.3 Exogenous executive function

The findings that I discussed in the previous section suggested that scaffolding can teach executive function. Looking more closely at how the scaffolding functions, and what it provides, reveals that

scaffolding also appears to provide executive function during the action. Bibok et al say that, '[b]y commenting on their children's activities in a temporally contingent manner... parents facilitate their children in either examining their own activity or switching to a new activity. That is, contingent parental utterances assist children in not perseverating in erroneous activities.' (2009, p25) As noted earlier, in tracking the development of EF in childhood, researchers look for the cessation of perseveration (eg Somsen, 2007). In encouraging children not to persevere, parents are encouraging the use of EF skills, and thus the development of these skills. Bibok et al see this as a mark of attention-switching EF, but it seems right to suggest that there is also an element of inhibition – considering and attempting alternative strategies requires also resisting the inclination to continue with what one is already doing. Ceasing to persevere is both trying something new and stopping repeated efforts at the previous tactic. What the scaffolding does is offer the prompts associated with EF, demonstrating EF skills, and leading to the outcomes that EF allows. Without these prompts, the children are likely to persevere in their unsuccessful attempts.

We can go further, though, than thinking of this as providing guidance in EF. Bibok et al say:

Parental utterances can help children detect an erroneous puzzlesolving activity, or help them determine why a successful activity was productive (that is, why a successful activity was not in error). Parents therefore serve as an auxiliary and exogenous form of attention-switching EF for their children.

(Bibok, Carpendale, and Müller 2009, p25).

This suggests that we can think of what happens in at least some instances of scaffolding as a kind of 'outsourced' executive function. That is, what the scaffolding caregiver provides is not merely guidance and teaching of EF, or examples of EF, but the EF itself. The child does not have the capacity for attention-switching EF to the level necessary to solve the puzzle presented to them, but the parent does. In helping the child, part of what the parent is offering is their own EF, directed at the end of the child's puzzle-solving. This thought is not only expressed by Bibok et al. It is present in Bernier et al's (2010) terminology of 'autonomy support', which suggests that what is provided is something that strengthens the child's own capacity for autonomous action — the outsourced EF supports the autonomy of the child. When they discuss what is provided by autonomy support, they state it 'involved externally guided problem solving' (ibid, p335) — that is, the caregiver's scaffolding is providing what would ordinarily occur internally, the EF needed for solving the problem at hand.

They also describe parents as ‘external regulators’ (ibid) for the child in infancy and the acquisition of EF as the transition ‘from being externally regulated to self-regulated’ (ibid, p328). That is, they suggest that what is provided by the parent becomes what the child can do for themselves, and so again we have the idea that scaffolding involves outsourced EF. Munakata et al (2012) set out what they take to be the three key transitions in the development of cognitive control, a notion clearly related to that of EF. The third of these is the transition from ‘Externally Driven to Self-Directed Control’ (ibid, p74). Although some of the exogenous support of attention switching can come from features of the external environment, it is also clear that part of this is the provision of goals and direction from caregivers – for instance, ‘stopping playing and putting away toys when told’ (ibid) and they discuss the impact of different kinds of guidance in switching behaviours — being told which rule to switch to in a card-sorting task is more likely to end perseveration with an old rule than merely being told to change rules with no further instructions. Offering a specific rule to switch to is a kind of ‘reduction in degrees of freedom’ ie an instance of scaffolding. Again, then, it seems that what we have is control being provided first externally and then by the child themselves, we seem to have outsourced the EF provided by scaffolding.

What is my claim here? It is that empirical research points in favour of the claim that we can understand at least some instances of scaffolding as cases in which some of what is necessary for the completion of an action is provided by another agent. Scaffolding, as I have been considering it, is a feature of particular teaching contexts. It enables children to better solve the problem at hand when they are being taught to solve that problem. Although solving a particular concrete problem, or developing a particular skill, is the primary purpose of the teaching, it is not the only thing that is taught. Elements of scaffolding correlate with features of EF, and one of the things that is taught in scaffolding is component abilities of executive function. These are taught by demonstration, as the child learns what the application of these abilities allows. The caregiver suggests alternative courses of action, for instance, which demonstrates what attention-switching can allow, and they remind the child of things that have happened before in maintaining their direction, demonstrating working-memory and how to utilize it. The child is able to solve the puzzle that they are tackling because of the contribution of this teaching. The further claim is that the contribution of this teaching is itself EF that is used by the child, but EF that is external to the child’s own processes. This is not an unusual idea, it is just that the outsourced processes are often more straightforwardly physical — dialysis outsources the purification of the blood, for instance. The process happens outside the body, but it also supports the internal processes that rely on purified blood, such as the normal functioning of the

circulatory system. In adults, who have fully-developed EF, it is ordinarily internal and integrated with other cognitive processes. When children are in the process of developing executive function and learning the skills of executive function through scaffolding, they can act in ways that require executive function on the back of caregiver's directions that stem from the caregiver's own executive function.

7.5 Scaffolded action as dependent action

In the preceding sections, I have argued that we can understand scaffolding not only as supporting the development of executive function in early childhood, but also as providing exogenous executive function for the actions of the young child. What I need to do now, is clarify how this ties in to my broader discussion of dependent intentional action. In particular, I want to suggest that we can sometimes understand the scaffolded child as engaging in a dependent intentional action. Dependent intentional action has the following form:

Dependent Intentional Action

An event X is a dependent intentional action of an agent A , supported by an agent B , iff

- (D1) X is an event of A 's ϕ -in(g)
- (D2) A is acting with an intention to ϕ
- (D3) B is acting with an intention to help A ϕ
- (D4) B is ϕ -ing because their ϕ ing contributes to the realization of their intention help A ϕ
- (D5) A guides X such that X occurs in keeping with their direction
- (D6) B is attuned to and responsive to A 's guidance and to A 's actions in realizing their intention to ϕ

I will first elaborate on the kinds of scaffolding detailed by Wood et al, which I mentioned in §7.2, suggesting their compatibility with the conditions of dependent intentional action, before discussing this more directly. Those kinds of scaffolding are:

1. Recruitment: This is not simply trying to get the child interested in the task. It is trying to get them to be interested in and to follow the *requirements* of the task. The child can want to do something, but not know how to, and here recruitment is about directing them to its necessary features, about showing them what it is to ϕ and

encouraging them to do exactly that. This is consistent with their being the one who showed initial interest in ϕ -ing, in keeping with (D2)

2. *Reduction in degrees of freedom:* This is something we do all the time in helping children achieve their aims, it is the reduction of the number of constituent acts. Sometimes this means suggesting a simpler version of the task, sometimes this means doing some of the constituent acts ourselves. We provide the materials, we suggest how they might carry it out. The child has painted the picture even if you gave them the paints, cleaned their brushes between colours, and suggested a less complicated design. This is in keeping with (D3) and (D4) — these actions are directed at helping the child with what they set out to do.

3. *Direction maintenance:* This is the act of keeping the child on task. This is trickier than some of the other features of scaffolding, if we want to think of ourselves as nonetheless supporting rather than doing for the child. If we persist in keeping the child ‘on track’ after they have expressed a firm desire to cease the activity, we would not obviously be supporting their efforts. However, if what we do is keep them focussed on something they desire to do, but in which their interest fades as it becomes difficult or time-consuming, we are supporting their action. This is not to say that if they no longer wanted to carry out this action, and we kept them doing it, the action would not be their own – they would still have done it, begrudgingly. As the preceding discussion has explored, young children are still in the process of developing executive function, which is integral to keeping oneself on-task. Without these, there can be a substantive disconnect between what it is one wants to do, and what one does, and direction-maintenance from a supporting agent helps close this gap. Direction-maintenance is a clear instance of exogenous executive function, since it is a way of getting the child to persevere (not perseverate) in the way that the adult (with ordinary executive function) would be able to themselves. This involves a responsiveness to the child in keeping with (D6), and is compatible with (D2).

