HOW MANY THEORIES OF ACT INDIVIDUATION ARE THERE?

by

Joseph W. Ulatowski

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Philosophy

The University of Utah

May 2008

ABSTRACT

The problem of action individuation is a debate about the identity conditions of human action. The fumdamental question about action individuation is: how do we distinguish between actions? By "we," action theorists have sought the ordinary conception of how people distinguish between actions. The aim of this dissertation is to show that discovering what ordinary intuitions about action individuation are may assist us in coming to terms with action theory.

Three views of action individuation have dominated the action theory literature. Donald Davidson and G.E.M. Anscombe have argued that a number of different descriptions may refer to a single action. Alvin Goldman has argued that each description designates a distinct action. Others, e.g., Irving Thalberg and Judith Jarvis Thomson, have averred that some acts are sequences of causally related events, which include both a primitive bodily action and some of its effects. These action theorists have assumed that a simple invariant account of action individuation captures how ordinary people distinguish between actions. In my dissertation, I devised an experiment to test the action theorists' assumption. My data show that people's intuitions seem to depend on the valence of the consequences of the action under consideration. So, a simple invariant account is not possible. In light of the

empirical results, I argue that if we seek a folk account of action individuation, then that account should be able to explain the variability that seems to be present in people's intuitions about different cases.

For Mom and Dad

TABLE OF CONTENTS

INTRODUCTION

Recent contributions in philosophy have challenged the presumption that the philosopher's own intuitions accurately represent ordinary people's intuitions. These contributions have used empirical assessments to mine ordinary intuitions. The results of these empirical assessments have been surprising because ordinary intuitions have failed to correspond with philosophers' intuitions about ordinary intuitions.

Many philosophical problems find their source in intuitions,[1] for example, "what is free will?", "what is knowledge?", or "what is intentional action?" Some philosophers, such as Edmund Gettier, use counterexamples to show how inadequate our theories of free will, knowledge, or intentional action are. For Gettier, his counterexamples had attempted to undermine theories of knowledge. His project, however, did not analyze concepts. Sometimes philosophers do not use counterexamples; instead, philosophers tend to engage in analyzing concepts (Jackson 1998).

---

[1] I ought to mention at the outset that some philosophers deny that we should depend on intuitions. For example, John Stuart Mill, especially in his work in the epistemology of logic, argues that the reliability of intuition is itself an *a posteriori* question; the belief that an appeal to intuition must rest on *a priori* properties is mistaken.

Conceptual analysis has tried to resolve complex problems by breaking them down into simpler parts. For example, Jonathan Bennett writes,

> Typically, [conceptual analysis] equates a thought with a conjunction of thoughts which are, in a sense, its conceptual parts; and parts must be simpler than the structured whole which they compose; so the analysis resolves a complex into items that are simpler. (Bennett 1995, 3f.)

Conceptual analysis addresses something reasonably close to our ordinary conception of some philosophically important topic (Jackson 1998, Gibbard 1990, Lewis 1972). By "our ordinary conception," these conceptual analysts mean what any ordinary person off the street would say. If the goal is to analyze folk concepts, one might expect philosophers to ask people for their opinions. But conceptual analysts claim that such polls are largely unnecessary.

The philosophy of action, or action theory, depends on intuitions. Just as in other areas of philosophical research, action theorists have assumed that they need not ask people for their intuitions.[2] As a result, they believe that their own intuitions are sufficient to systematize ordinary intuitions about action.

One aim of this project is to question the action theorists' assumption that they know what people's intuitions about action individuation are. I will use an empirical

---

[2] Alvin Goldman asserts that he wants to explain the common sense conceptual scheme in *A Theory of Human Action* (p. vi). There are other examples of action theorists calling on the common sense conceptual scheme in the literature, and I will devote some time to each of their explanations in Chapter 2.

assessment to test people's intuitions about individuating action. The results of my empirical assessment will show something very surprising. Action theorists have assumed too much about people's intuitions. People's intuitions fail to correspond with what action theorists have believed they are. Therefore, the problem of action individuation may be re-opened for further investigation.

Before I summarize the argument and structure of my dissertation, I will briefly outline what the problem of action individuation is. Moreover, I will explore the role of ordinary intuitions in action theory. If action theorists want a folk account of action individuation, then it is incumbent upon them to listen to the people.

The Problem of Action Individuation

A fundamental question in action theory is "what are actions?" This question suggests another: "how are actions distinguished?" The latter question raises the problem of action individuation.

An example will help clarify the problem of action individuation. Consider a rather bloody example: at a certain instance, Smith, who takes a gun, moves his finger, pulls the trigger, shoots Jones, and kills Jones. How many particular actions did Smith perform? Whether Smith's moving his finger is identical with shooting Jones will lead to an understanding of the nature of actions.

Let me begin by outlining two reasons why people have addressed the question of action individuation. First, interest in action individuation arose out of

concern for the nature of acts.[3] Alvin Goldman believes that an account of how many actions one performs "will provide a partial answer to the question of the nature of acts" (Goldman 1970, 1). Similarly, Judith Jarvis Thomson says that whether two act descriptions designate either distinct acts or the same action "will bring us to what seems to me a good place at which to begin an investigation into the nature of acts" (Thomson 1977, 14). Both Goldman and Thomson suggest that a solution to the problem of action individuation will lead to a satisfactory answer of the question "what are actions?"[4]

Second, action theorists had become dissatisfied with the accepted solution to the problem. The accepted solution argued that a variety of action descriptions designate the same action. To use Anscombe's example, suppose a man operates a lever that pumps water into a house; he does so by moving his arm, and by doing so he introduces poison into the water supply and poisons the inhabitants. On Anscombe's view, the descriptions "the man operated the lever," "the man moved his arm," and "the man poisoned the inhabitants" all designate the same action. Both Anscombe (1957/2000) and Davidson (2001, 43-61) concluded that "[the] welter of related descriptions corresponds to a single descriptum" (Davidson 2001, 59). Other

---

[3] In Chapter 1, I will consider the adequacy of this motivation for the problem of action individuation. I will give reasons for thinking that alternative motivations will assist us in thinking the problem of action individuation is one we ought not ignore. I will offer at least two motivations for taking up the problem of action individuation. The motivations will be outlined briefly in a later section of this introduction.

[4] There are many others that use the nature of acts to motivate a discussion of the problem of action individuation. For example, Hornsby (1980) and von Wright (1963, 1971) use the question about the nature of action as a motivational tool for discussing action individuation.

action theorists, however, such as Goldman and Thomson, were dissatisfied with Anscombe and Davidson's view because – for one – the view failed to appreciate that each description of the man's actions had different properties.

Goldman challenged, what they called, Anscombe and Davidson's "identity theory" of act individuation (Goldman 1970, 1; Goldman 1971, 761; Thomson 1971b, 775).[5] Goldman argued that we differentiate actions because "He replenished the water supply by operating the pump" is true, but "He replenished the water supply by poisioning the inhabitants" is false. The act of poisoning the inhabitants is different than the act of operating the pump. If the former proposition is true and the latter is false, the two sentences cannot identify the same descriptum. Goldman concluded that different descriptions designate distinct actions.

The Anscombe/Davidson "identity theory" of act individuation and Goldman's "multiplier" view represent two clear positions in the act individuation debate. Almost simultaneous to the time that Goldman devised his multiplier account a new view appeared. The new view does not have one coherent thesis around which it revolves. But it still is a view that some have considered a clear position to adopt.

Once there were clear alternatives in the debate, the discussion of act individuation stalled. The debate sputtered out when the middle view appeared in the work of Irving Thalberg and Carl Ginet. Carl Ginet says, "Though sufficiently interesting in its own right… there is no other significant question in the philosophy

---

[5] Other action theorists have alternative names for Anscombe and Davidson's account of action individuation. See chapter 4.

of action that depends on [a solution to the problem of act individuation]." (Ginet 1990, 70)[6] If the problem of act individuation is irrelevant to action theory, then there is no reason to continue arguing over the details of act individuation.

The irrelevance of act individuation causes Ginet to argue that the problem is merely a verbal dispute. Verbal disputes should not have any bearing on the debate about action (ibid.). Ginet contends that each of the theories of act individuation can find some support in our ordinary talk.

But he does not attempt to confirm his assumption that each of the theories will find at least some support in our ordinary talk about actions.