4. *Marking critical features:* This is the pointing out of the important elements of the task, as a way of guiding the child to recognize the problems with their current strategies. When the child is distracted by the wrong bits of a task, because they do not fully understand the act-type, or the object on which they are acting, we can refocus them to what matters. This responsiveness to the child’s distraction is in keeping with (D6).

5. Frustration control: This is in many ways a feature that runs through all of the other features – by providing scaffolding, we reduce the likelihood of the child losing interest in the task, and giving up. We can suggest where to refocus attention, remove distractions, and give guidance, and in doing this we make the task less frustrating. This, again, is in keeping with (D6)

6. Demonstrating: This is not ϕ -ing on the child's behalf, but a way of showing the child how to ϕ themselves, when other kinds of explanation do not succeed. This is something that we might do only for a small part of what they are doing (like a few stitches of their knitting), or with a distinct object (such as painting on another piece of paper). We might do and then undo something, such as placing a puzzle piece in the right place and then taking it back out. Demonstrating is not the same as doing something for the child, it is showing them how to do something for themselves, and so is in keeping with (D4).

With this kind of picture of scaffolding in sight, and the suggestions I have made about its compatibility with dependent intentional action, I want to tackle this more directly, considering whether conditions (D1-D6) can be met by a scaffolded action. I will set aside (D1) and (D2) initially. (D3) and (D4) are fairly straightforwardly satisfied by scaffolding as it has been described, since the adult is acting to help the child solve the puzzle. Even in the instances in which the adult's behaviours are directive rather than elaborative, they are doing this because they understand this as a way of supporting the child in what they are doing. (D5) and (D6) effectively require that the adult does not overrule the child, but help them to do that which they are trying to do. For this, it seems that we need support that is at least primarily elaborative rather than direction – Bibok et al, drawing on Landry et al. (2000) suggest that 'directive utterances may require children to first abandon their present activity, and then orient elsewhere' (Bibok, Carpendale, and Müller 2009, p23) which would mean that the child is not guiding but instead being guided. If we look at Figure 1 in §7.2, all the examples of non-scaffolding behaviours involve interfering with the child's direction of the action — preventing them from guiding it. In scaffolding, however, the child is in control of the solving of the puzzle — they use the suggestions given by the adult, but that is not the same as following instructions given by the adult. Bibok et al's findings lend 'support to the notion that scaffolding is a process that is led by the child, with tutors adjusting and accommodating to the child's current level of performance' (ibid, p30). Given that scaffolding is responsive to what the child is doing, and consists of behaviours designed to support those actions and help the child to work out themselves how to solve the problem, we have

conditions (D5) and (D6). These entail that (D1) holds — what the adult is doing is clearly supporting the child's solving of the puzzle not solving it themselves. This entails that we have a dependent action of a kind, but without (D2), it is not a dependent intentional action. What needs to be the case for this is that the child is not merely solving the puzzle, but acting with an intention to solve the puzzle.

Acting with an intention to ϕ does not require having some particular mental state that we can identify prior to the beginning of the agent's ϕ -ing. Instead, what is needed is that the agent is acting intentionally, doing the things that they are because they contribute to realizing an intention to ϕ with which they are acting. It might be objected that the child is not acting intentionally, but merely purposively. What I want to suggest is that, in virtue of the executive function that the adult provides, we can understand the child's action as intentional. It is clear that the child at least wants to do this. We can see this has to be the case when we look at the methodology of developmental psychology research, which shows that only those children who want to carry out these kinds of tasks are involved in the experiments. Schünemann et al (2022, p7) excluded children because they were 'uncooperative', Liskai-Peres et al (2021, p5) excluded 16 infants for 'passivity' and 11 for 'fussiness', Kamps et al (Kamps et al. 2020, p7) excluded 6 children because they 'refused to take the sticker' needed for the experiment and Allen & Bickhard (2018, p84) excluded 3 children because 'they did not want to play the games'. These reasons for exclusion indicate the way in which those children who *do* participate in such experiments do so volitionally, since they could simply opt out. What I want to suggest is that they are not only acting with a desire to solve the puzzle, but with an intention, because they are committed to doing this, they are changing strategies when they do not work, they are doing the things that we do when we act intentionally. Although the task is presented to them, the intention to participate in it is theirs. That the children sometimes require direction-maintenance from a supporting adult does not mean that they are not intentionally acting — the need for this is due to their still-developing executive function, not a lack of interest in the task. The support that the adult provides enables the child to act intentionally.

Understanding scaffolded actions as dependent intentional actions chimes with Bibok et al's conclusion:

What the findings of this study globally suggest, therefore, is that both scaffolding and cognitive development are inherently active processes... For example, simply exposing children to elaborative utterances does not guarantee that children will

developmentally benefit from such utterances. Rather, such utterances must temporally occur in relation to children's activities such that children are able to construct the meaning of the utterances by relating them to their own activities. Thus scaffolding utterances are defined by how children cognitively use such utterances, and not by the purposes for which the tutor intended them. That is, the meaning of instructive utterances does not come from outside the children's own activity. Tutors can therefore facilitate children's development by timing their utterances so that they feed into the children's constructivist activities. Tutors cannot, however, circumvent the children's constructivist activity, nor externally or environmentally cause the children to develop.

(Bibok, Carpendale, and Müller 2009, p30)

It is not just that we *can* so understand scaffolded actions, but rather that we *must* so understand them, if we are to account for their character properly. Scaffolding does not work by the adult directing the child, but rather by the child reacting to the adult's utterances. Their role cannot be imposed on the child, the adult cannot make the child act in the desired way, nor can they act in this way for the child. The scaffolding mechanism only works in virtue of the child aiming at solving the puzzle. Bibok et al note that 'it was the children's level of performance (that is, puzzle placement errors) that determined when parental elaborative utterances would be most predictive of children's attention-switching EF' (Bibok, Carpendale, and Müller 2009, p30). That is, it is only in virtue of what it is the child is doing and how they respond to the adult that the scaffolding has effect. Given this, it seems that we can understand young children as engaging in dependent intentional action, even if we are unsure of their ability to engage in individual intentional action.

In this chapter, I have argued that young children are able to engage in intentional action when acting dependently. The core idea is that young children can be dependently intentional. I began in §7.1 by considering an argument offered by Rödl, that suggests that caregivers can provide determinate content to the intention of a young child, making the child's action intentional even when they understand little of what they are trying to do. In §7.2, I introduced the notion of scaffolding as used in developmental literature, a kind of support that facilitates learning and extends the abilities of the scaffolded child. I then turned in §7.3 to a particular element of agency that is often scaffolded, explaining the role of executive function in action, and in §7.3.1 discussing when we develop this ability. In §7.4, I combined these considerations, laying out empirical research on the role that scaffolding plays in the development of executive function. I built on this in §7.4.1, arguing that one of

the forms that scaffolding takes is not only the teaching of executive function, but the provision of executive function for the ongoing action. I then argued in §7.5, that scaffolded actions are dependent actions, and that, in particular, the executive function provided by scaffolding can enable a child who could not engage in individual intentional action to engage in dependent intentional action.