A gap in the literature remains since action theorists have not attempted to test empirically whether ordinary people do support each of the theories of act individuation. If testing ordinary people's intuitions about action individuation shows no clear support for one theory or other, then Ginet's assumption may be affirmed. The lack of empirical data should make us wary of Ginet's claim because other experiments have shown Gettier's counterexamples are insufficient (Swain, Alexander, and Weinberg 2005) or Kripke's causal-historical conception of meaning fails to apply to non-western cultures (Machery, Mallon, Nichols, and Stich 2004). If other experiments have shown the insufficiency of Gettier's counterexamples or Kripke's causal-historical theory of meaning, then perhaps an experiment will help us

---

[6] It should be noted that Alfred Mele (1997, 2) supports the view that the debate about action individuation is insignificant.

out with act individuation. Therefore, we should ask people for their intuitions about action individuation not to settle the matter but to bring us to a point where we can move forward in our analysis of the problem of action individuation.

The Folk Account of Action Individuation

Even if we ask people for their intuitions, Ginet argues that an account of act individuation that is consistent with ordinary folk intuitions would be too imprecise to be useful (Ginet 1990, 46-52). Ginet is not wrong, but his reluctance to ask people for their opinions is misplaced. What we should do is supplement our analysis of individuating actions with an empirical assessment of the folks' views. After all, action theorists want to analyze what everyone means by "action," not just what philosophers or experts mean by it. The supplementary information will provide us with data for a clearer analysis of our ordinary conception of "action."

I dismiss any action theorist who does not try to give a folk account of action or motivation-based theory because they are not concerned with ordinary intuitions. But I cannot dismiss Ginet's argument so easily since a foundation of his view assumes ordinary people would agree with *any* account of action individuation (Ginet 1990, 71). He suggests that they would agree with his own "concrete" account because it is "better supported," but he has no empirical support either for his preferred account or for any of the other accounts. Given that he has no evidence supporting his claim, his theory remains a target of this project.

Ginet argues that a theory of action individuation should be more refined and more regimented than anything demonstrable by ordinary talk.[7] Ginet writes, "If there are reasons for preferring [my] account, they are reasons, not so much for a judgment as to what we already plainly mean by an action, but for a decision as to what more precise thing we should mean when precision about the individuation of actions is called for" (Ginet 1990, 71). Ginet's account will be preferred by those who want what someone should mean by action when precision is important to them. Ginet's view implies that people sacrifice precision for pragmatic reasons.

There are cases where people generally do not sacrifice precision. For example, jurors who must determine the guilt or innocence of a defendant are sometimes required to be very precise in their reasons for judgment about a defendant's actions. Jurors may not have to explain what the necessary and sufficient conditions of an action are, but they may engage in discussion about actions and how to individuate them. For example, jurors may contemplate whether a defendant had "intentionally planned to kill" someone or had "defended himself." Determining the guilt or innocence of a person on trial for murder may require the jurors to look at evidence supporting one of these ways of individuating the defendant's actions. In some ways, ordinary people are precise in terms of debating about action.

According to Ginet, precise necessary and sufficient conditions are a part of any good theory of action individuation. He writes, "To establish that [a complex

---

[7] Ginet's *concrete account* gives a criterion of action individuation. Ginet explains that necessary and sufficient conditions for two different canonical action-designators to designate one and the same concrete action. More will be said of Ginet's concrete account in chapter 4.

action is a layered structure comprising core action and outer layer(s) of consequence or circumstance], I need to give a criterion of individuation for concrete actions. That is, I need to explain what is necessary and sufficient for it to be the case that two different canonical action-designators designate one and the same concrete action" (Ginet 1990, 65). He believes that people do not provide precise necessary and sufficient conditions for anything they seek to explain. In fact, ordinary people are unaware of the importance of necessary and sufficient conditions. Only experts, such as Ginet or action theorists, have the wherewithal to alert ordinary people of the importance of necessary and sufficient conditions. So, people are unreliable sources of information for an adequate or good theory of action individuation.

People are often imprecise when they try to explain some difficult concept. That people are imprecise in their explanations does not mean we should ignore them or their views, which is what Ginet wants us to do. People's intuitions should not be ignored because they may be informative. Moreover, if action theorists seek an account of what action is (which is what most action theorists are concerned with), they do not seek just what action is *for the learned philosopher* but for what action is – period.

People's views matter. People use terms like "action," "knowledge," "right," and "wrong." Their use of the terms works sufficiently well to get along in the world. The analysis of concepts like "action," "knowledge," "right," or "wrong" could be assisted if we were to learn how people use these concepts in their everyday talk. The folks' may assist us in working out the necessary and sufficient conditions for some

difficult concept, such as action. Analogously, the assessment of people's intuitions should be helpful in determining the necessary and sufficient conditions of an account of action individuation.

This argument treads on very controversial grounds. Many empirically-minded people (and not just experimental philosophers) have denied that concepts have the kind of structure that permits traditional analyses in terms of necessary and sufficient conditions (Cf. Fodor 1998, Rosch 1978, Stich 1996, Wittgenstein 1953). So, there may be reason to deny Ginet's account without moving into the domain of experimental philosophy.

Ginet's account becomes problematic from an experimental standpoint when he tried to establish that none of the accounts are better than any of the other accounts. First, he has claimed that all of the accounts are equally compelling and equally coherent. Second, he has tried to establish that each of the accounts could find support in ordinary talk about actions. The second is my sticking point.

Since he does not ask, there is no way of knowing whether people agree with one account or another. If we ask people for their intuitions and a statistically significant majority of people support one theory over another or if the assessment determines that people individuate action depends on the valence of the consequences of action, then Ginet has mistakenly asserted that a more regimented and more refined theory gives us something more than what ordinary talk offers. He would be wrong to suggest that ordinary people's intuitions are unimportant in the action individuation

debate, and he would be wrong to assume that such a survey would not be useful. To my mind, Ginet has missed an opportunity here.

Not all of this project will involve the misfortune of Ginet's assumption. In fact, the target of my experimental data is not Ginet. What I have tried to show in the last few pages is that Ginet – and other action theorists – have missed an opportunity to become better acquainted with what "action" is or how to distinguish between action when they assume that their own intuitions find support in ordinary talk.

Overview of the Dissertation

In my dissertation, I ask ordinary people for their intuitions on individuating action. Ordinary people and academic philosophers think about problems in significantly different ways precisely because people have not had the analytic training philosophers have. If ordinary intuitions differ from the philosopher's own intuitions, then the philosopher's own intuitions cannot reflect what the folk theory is.

The empirical data assists us in coming to terms with individuating actions. Folk intuitions enable us to arrive at a more informed understanding of action. Thus, we are better equipped to talk about the folk theory of action individuation once we have run some experiments testing for people's intuitions about action individuation.

Structure of the Dissertation

I have outlined what this project will address in the last few pages, so I would now like to discuss the structure of the project.

The first chapter will discuss motivations for undertaking an investigation of the problem of action individuation. Very few action theorists directly address the motivation question. Even fewer action theorists have an argument motivating their discussion of action individuation. The standard motivation is that action individuation will help us come to terms with the nature of action. I argue that the standard motivation for exploring action individuation may be supplemented by motivations we find outside action theory. I will show that there is an important reason to take up the problem that is derived from moral theory, particularly virtue ethics, from applied ethics, such as the doctrine of double effect, and from practical reasoning, e.g., planning theories of practical reason.

The second chapter will outline what experimental philosophy is. Experimental philosophy is a recent methodological addition to the philosophical landscape. Given that philosophy is an emerging sub-discipline within philosophy and given that I use the methods it employs in this project, some space should be allotted to discussing what it is. The first part of the second chapter will discuss what intuitions are and offer some general characteristics of experimental philosophy. Then I give a few advantages of experimental philosophy has over its methodological counterparts. Finally, I argue that the preferred type of experimental philosophy for this project is a form of what has been called by Alexander and Weinberg (2006) "the proper foundational view."

The third chapter will argue that empirical data should play a methodological role in philosophical research. Chapter three is largely a defense of experimental philosophy against its detractors. Since experimental philosophers have used empirical data to undermine philosophers' intuitions and the raw data must be analyzed for it to be useful in philosophical projects, I will outline why this development is important for my dissertation topic in particular and for philosophy in general.

The fourth chapter will summarize the dominant views in the act individuation debate. The summary will include a comprehensive outline of the alternatives. The summary will include the arguments of the *minimizing* view, the *maximizing* view, and the *componential* view. Also, the chapter will open with a brief historical look at the problem of action individuation. The problem seemed to arise in the work of Georg Henrik von Wright and H.L.A. Hart. Although these philosophers were not concerned with developing an account of act individuation, their studies of "action" generally paved the way for the three dominant views.