Conclusion

In this thesis, I have been concerned with the ways that human agents can be involved in intentional action. In particular, I have worked to show that there is a kind of involvement that an agent can have in the action of another that enables and extends their agency. Although this framing suggests that the interest is on the agent that provides this support, I have endeavoured to focus on the supported agent throughout, showing that there are ways of being supported that do not undermine an agent's status in relation to their own action. That dependent intentional action is action of the supported agent, and that it enhances that which they can do when acting alone, is something I have aimed to keep in sight throughout. Being helped does not entail a lack or loss of agency.

I explored these themes across three broad parts. In the first part, consisting of chapters 1-3, I considered the nature of intentional action, both individual and joint, setting out criteria for these kinds of action that are compatible with the approaches found in the existing literature. In chapter 3, I drew out the central way of being related to an action, that of being its agent. Here, my considerations were primarily those of moral psychology, considering what is necessary for an agent's being morally responsible for an action. In arguing that this requires identity with the agent of the action, however, I hope to have shown the importance of this relation to an action. The considerations of this chapter are ones I hope will have resonated through the rest of the thesis. If, as I argued, moral responsibility for an action requires being its agent, then we can infer the other way. When we identify an agent as morally responsible for an action as an action, we are implicitly making a judgement that they are its agent. Of course, we need to take care here to ensure that we are judging them action-responsible, rather than holding them responsible for its consequences or for coercing another to so act, but if we truly deem an agent to be so-responsible, we are judging them to be its agent, even if others are involved.

In the second part of the thesis, chapters 4 and 5, I characterized *dependent intentional action*. Although this terminology is new, I endeavoured to show that this kind of action is familiar. The work of these chapters was in part a kind of conceptual engineering, carving this category out from proximate phenomena as a distinct and important kind of action. This is a kind of action that necessarily

involves active involvement of multiple agents, but that is correctly understood as the action of one of these agents in particular. I identified the features of this kind of action, showing their importance to it being action of this particular kind, and arguing for the coherence of understanding this as the action of the agent that is acting with help. Central to this discussion is the claim I drew out in the first parts of chapter 5, that we cannot assimilate dependent intentional actions into joint action. Centring the way that the agents involved understand what they are up to enabled me to show that this is a way for an agent to do something with help, not to act jointly with another.

Finally, I turned in chapters 6 and 7 to consider the agency of young children, something that, like dependent intentional action, has been paid less attention than I believe it warrants. I began this part of the thesis by stepping away from dependent intentional action, considering instead whether there are barriers to young children engaging in individual and joint intentional action. By arguing that no such insurmountable barriers exist, however, I opened up the possibility that they can engage in dependent intentional action, and in the final chapter I discussed the ways that dependent intentional action can expand the agency of young children. This, I argued, occurs not merely through teaching them new act-types, but through providing the support needed to act in ways they could not alone. Part of this support, I suggested, enables the actions of young children to be intentional in ways that they would not otherwise be, by providing the executive function that is often integral to realizing an intention. As well as these being dependent intentional actions, they are dependently intentional actions – their being intentional depends on the support involved. Looking at young children, as we find them, amongst their caregivers, allows us to recognize their abilities as farther reaching than when we look at them alone.

At the core of this thesis is the idea that what human agents can do is what they can do when helped, not only what they can do in isolation. This is what dependent intentional agency shows us, that we can do things with support that we cannot do alone, and that this support does not undermine but rather enhances our agency. My arguments have not been moral or political, but if we think that we have a duty to organize society such that people can flourish, can achieve the things they desire, then the need to provide people with the support needed for dependent intentional action is clear. I started this thesis with an insight from the disability movement and it is where I want to close things — with recognition of the value of the support that enables people to do the things they intend to do.

Bibliography

Cases, Laws and Conventions

Alexander v. State 102 So. 597, 598 (Ala. Ct. App. 1925)

Obiter in R v. Giannetto [1997] 1. Cr. App. R. 1

State v. Duran 526 P.2D 188 (N.M. Ct. App. 1974)

United States v. Ortega, 44 F.3d 505 (7th Cir. 1995)

Wilcox v. Jeffrey [1951] 1 All E.R. 464 (K.B.).

OPINIONS OF THE LORDS OF APPEAL FOR JUDGMENT IN THE CAUSE R v JTB [2009] UKHL 20
R v Cogan and Leak [1976] QB 217

Children and Young Persons Act 1933. (23 & 24 Geo. 5. c. 12). <https://www.legislation.gov.uk/ukpga/Geo5/23-24/12/contents>

Children and Young Persons Act 1963. (c. 37). <https://www.legislation.gov.uk/ukpga/1963/37/contents>.

UN General Assembly. 2007. *Convention on the Rights of Persons with Disabilities : Resolution. A/RES/61/106.*
<https://www.refworld.org/docid/45f973632.html>.

Texts

Allen, Jedediah W.P., and Mark H. Bickhard. 2018. 'Stage Fright: Internal Reflection as a Domain General Enabling Constraint on the Emergence of Explicit Thought'. *Cognitive Development* 45 (January): 77–91.
<https://doi.org/10.1016/j.cogdev.2017.12.005>

Alvarez, Maria, and John Hyman. 1998. 'Agents and Their Actions'. *Philosophy* 73 (284): 219–45.

Anderson, Peter. 2002. 'Assessment and Development of Executive Function (EF) During Childhood'. *Child Neuropsychology* 8 (2): 71–82. <https://doi.org/10.1076/chin.8.2.71.8724>

Anscombe, G.E.M. 1966. 'A Note on Mr. Bennett'. *Analysis* 26 (6): 208–208.
<https://doi.org/10.1093/analys/26.6.208>

———. 1979. 'Under a Description'. *Noûs* 13 (2): 219. <https://doi.org/10.2307/2214398>

———. (1957) 2000. *Intention*. 2nd ed. London: Harvard University Press.

- . 2005. 'Practical Inference'. In *Human Life, Action and Ethics: Essays by G.E.M. Anscombe*, edited by Mary Geach and Luke Gormally, 109–47. Exeter: Imprint Academic.
- Apperly, Ian A., and Stephen A. Butterfill. 2009. 'Do Humans Have Two Systems to Track Beliefs and Belief-Like States?' *Psychological Review* 116 (4): 953–70. <https://doi.org/10.1037/a0016923>
- Aquinas, Thomas. (1485) 1920. *The Summa Theologiae of St Thomas Aquinas*. Translated by Fathers of the English Dominican Province. <https://www.newadvent.org/summa/>
- Aristotle. 1999. *Nicomachean Ethics*. Translated by Terence Irwin. 2nd ed. Indianapolis: Hackett.
- Armstrong, D. M. 1973. 'Acting and Trying'. *Philosophical Papers* 2 (1): 1–15. <https://doi.org/10.1080/05568647309506426>
- Asarnow, Samuel. 2020. 'Shared Agency without Shared Intention'. *The Philosophical Quarterly* 0 (0): 1–24. <https://doi.org/10.1093/pq/pqaa012>
- Baier, Annette C. 1970. 'Act and Intent'. *The Journal of Philosophy* 67 (19): 648–58.
- . 1997. 'Doing Things With Others: The Mental Commons'. In *Commonality and Particularity in Ethics*, edited by Lilli Alanen, Sara Heinämaa, and Thomas Wallgren. London: Macmillan Press LTD. <https://doi.org/10.1007/978-1-349-25602-0>
- Bakeman, Roger, and Lauren B. Adamson. 1986. 'Infants' Conventionalized Acts: Gestures and Words with Mothers and Peers'. *Infant Behavior and Development* 9 (2): 215–30. [https://doi.org/10.1016/0163-6383\(86\)90030-5](https://doi.org/10.1016/0163-6383(86)90030-5)
- Bakhurst, David. 2011. *The Formation of Reason*. Oxford: Wiley-Blackwell.
- Banks, Sarah, and Kirsten Nøhr, eds. 2013. 'The Reluctant Vegan: The Case of an Older Man in a Swedish Nursing Home'. In *Practicing Social Work Ethics Around the World: Cases and Commentaries*, 0 ed. New York: Routledge. <https://doi.org/10.4324/9780203807293>
- Barnes, Colin and British Council of Organizations of Disabled People. 1993. *Making Our Own Choices: Independent Living, Personal Assistance and Disabled People: Report of the BCOOP Seminar on Independent Living and Personal Assistance*. Belper: British Council of Organisations of Disabled People. <https://archive.org/details/makingourownchoi0000unse>
- Bartsch, Karen, and Henry M. Wellman. 1995. *Children Talk about the Mind*. New York: Oxford University Press.
- Bellagamba, Francesca, and Michael Tomasello. 1999. 'Re-Enacting Intended Acts: Comparing 12- and 18-Month-Olds'. *Infant Behavior and Development* 22 (2): 277–82. [https://doi.org/10.1016/S0163-6383\(99\)00002-8](https://doi.org/10.1016/S0163-6383(99)00002-8)
- Bernier, Annie, Stephanie M. Carlson, and Natasha Whipple. 2010. 'From External Regulation to Self-Regulation: Early Parenting Precursors of Young Children's Executive Functioning'. *Child Development* 81 (1): 326–39. <https://doi.org/10.1111/j.1467-8624.2009.01397.x>
- Best, John R., and Patricia H. Miller. 2010. 'A Developmental Perspective on Executive Function: Development of Executive Functions'. *Child Development* 81 (6): 1641–60. <https://doi.org/10.1111/j.1467-8624.2010.01499.x>

- Best, John R., Patricia H. Miller, and Lara L. Jones. 2009. 'Executive Functions after Age 5: Changes and Correlates'. *Developmental Review* 29 (3): 180–200. <https://doi.org/10.1016/j.dr.2009.05.002>
- Bibok, Maximilian B., Jeremy I. M. Carpendale, and Ulrich Müller. 2009. 'Parental Scaffolding and the Development of Executive Function'. *New Directions for Child and Adolescent Development* 2009 (123): 17–34. <https://doi.org/10.1002/cd.233>
- Bishop, John. 1989. *Natural Agency: An Essay on the Causal Theory of Action*. Cambridge Studies in Philosophy. Cambridge [England]; New York: Cambridge University Press.
- Boyle, Matthew. 2012. 'Essentially Rational Animals'. In *Rethinking Epistemology*, edited by Günter Abel and James Conant. Berlin, Boston: DE GRUYTER. <https://doi.org/10.1515/9783110277944.395>
- . 2016. 'Additive Theories of Rationality: A Critique'. *European Journal of Philosophy* 24 (3): 527–55. <https://doi.org/10.1111/ejop.12135>
- Bratman, Michael E. 1984. 'Two Faces of Intention'. *The Philosophical Review* 93 (3): 375–405. <https://doi.org/10.1017/CBO9781107415324.004>
- . 1987. *Intention, Plans, and Practical Reason*. *The Philosophical Review*. Cambridge, MA: Harvard University Press. <https://doi.org/10.2307/2185304>
- . 1992. 'Shared Cooperative Activity'. *The Philosophical Review* 101 (2): 327–41. <https://doi.org/10.2307/2185537>
- . 1993. 'Shared Intention'. *Ethics* 104 (1): 97–113.
- . 2009. 'I Intend That We J'. *Faces of Intention*, 142–62. <https://doi.org/10.1017/cbo9780511625190.008>
- Brewer, Bill. 1993. 'The Integration of Spatial Vision and Action'. In *Spatial Representation: Problems in Philosophy and Psychology*, 294–316. Oxford: Blackwell.
- Brooks, Anna. 2018. 'Koko, the Beloved Gorilla Who Communicated through Sign-Language, Dies at Age 46'. *Popular Science*, 22 March 2018. <https://www.popsci.com/koko-sign-language-gorilla/>
- Brownell, Celia A., and Michael Sean Carriger. 1990. 'Changes in Cooperation and Self-Other Differentiation during the Second Year'. *Child Development* 61 (4): 1164. <https://doi.org/10.2307/1130884>
- Bruner, Jerome S. 1973. 'Organization of Early Skilled Action'. *Child Development* 44 (1): 1–11.
- Burge, Tyler. 1979. 'Individualism and the Mental'. *Midwest Studies in Philosophy* 4: 73–121. <https://doi.org/10.1111/j.1475-4975.1979.tb00374.x>
- . 2003. 'Perceptual Entitlement'. *Philosophy and Phenomenological Research* 67 (3): 503–48. <https://doi.org/10.1111/j.1933-1592.2003.tb00307.x>
- Butler, Joseph. 2008. 'Of Personal Identity'. In *Personal Identity*, edited by John Perry, 2nd ed., 99–105. London: University of California Press.
- Buttelmann, Frances, Janina Suhrke, and David Buttelmann. 2015. 'What You Get Is What You Believe: Eighteen-Month-Olds Demonstrate Belief Understanding in an Unexpected-Identity Task'. *Journal of Experimental Child Psychology* 131: 94–103. <https://doi.org/10.1016/j.jecp.2014.11.009>

- Butterfill, Stephen A., and Ian A. Apperly. 2013. 'How to Construct a Minimal Theory of Mind'. *Mind & Language* 28 (5): 606–37. <https://doi.org/10.1111/mila.12036>
- Butterworth, George, and Edward Cochran. 1980. 'Towards a Mechanism of Joint Visual Attention in Human Infancy'. *International Journal of Behavioral Development* 3 (3): 253–72. <https://doi.org/10.1177/016502548000300303>
- Campbell, Lucy. 2019. 'Propositionalism about Intention: Shifting the Burden of Proof'. *Canadian Journal of Philosophy* 49 (2): 230–52. <https://doi.org/10.1080/00455091.2018.1512819>
- Carlson, Stephanie M. 2005. 'Developmentally Sensitive Measures of Executive Function in Preschool Children'. *Developmental Neuropsychology* 28 (2): 595–616. https://doi.org/10.1207/s15326942dn2802_3
- Carpenter, Malinda, Nameera Akhtar, and Michael Tomasello. 1998. 'Fourteen- through 18-Month-Old Infants Differentially Imitate Intentional and Accidental Actions'. *Infant Behavior and Development* 21 (2): 315–30. [https://doi.org/10.1016/S0163-6383\(98\)90009-1](https://doi.org/10.1016/S0163-6383(98)90009-1)
- Center for Disease Control and Prevention. 2022. 'CDC's Developmental Milestones'. 17 August 2022. <https://www.cdc.gov/ncbddd/actearly/milestones/index.html>
- Center on the Developing Child. 2007. 'A Science-Based Framework for Early Childhood Policy'. Center on the Developing Child at Harvard University. <https://developingchild.harvard.edu/resources/a-science-based-framework-for-early-childhood-policy/>
- . 2023. 'What Is Early Childhood Development? A Guide to the Science (ECD 1.0)'. *Center on the Developing Child at Harvard University* (blog). 2023. <https://developingchild.harvard.edu/guide/what-is-early-childhood-development-a-guide-to-the-science/>
- Center on the Developing Child, and National Forum on Early Childhood Policy and Programs. 2011. 'Building the Brain's "Air Traffic Control" System: How Early Experiences Shape the Development of Executive Function: Working Paper No. 11'. Center on the Developing Child at Harvard University. <http://www.developingchild.harvard.edu>
- Chisholm, Roderick M. 1964. 'The Descriptive Element in the Concept of Action'. *The Journal of Philosophy* 61 (20): 613. <https://doi.org/10.2307/2023443>
- Clark, Andy, and David Chalmers. 1998. 'The Extended Mind'. *Analysis* 58 (1): 7–19.
- Cortes Barragan, Rodolfo, and Carol S. Dweck. 2014. 'Rethinking Natural Altruism: Simple Reciprocal Interactions Trigger Children's Benevolence'. *Proceedings of the National Academy of Sciences* 111 (48): 17071–74. <https://doi.org/10.1073/pnas.1419408111>
- Danto, Arthur Coleman. 1981. *The Transfiguration of the Commonplace: A Philosophy of Art*. Vol. 40. Cambridge, Mass.: Harvard University Press.
- Davidson, Donald. 2001a. 'Actions, Reasons, and Causes'. In *Essays on Actions and Events*, 2nd ed, 3–20. Oxford: New York: Clarendon Press; Oxford University Press. <https://doi.org/10.1093/0199246270.003.0001>
- . 2001b. 'Mental Events'. In *Essays on Actions and Events*, 2nd ed., 207–24. Oxford: New York: Clarendon Press; Oxford University Press.