Finally, chapter five will discuss my experiment, including the methods and procedures used, and an explanation of the data accumulated. Philosophy dissertations (or projects for that matter) rarely use empirical data to support arguments or counterarguments. In this chapter, I will show how the collected experimental data assist me in reaching conclusions about action individuation. The final part of this chapter will be a bit speculative. In it, I will suggest that the data I collected may imply that how people individuate actions depends upon the valence of

the consequences of some action. When subjects were given a vignette, they judge that different descriptions refer to the same act when the consequence of the act is harmful. Furthermore, they judge that different descriptions refer to distinct acts when the consequence of the act is beneficial. What should be left for future research is an explanation of why we individuate actions based upon the valence of the consequences of action.

Chapter 1: Why Action Individuation Matters

A philosophical project should be compelling. A compelling project is well motivated. So, a philosophical project should be well motivated.

The problem of act individuation has been discussed without an explicit account of why it is an important problem for action theory. Action theorists have not sufficiently motivated the action individuation debate. The goal of this chapter is to motivate an investigation of the problem of act individuation. In this chapter, I contend with the motivations from moral theory and practical reasoning.

First, I will discuss the motivation action theorists have accepted. Most action theorists have suggested that action individuation must be discussed in order to devise an account of the nature of acts. The argument comes from Davidson's causal theory of action. If we can understand mental states as causes, then they are a non-mysterious part of the physical world. Beliefs and desires presumably cause actions. So we need a theory of action that allows us to see them in this role. This will

involve showing that there are actions and showing that they can be caused by the beliefs and desires that 'rationalize' them. The entailment relations between different action descriptions imply that each action description designates the same action. Subsequently, there are actions. This argument – or some variant of it – has motivated a large chunk of the literature about the problem of action individuation.

Some might resist the importance of that debate for a discussion of the individuation of action. Even action theorists may not be convinced that the nature of action question calls for an investigation of action individuation. The individuation of action will only tell us whether two action descriptions refer to the same or different acts. On this reading, act individuation seems ill-equipped to enlighten the debate on the nature of action. Even if a theory of action individuation informs the nature of actions, the motivation is persuasive for action theorists and not other philosophers or ordinary people.[8]

The remainder of the chapter unfolds by offering a few more motivations for taking up the problem of action individuation. First, I will contend that a motivation for exploring the problem of action individuation arises from interest in the doctrine of double effect. Then, I will show why planning theories of practical reasoning might warrant an exploration of action individuation. Next, I will argue that the unity of virtue, of the virtues, in virtue ethics provides impetus for the debate on the individuation of action. Finally, I will briefly consider some other motivations from

---

[8] Even action theorists may not be convinced that this is a good reason to take up the action individuation debate. Given that some might find the motivation uncompelling, I want to give a few other reasons for finding the debate compelling.

outside philosophy. One is technical and comes from biology, and the other is practical and comes from everyday life.

Chapter 2: What is Experimental Philosophy?

The second chapter introduces a new wave of philosophical research that uses social scientific methods to shed light on deep philosophical problems. The new wave of philosophical research has been called "experimental philosophy." Its primary purpose is to give philosophers a new outlet for coming to grips with ordinary intuitions.

The chapter will outline what experimental philosophy is. The first section of the chapter will discuss a few characteristics of experimental philosophy. Experimental philosophy's concern for ordinary intuitions calls for an exploration of what an intuition is. So, I will discuss what experimental philosophers have said about "intuitions." Then, I will introduce the two prominent projects within experimental philosophy. More will be said about each of these projects later in the chapter. One is the restrictionist view and the other is the proper foundational view, following on the terminology devised in Alexander and Weinberg (2006). In the final part of the first section, I offer a few advantages for undertaking an *experimental* philosophy.