- . 2001c. ‘The Logical Form of Action Sentences’. In *Essays on Actions and Events*, 2nd ed., 105–21. Oxford: New York: Clarendon Press; Oxford University Press.
- . 2001d. ‘Freedom to Act’. In *Essays on Actions and Events*, 2nd ed., 63–82. Oxford: New York: Clarendon Press; Oxford University Press. <https://doi.org/10.1093/0199246270.003.0004>
- . 2001e. ‘How Is Weakness of the Will Possible?’ In *Essays on Actions and Events*, 2nd ed., 21–42. Oxford: New York: Clarendon Press; Oxford University Press. <https://doi.org/10.1093/0199246270.003.0001>
- Diamond, Adele. 2006. ‘The Early Development of Executive Functions’. In *Lifespan Cognition: Mechanisms of Change*, edited by Ellen Bialystok and Fergus I. M. Craik, 70–95. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195169539.001.0001>
- Dyson, Matthew. 2022. ‘The Contribution of Complicity’. *The Journal of Criminal Law* 86 (6): 389–419. <https://doi.org/10.1177/00220183221133439>
- Eckerman, Carol O. 1993. ‘Toddlers’ Achievement of Coordinated Action with Conspecifics: A Dynamic Systems Perspective’. In *A Dynamic Systems Approach to Development: Applications*, edited by Linda B. Smith and Esther Thelen, 333–57. The MIT Press. <https://doi.org/10.7551/mitpress/2523.001.0001>
- Evans, Gareth. 1982. *The Varieties of Reference*. Edited by John McDowell. Oxford [Oxfordshire]: New York: Clarendon Press; Oxford University Press.
- Foltz, Richard C. 2003. ‘Does Nature Have Historical Agency? World History, Environmental History, and How Historians Can Help Save the Planet’. *The History Teacher* 37 (1): 9–28. <https://doi.org/10.2307/1555594>
- Ford, Anton. 2018. ‘The Province of Human Agency’. *Nous* 52 (3): 697–720. <https://doi.org/10.1111/nous.12178>
- Francis, Leslie P. 2009. ‘Understanding Autonomy in Light of Intellectual Disability’. In *Disability and Disadvantage*, edited by Kimberley Brownlee and Adam Cureton, 200–215. Oxford: OUP.
- Frankfurt, Harry G. 1978. ‘The Problem of Action’. *American Philosophical Quarterly* 15 (2): 157–62.
- Gardner, John. 2007. ‘Complicity and Causality’. *Criminal Law and Philosophy* 1 (2): 127–41. <https://doi.org/10.1007/s11572-006-9018-6>
- Gibbard, Allan. 2008. *Thinking How to Live*. Cambridge (Mass.): Harvard university press.
- Gilbert, Margaret. 1990. ‘Walking Together: A Paradigmatic Social Phenomenon’. *Midwest Studies In Philosophy* 15 (1): 1–14. <https://doi.org/10.1111/j.1475-4975.1990.tb00202.x>
- . 2009. ‘Shared Intention and Personal Intentions’. *Philosophical Studies* 144 (1): 167–87. <https://doi.org/10.1007/s11098-009-9372-z>
- . 2014. *Joint Commitment: How We Make the Social World*. New York, NY: Oxford University Press.
- . 2017. ‘Joint Commitment’. In *The Routledge Handbook of Collective Intentionality*, edited by Marija Jankovic and Kirk Ludwig, 130–39. Abingdon: Routledge. <https://doi.org/10.4324/9781315768571-13>
- Girgis, Sherif. 2013. ‘The Mens Rea of Accomplice Liability: Supporting Intentions’. *Yale Law Journal* 123: 460–94.

- Gopnik, Alison, and Virginia Slaughter. 1991. 'Young Children's Understanding of Changes in Their Mental States'. *Child Development* 62 (1): 98–110.
- Grosse Wiesmann, Charlotte, and Victoria Southgate. 2021. 'Early Theory of Mind Development: Are Infants Inherently Altercentric?' In *The Neural Basis of Mentalizing*, edited by Michael Gilead and Kevin N. Ochsner, 49–66. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-51890-5_3
- Grünbaum, Thor. 2008. 'Trying and the Arguments from Total Failure'. *Philosophia* 36 (1): 67–86. <https://doi.org/10.1007/s11406-007-9096-2>
- Haddock, Adrian. 2005. 'At One with Our Actions, but at Two with Our Bodies: Hornsby's Account of Action'. *Philosophical Explorations* 8 (2): 157–72. <https://doi.org/10.1080/13869790500095939>
- Hammond, Stuart I., Ulrich Müller, Jeremy I. M. Carpendale, Maximilian B. Bibok, and Dana P. Liebermann-Finestone. 2012. 'The Effects of Parental Scaffolding on Preschoolers' Executive Function.' *Developmental Psychology* 48 (1): 271–81. <https://doi.org/10.1037/a0025519>
- Harman, Gilbert. 1976. 'Practical Reasoning'. *The Review of Metaphysics* 29 (3): 431–63.
- Hinshelwood, Alec James. 2018. 'Thought in Motion : An Essay on the Possibility of Practical Reason'. Doctoral Thesis, London: UCL.
- Hoare, Philip. 2023. 'Orcas Are Ramming Yachts off the Spanish Coast – Is the Whale World Rising Up?' *The Guardian*, 25 May 2023, Online edition. <https://www.theguardian.com/commentisfree/2023/may/25/orcas-ramming-yachts-spanish-whale-behaviour-trauma-humans>
- Hornsby, Jennifer. 1980. *Actions*. International Library of Philosophy. London: Routledge and Kegan Paul.
- Huizinga, Mariëtte, Conor V. Dolan, and Maurits W. van der Molen. 2006. 'Age-Related Change in Executive Function: Developmental Trends and a Latent Variable Analysis'. *Neuropsychologia* 44 (11): 2017–36. <https://doi.org/10.1016/j.neuropsychologia.2006.01.010>
- Hume, David. (1748) 2007. *Hume: An Enquiry Concerning Human Understanding: And Other Writings*. Edited by Stephen Buckle. Cambridge Texts in the History of Philosophy. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511808432>
- Hyman, John. 2017. *Action, Knowledge, and Will*. Oxford: OUP.
- Jaworska, Agnieszka. 1999. 'Respecting the Margins of Agency: Alzheimer's Patients and the Capacity to Value'. *Philosophy and Public Affairs* 28 (2): 105–38. <https://doi.org/10.1111/j.1088-4963.1999.00105.x>
- Jaworski, Peter M., and David Shoemaker. 2018. 'Me and Mine'. *Philosophical Studies* 175 (1): 1–22. <https://doi.org/10.1007/s11098-016-0853-6>
- Kaisermann, Alex. 2021. 'Against Accomplice Liability'. In *Oxford Studies in Philosophy of Law Volume 4*, edited by John Gardner, Leslie Green, and Brian Leiter. New York: Oxford University Press.