The second part of the chapter will discuss the different types of experimental philosophy. Given that standard philosophical practice, e.g., conceptual analysis and

reflective equilibrium, appeal to intuitions as evidence for or against some philosophical claim, I argue that claims about the distribution of intuitions should be empirically informed since these claims are empirical claims. So, one type of experimental philosophy, the proper foundational view, attempts to supplement standard philosophical practice by providing a proper evidentiary foundation for certain philosophical claims. The other type of experimental philosophy, the restrictionist view, calls for the radical restriction of the employment of intuitions as evidence. The experimental evidence, according to this view, seems to point to the unsuitability of intuitions to serve as evidence at all. Whereas the restrictionist view attempts to convince us of giving up altogether on intuitions, the proper foundational view tries to uncover ordinary intuitions about philosophical problems so that we can use these intuitions to formulate a new – more empirically informed – understanding of the problem under consideration.

The final part of the chapter explores the advantages experimental philosophy has over its methodological competitors. I will suggest that the introduction of innovative methods into philosophy will contribute to greater creativeness among those interested in philosophy. Then I will argue that the challenge experimental philosophy poses to standard philosophical practice is not something that will harm philosophy because the methods philosophers use deserves critical scrutiny.

Chapter 3: Defending Experimental Philosophy

The third chapter challenges the traditional methods philosophers use, such as conceptual analysis and reflective equilibrium. Philosophers sometimes depend on ordinary intuitions to construct folk accounts of a concept. When they do, they assume the reliability of their own intuitions and they assume their own intuitions apply to ordinary people.

First, I will argue that recent empirical work demonstrates how philosophers' intuitions fail to represent ordinary intuitions. Some philosophers rely on intuitions to test philosophical hypotheses. A hypothesis is good insofar as it accords with our intuitions, and a hypothesis is bad insofar as it does not. A number of philosophers have used empirical methods to figure out what people think about particular hypothetical cases. Individuals are given questionnaires designed to elicit candid responses. Their responses have failed to support the philosophers' practice of appealing to intuitions. Thus, the empirical research fails to legitimize the philosophers *a priori* appeal to intuitions.

Second, empirical studies provide good evidence to criticize philosophers' *a priori* intuitions if the data show that the appeal to intuition is unrepresentative of the folks' views. The empirical evidence is good evidence of ordinary intuitions because it is gleaned from what people's intuitions are. Philosophers who attribute intuitions to people without checking risk ascribing their own intuitions to the folk. Philosophers need to remember that they seek a folk account. Folk accounts should accurately represent ordinary intuitions. Hence, we should support a philosopher who

has not surveyed people for their intuitions *only if* that philosopher has some other equally powerful empirical data.

Conceptual analysis is helpful if we want to know about a philosopher's intuitions, but it does not assist us in getting at ordinary intuitions. I will argue that conceptual analysis has failed to capture folk conceptions of mental properties. Those who engage in a bit of conceptual analysis are very unlikely to have incorporated empirical data at all. If philosophers make claims like, "ordinarily, people would say *x*…", then it seems counterintuitive not to check what people say.

The final part of the chapter challenges action theorists' employment of conceptual analysis. If action theorists rely upon conceptual analysis to devise an account of act individuation, then my use of the experimental method will challenge these action theorists' claims.

Views of action individuation have refused to ask people for their intuitions. Jonathan Bennett, for example, writes:

> [Questions of act individuation] are conceptual questions: no empirical work such as geographers or physicists or psychologists do is relevant to any of them. If you want… to know whether two events can fill the same region of space-time or the like, you must get your answer by thinking; there is no other way. Since your thinking will include reflection on your own patterns of thought and speech, it will be necessary to include empirical inquiry of a kind, but a kind that need not get you out of your armchair (Bennett 1988, 1).

If empirical research has shown that a philosopher's own intuitions do not track ordinary intuitions with any degree of regularity, then action theorists should consider getting out of our armchairs too. Even in action theory it is worthwhile to ask people for their opinions.

Chapter 4: Action Individuation: The Current State of Play

Chapter 4 reviews the current state of play in the debate on the individuation of action. First, I outline the work of two figures, G.H. von Wright (1963) and H.L.A. Hart (1948, 1949), who play a minor but significant role in the genesis of the problem of action individuation. Von Wright argued that action descriptions correspond to the agent's intentions. His reason for thinking so stems from a careful analysis of the distinction between the 'consequences' and the 'results' of an action. Hart argued that statements of the kind "Beatrice did *x*" are not descriptive, but ascriptive or attributive statements. Their primary function is to ascribe responsibility to an agent.[9] I show why von Wright's and Hart's theories were important for the development of the problem of act individuation, and I point out some shortcomings of each theory.