- Kampis, Dora, Charlotte Grosse Wiesmann, Sarah Koop, and Victoria Southgate. 2020. 'Understanding the Self in Relation to Others: Infants Spontaneously Map Others' Face to Their Own at 16 to 26 Months'. Preprint. PsyArXiv. <https://doi.org/10.31234/osf.io/ky35p>
- Khoury, Andrew C. 2013. 'Synchronic and Diachronic Responsibility'. *Philosophical Studies* 165 (3): 735–52. <https://doi.org/10.1007/s11098-012-9976-6>
- Khoury, Andrew C., and Benjamin Matheson. 2018. 'Is Blameworthiness Forever?' *Journal of the American Philosophical Association* 4 (2): 204–24. <https://doi.org/10.1017/apa.2018.17>
- Köhler, Sebastian. 2021. 'Moral Responsibility Without Personal Identity?' *Erkenntnis* 86 (1): 39–58. <https://doi.org/10.1007/s10670-018-0092-7>
- Kukla, Quill R. 2021. 'A Nonideal Theory of Sexual Consent'. *Ethics* 131 (2): 270–92. <https://doi.org/10.1086/711209>
- Kutz, Christopher. 2000. 'Acting Together'. *Philosophy and Phenomenological Research* 61 (1): 1–31. <https://doi.org/10.2307/2653401>
- . 2007. 'Causeless Complicity'. *Criminal Law and Philosophy* 1 (3): 289–305.
- Landry, Susan H., Karen E. Smith, Paul R. Swank, and Cynthia L. Miller-Loncar. 2000. 'Early Maternal and Child Influences on Children's Later Independent Cognitive and Social Functioning'. *Child Development* 71 (2): 358–75. <https://doi.org/10.1111/1467-8624.00150>
- Laurence, Ben. 2011. 'An Anscombian Approach to Collective Action'. In *Essays on Anscombe's Intention*, edited by Anton Ford, Jennifer Hornsby, and Frederick Stoutland, 270–96. Cambridge, Mass: Harvard University Press.
- Lavin, Douglas. 2013. 'Must There Be Basic Action?' *Notus* 47 (2): 273–301.
- . 2015. 'Action as a Form of Temporal Unity: On Anscombe's *Intention*'. *Canadian Journal of Philosophy* 45 (5–6): 609–29. <https://doi.org/10.1080/00455091.2015.1125589>
- Lewis, David. 1976. 'The Paradoxes of Time Travel'. *American Philosophical Quarterly* 13 (2): 145–52.
- Lillard, Angeline S., and John H. Flavell. 1992. 'Young Children's Understanding of Different Mental States'. *Developmental Psychology* 28 (4): 626–34. <https://doi.org/10.1037/0012-1649.28.4.626>
- Lindemann, Hilde. 2010. 'Holding One Another (Well, Wrongly, Clumsily) in a Time of Dementia'. In *Cognitive Disability and Its Challenge to Moral Philosophy*, 161–69. <https://doi.org/10.1002/9781444322781.ch9>
- . 2014. *Holding and Letting Go: The Social Practice of Personal Identities*. Oxford: OUP.
- Liszkai-Peres, Krisztina, Dora Kampis, and Ildikó Király. 2021. 'The Flexibility of Early Memories: Limited Reevaluation of Action Steps in 2-Year-Old Infants'. *Journal of Experimental Child Psychology* 203 (March): 105046. <https://doi.org/10.1016/j.jecp.2020.105046>
- Mackenzie, Catriona, and Natalie Stoljar. 2000a. 'Introduction: Autonomy Reconfigured'. In *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*, edited by Catriona Mackenzie and Natalie Stoljar, 3–31. Oxford: Oxford University Press. <https://doi.org/10.1017/CBO9781107415324.004>

- , eds. 2000b. *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. Oxford: Oxford University Press. <https://doi.org/10.1017/CBO9781107415324.004>
- MacMurray, John. 1961. *Persons in Relation*. Vol. 2. 2 vols. The Form of the Personal. Faber and Faber. <https://www.giffordlectures.org/lecturers/john-macmurray>
- Madden, Rory. 2011. 'Intention and the Self'. *Proceedings of the Aristotelian Society* CXI (PART 3): 325–51. <https://doi.org/10.1111/j.1467-9264.2011.00313.x>
- . 2014. 'Self-Consciousness and Its Linguistic Expression'. In *Communicative Action*, edited by Tzu-Wei Hung, 75–92. Singapore: Springer Singapore. https://doi.org/10.1007/978-981-4585-84-2_5
- Marcovitch, Stuart, and Philip David Zelazo. 2009. 'A Hierarchical Competing Systems Model of the Emergence and Early Development of Executive Function'. *Developmental Science* 12 (1): 1–18. <https://doi.org/10.1111/j.1467-7687.2008.00754.x>
- Marx, Karl, and Friedrich Engels. (1848) 2011. *The Communist Manifesto*. Translated by Samuel Moore. New York: Penguin Books.
- Matheson, Benjamin. 2014. 'Compatibilism and Personal Identity'. *Philosophical Studies* 170 (2): 317–34. <https://doi.org/10.1007/s11098-013-0220-9>
- Mayr, Erasmus. 2011. 'Deviant Causal Chains'. In *Understanding Human Agency*, 104–41. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199606214.001.0001>
- McDowell, John. 1996. *Mind and World: With a New Introduction*. 1st Harvard University Press paperback ed. Cambridge, Mass: Harvard University Press.
- . 2015. 'Acting As One Intends'. In *Philosophy of Action: An Anthology*, edited by Jonathan Dancy and Constantine Sandis, 145–57. Oxford: Wiley-Blackwell.
- McGeer, Victoria. 2018. 'XV—Intelligent Capacities'. *Proceedings of the Aristotelian Society*, October. <https://doi.org/10.1093/arisoc/aoy017>
- . 2019. 'Scaffolding Agency: A Proleptic Account of the Reactive Attitudes'. *European Journal of Philosophy* 27 (2): 301–23. <https://doi.org/10.1111/ejop.12408>
- Meltzoff, Andrew N. 1995. 'Understanding the Intentions of Others: Re-Enactment of Intended Acts by 18-Month-Old Children.' *Developmental Psychology* 31 (5): 838–50. <https://doi.org/10.1037/0012-1649.31.5.838>
- Miller, Edie (@ediemill). 2022a. 'if betty's hat falls off when she's on my back i stop a friendly looking person and ask them to put it back on her because i can't reach. this system has now broken down because she's learned that if she takes her hat off and throws it on the ground she gets to make a new friend'. Tweet. <https://twitter.com/ediemill/status/1595824736625651717>
- . 2022b. 'not saying my toddler knows how to use a pub but he did just wake up from his nap and head towards the bar shouting "i get.....drink!!!" Tweet. <https://twitter.com/ediemill/status/1539988435162877952>

- Moses, Louis J. 1993. 'Young Children's Understanding of Belief Constraints on Intention'. *Cognitive Development* 8 (1): 1–25. [https://doi.org/10.1016/0885-2014\(93\)90002-M](https://doi.org/10.1016/0885-2014(93)90002-M)
- Mourelatos, Alexander P. D. 1978. 'Events, Processes, and States'. *Linguistic and Philosophy* 2.
- Müller, Ulrich, and Kimberly Kerns. 2015. 'The Development of Executive Function'. In *Handbook of Child Psychology and Developmental Science*, edited by Richard M Lerner, 1–53. Hoboken, NJ, USA: John Wiley & Sons, Inc. <https://doi.org/10.1002/9781118963418.childpsy214>
- Munakata, Yuko, Hannah R. Snyder, and Christopher H. Chatham. 2012. 'Developing Cognitive Control: Three Key Transitions'. *Current Directions in Psychological Science* 21 (2): 71–77. <https://doi.org/10.1177/0963721412436807>
- O'Brien, Lucy. 2007a. *Self-Knowing Agents*. Oxford: OUP.
- . 2007b. 'What Are Actions?' In *Self-Knowing Agents*, Chapter 8, 129–155. Oxford: OUP.
- . 2017. 'Actions as Prime'. *Royal Institute of Philosophy Supplement* 80 (May): 265–85. <https://doi.org/10.1017/S135824611700011X>
- . forthcoming. 'I, Myself, Move'. In *Beings and Doings*.
- Onishi, Kristine H., and Renée Baillargeon. 2005. 'Do 15-Month-Old Infants Understand False Beliefs?' *Science* 308 (5719): 255–58. <https://doi.org/10.1126/science.1107621>
- O'Shaughnessy, Brian. 1973. 'Trying (As the Mental "Pineal Gland")'. *The Journal of Philosophy* 70 (13): 365–86.
- . 2008. *The Will: Volume 2, a Dual Aspect Theory*. Cambridge University Press.
- Oxford English Dictionary. 2023. 'Own, Adj. & Pron.' In Oxford: Oxford University Press. Oxford English Dictionary. <https://doi.org/10.1093/OED/9248683774>
- Peacocke, Christopher. 1979. *Holistic Explanation: Action, Space, Interpretation*. Oxford: New York: Clarendon Press; Oxford University Press.
- Pears, David. 1975. 'The Appropriate Causation of Intentional Basic Actions'. *Crítica: Revista Hispanoamericana de Filosofía* 7 (20): 39–72.
- Pennick, Katie (@KatiePennick). 2021. 'Fellow wheelchair users - I am curious! In general conversation about getting somewhere, do you say "I walked there"/ "I went for a walk". Or do you use wheelchair language? ("I wheeled there/ I went for a roll/ Shall we push to the shops?") (If other - comment below!) Walk 70.1% Wheel 12.5% Roll 12.7% Push 4.7%'. 8 January 2021. <https://twitter.com/KatiePennick/status/1347591732700061696>
- Perner, Josef, and Johannes Roessler. 2010. 'Teleology and Causal Understanding in Children's Theory of Mind'. In *Causing Human Actions*, edited by Jesús H. Aguilar and Andrei A. Buckareff, 199–225. The MIT Press. <https://doi.org/10.7551/mitpress/9780262014564.003.0103>
- Perner, Josef, Petra Zauner, and Manuel Sprung. 2005. 'What Does "That" Have to Do with Point of View? Conflicting Desires and "Want" in German'. In *Why Language Matters for Theory of Mind*, edited by Janet Wilde Astington and Jodie A. Baird, 220–44. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195159912.003.0011>

- Piaget, Jean. 1983. *The Moral Judgement of the Child*. Translated by Marjorie Gabain. New York: Penguin Books.
- Piñeros Glasscock, Juan S. 2021. 'The Puzzle of Learning by Doing and the Gradability of Knowledge-how'. *Philosophy and Phenomenological Research*, August, phpr.12832. <https://doi.org/10.1111/phpr.12832>
- Piñeros Glasscock, Juan S., and Sergio Tenenbaum. 2023. 'Action'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Spring 2023. <https://plato.stanford.edu/archives/spr2023/entries/action/>
- Pullum, Geoffrey K. 2018. 'Koko Is Dead, but the Myth of Her Linguistic Skills Lives On'. *The Chronicle of Higher Education*, 27 July 2018. <https://www.chronicle.com/blogs/linguafranca/koko-is-dead-but-the-myth-of-her-linguistic-skills-lives-on>
- Repacholi, Betty M., and Alison Gopnik. 1997. 'Early Reasoning about Desires: Evidence from 14- and 18-Month-Olds.' *Developmental Psychology* 33 (1): 12–21. <https://doi.org/10.1037/0012-1649.33.1.12>
- Richmond, Sarah. 2000. 'Feminism and Psychoanalysis: Using Melanie Klein'. In *The Cambridge Companion to Feminism in Philosophy*, edited by Miranda Fricker and Jennifer Hornsby, 1st ed., 68–86. Cambridge University Press. <https://doi.org/10.1017/CCOL0521624517.005>
- Rödl, Sebastian. 2015. 'Joint Action and Recursive Consciousness of Consciousness'. *Phenomenology and the Cognitive Sciences* 14 (4): 769–79. <https://doi.org/10.1007/s11097-015-9423-1>
- . 2016. 'Education and Autonomy'. *Journal of Philosophy of Education* 50 (1): 84–97. <https://doi.org/10.1111/1467-9752.12175>
- . 2018. 'Joint Action and Plural Self-Consciousness'. *Journal of Social Philosophy* 49 (1): 124–36. <https://doi.org/10.1111/josp.12226>
- Roessler, Johannes. 2020. 'Plural Practical Knowledge'. *Inquiry*, July, 1–20. <https://doi.org/10.1080/0020174X.2020.1787221>
- Rumfitt, Ian. 1994. 'Frege's Theory of Predication: An Elaboration and Defense, with Some New Applications'. *The Philosophical Review* 103 (4): 599–637.
- Ryle, Gilbert. 1945. 'Knowing How and Knowing That: The Presidential Address'. *Proceedings of the Aristotelian Society* 46: 1–16.
- . 2009. *The Concept of Mind*. London ; New York: Routledge.
- Satne, Glenda Lucila. 2020. 'Practical Knowledge and Shared Agency: Pluralizing the Anscombean View'. *Inquiry*, November, 1–28. <https://doi.org/10.1080/0020174X.2020.1837236>
- Scaife, M., and J. S. Bruner. 1975. 'The Capacity for Joint Visual Attention in the Infant'. *Nature* 253 (5489): 265–66. <https://doi.org/10.1038/253265a0>
- Schechter, Elizabeth. 2018. *Self-Consciousness and 'Split' Brains: The Minds' I*. First edition. Oxford, United Kingdom: Oxford University Press.
- Schechtman, Marya. 2007a. *The Constitution of Selves*. 1. print., Cornell paperbacks. Cornell Paperbacks. Ithaca: Cornell Univ. Press.

- . 2007b. ‘Stories, Lives, and Basic Survival: A Refinement and Defense of the Narrative View’. *Royal Institute of Philosophy Supplements* 60 (May): 155–78. <https://doi.org/10.1017/S1358246107000082>
- Schlosser, Markus. 2019. ‘Agency’. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2019. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2019/entries/agency/>
- Schmid, Hans Bernhard. 2016. ‘On Knowing What We’re Doing Together’. In *The Epistemic Life of Groups*, edited by Michael S. Brady and Miranda Fricker, 51–72. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198759645.003.0004>
- Schünemann, Britta, Lydia Paulin Schidelko, Marina Proft, and Hannes Rakoczy. 2022. ‘Children Understand Subjective (Undesirable) Desires before They Understand Subjective (False) Beliefs’. *Journal of Experimental Child Psychology* 213 (January): 105268. <https://doi.org/10.1016/j.jecp.2021.105268>
- Setiya, Kieran. 2003. ‘Explaining Action’. *Philosophical Review* 112 (3): 339–93.
- . 2007. ‘Explaining Action’. In *Reasons without Rationalism*, 21–67. Princeton University Press. <http://www.jstor.org/stable/j.ctt7sxdm.5>
- Shoemaker, David. 2012. ‘Responsibility Without Identity’. *The Harvard Review of Philosophy* XVIII: 109–32.
- . 2016. ‘The Stony Metaphysical Heart of Animalism’. In *Animalism: New Essays on Persons, Animals, and Identity*, edited by Stephan Blatti and Paul F. Snowdon, 1st edition, 303–28. New York, NY: Oxford University Press.
- Shoemaker, David W. 2007. ‘Personal Identity and Practical Concerns’. *Mind* 116 (462): 317–57. <https://doi.org/10.1093/mind/fzm317>
- . 2011. ‘Moral Responsibility and the Self’. In *The Oxford Handbook of the Self*, edited by Shaun Gallagher. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199548019.003.0022>
- Silver, Sarah. 2016. ‘Review: Koko: A Talking Gorilla’. *Reverse Shot* (blog). 9 November 2016. <https://reverseshot.org/reviews/entry/1472/koko>
- Small, Will. 2014. ‘The Transmission of Skill’. *Philosophical Topics* 42 (1): 85–111.
- Smith, Thomas H. 2011. ‘Playing One’s Part’. *Review of Philosophy and Psychology* 2 (2): 213–44. <https://doi.org/10.1007/s13164-011-0059-y>
- Somsen, Riek J.M. 2007. ‘The Development of Attention Regulation in the Wisconsin Card Sorting Task’. *Developmental Science* 10 (5): 664–80. <https://doi.org/10.1111/j.1467-7687.2007.00613.x>
- Southgate, Victoria. 2020. ‘Are Infants Altercentric? The Other and the Self in Early Social Cognition.’ *Psychological Review* 127 (4): 505–23. <https://doi.org/10.1037/rev0000182>
- Steward, Helen. 2012. ‘“Up-to-Us-Ness”, Agency, and Determinism’. In *A Metaphysics for Freedom*, 24–42. Oxford: OUP. <https://doi.org/10.1093/acprof>
- Thalberg, Irving. 1984. ‘Do Our Intentions Cause Our Intentional Actions?’ *American Philosophical Quarterly* 21 (3): 249–60.

- Thompson, Michael. 2008. *Life and Action: Elementary Structures of Practice and Practical Thought*. London: Harvard University Press.
- Tollefsen, Deborah. 2005. 'Let's Pretend! Children and Joint Action'. *Philosophy of the Social Sciences* 35 (1): 75–97. <https://doi.org/10.1177/0048393104271925>
- Tomasello, Michael, and Michael Jeffrey Farrar. 1986. 'Joint Attention and Early Language'. *Child Development* 57 (6): 1454. <https://doi.org/10.2307/1130423>
- Travis, Charles. 2013. *Perception: Essays after Frege*. First edition. Oxford: Oxford University Press.
- Tuomela, Raimo. 2005. 'We-Intentions Revisited'. *Philosophical Studies* 125: 327–69.
- UNICEF. 2022. 'Early Childhood Development'. *UNICEF Data* (blog). May 2022. <https://data.unicef.org/topic/early-childhood-development/overview/>
- Velleman, J. David. 1997. 'How To Share An Intention'. *Philosophy and Phenomenological Research* 57 (1): 29–50.
- . 2000. *The Possibility of Practical Reason*. Oxford: New York: Clarendon; Oxford University Press.
- Vendler, Zeno. 1967. 'Verbs and Times'. In *Linguistics in Philosophy*, 97–121. <https://doi.org/10.2307/2182828>
- Vygotsky, L. S. (1935) 2011. 'The Dynamics of the Schoolchild's Mental Development in Relation to Teaching and Learning'. Translated by Alex Kozulin. *Journal of Cognitive Education and Psychology* 10 (2): 198–211.
- Wallage, Martijn. 2020. 'Charles Travis on Truth and Perception'. *Canadian Journal of Philosophy* 50 (7): 878–89. <https://doi.org/10.1017/can.2020.36>
- Warneken, Felix, Frances Chen, and Michael Tomasello. 2006. 'Cooperative Activities in Young Children and Chimpanzees'. *Child Development* 77 (3): 640–63. <https://doi.org/10.1111/j.1467-8624.2006.00895.x>
- Warneken, Felix, and Michael Tomasello. 2006. 'Altruistic Helping in Human Infants and Young Chimpanzees'. *Science* 311 (5765): 1301–3. <https://doi.org/10.1126/science.1121448>
- . 2007. 'Helping and Cooperation at 14 Months of Age'. *INFANCY* 11 (3): 271–94.
- Wellman, Henry M., David Cross, and Julianne Watson. 2001. 'Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief'. *Child Development* 72 (3): 655–84. <https://doi.org/10.1111/1467-8624.00304>
- Wellman, Henry M., and Jacqueline D. Woolley. 1990. 'From Simple Desires to Ordinary Beliefs: The Early Development of Everyday Psychology'. *Cognition* 35 (3): 245–75. [https://doi.org/10.1016/0010-0277\(90\)90024-E](https://doi.org/10.1016/0010-0277(90)90024-E)
- Wiggins, David. 1976. 'Locke, Butler and the Stream of Consciousness: And Men as a Natural Kind'. *Philosophy* 51 (196): 131–131. <https://doi.org/10.1017/S003181910002057X>
- Willcutt, Erik G., Alysia E. Doyle, Joel T. Nigg, Stephen V. Faraone, and Bruce F. Pennington. 2005. 'Validity of the Executive Function Theory of Attention-Deficit/Hyperactivity Disorder: A Meta-Analytic Review'. *Biological Psychiatry* 57 (11): 1336–46. <https://doi.org/10.1016/j.biopsych.2005.02.006>
- Williams, Bernard, ed. 1981. 'Internal and External Reasons'. In *Moral Luck: Philosophical Papers 1973–1980*, 101–13. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781139165860.009>
- Williams, Glanville. 1983. *Textbook of Criminal Law*. 2nd ed. London: Stevens.

- Wilson, George, and Samuel Shpall. 2022. 'Action'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Winter 2022. Metaphysics Research Lab, Stanford University.
<https://plato.stanford.edu/archives/win2022/entries/action/>
- Wittgenstein, Ludwig. 2009. *Philosophical Investigations*. Edited by P.M.S. Hacker and Joachim Schulte. 4th ed. Oxford: Wiley-Blackwell.
- Wood, David, Jerome S. Bruner, and Gail Ross. 1976. 'The Role of Tutoring in Problem Solving'. *Journal of Child Psychology and Psychiatry* 17 (2): 89–100. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>
- Wood, David, and Heather Wood. 1996. 'Vygotsky, Tutoring and Learning'. *Oxford Review of Education* 22 (1): 5–16. <https://doi.org/10.1080/0305498960220101>