---

[9] Feinberg (1965) modified Hart's solution to the problem of action individuation. He modifies it to circumvent a problem that Geach and Pitcher noticed about the examples Hart used. Hart's examples attend to a statement of action which was equivalent to reproach. Hart himself later explicitly rejected the theses of his article (Hart 1973, Preface).

Second, I review the three dominant theories of action individuation: Davidson's (1963) and Anscombe's (1957) *minimizing* view, Goldman's (1970) *maximizing* view, and Thalberg's (1977) and Ginet's (1990) *componential* view.[10]

That there is but *one* act an agent performs which can be described in a variety of ways is called the *minimizing* view. For example, suppose that Smith does each of the following things: (1) Smith moves a certain pump up and down, (2) Smith pumps water into the house, and (3) Smith quenches the occupants' thirst. On the *minimizing* view, Smith performs only one action, of which three descriptions can be given.

Supporters of the *maximizing* view argue against the minimizers' claim that multiple descriptions designate the same action. Goldman, for instance, has averred that there is something that enters into the causal explanation of the action designated by (3) that does not enter into any causal explanation of the one designated by (1). It is true that Smith quenches the occupants' thirst, but it is not true that Smith moves the pump up and down *because* she quenches the occupants' thirst. If (3) and (1) are one and the same action, then one would expect them, if they cause anything at all, to cause the same set of events or states of affairs. They do not; so, nonequivalent descriptions represent distinct actions.

Finally, the *componential* view is the account of act individuation that there are multiple components of different actions because it sees the distinction between

---

[10] The authors cited are the most visible representatives of each view. In chapter 4, I will have to acknowledge the contributions of D'Arcy (1963), Shwayder (1965), Hornsby (1980), Davis (1970, 1979), McCann (1998) and Thomson (1971a, 1971b, 1977).

causal and non-causal generation as crucial. For example, quenching the occupants' thirst is caused by Smith's pumping the water into the house, (3) is caused by (2). Whereas the maximizing view implies that one action noncausally generates another, the componential view insists there is only one action. Smith's pumping the water into the house is the very same action as Smith's moving his hand up and down. This is contrary to the maximizing view. For the componential view, then, it is supposed to be possible for a single action to be of more than one type.

None of these accounts have captured ordinary intuitions about act individuation. The final chapter will show that is the case by using empirical data I collected. Ultimately, if what action theorists seek is an account of action individuation consistent with people's intuitions, then an account that incorporates empirical data will be a good means of generating that account.

Chapter 5: An Experimental Study of Individuating Actions

The fifth chapter reports the experiment itself and discusses the implications of the data collected. Accounts of act individuation have attempted to capture peoples' pre-theoretic intuitions of how to distinguish between actions. Proponents of the minimizing view have argued that a multitude of action descriptions designate only one act, while Goldman has averred that each action description refers to a distinct act. Following on studies by Joshua Knobe and others about intentional action, I subject these accounts of act individuation to experimentation. The data

indicate that people distinguish between actions differently depending upon the valence of the outcomes. Thus, the assumption that a single account of action individuation applies invariantly seems mistaken.

The experiment tests a subject's intuitions about act individuation. My hypothesis is that the lack of attention to people's intuitions shows that action theorists have not adequately considered the common sense conceptual scheme they had assumed to be correct. One experiment should suffice to assess whether this hypothesis is correct.

The experiment was inspired by an experiment Joshua Knobe used to test for ordinary intuitions about the concept 'intentional' (Knobe 2003a, 2003b, 2004). Knobe's experiment involves two vignettes. They are:

> **Harm**: The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, but it will also harm the environment."
>
> The chairman of the board answered, "I don't care at all about harming the environment. I just want to make as much profit as I can. Let's start the new program."
>
> They started the new program. Sure enough, the environment was harmed.
>
> **Help**: The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, and it will also help the environment."

> The chairman of the board answered, "I don't care at all about helping the environment. I just want to make as much profit as I can. Let's start the new program."
>
> They started the new program. Sure enough, the environment was helped.

In Knobe's original experiment, he distributes one of the two vignettes to a subject. He asks the subjects to answer the following question: "Did the chairman intentionally harm (help) the environment?" Since people tend to distinguish between actions performed intentionally and actions performed unintentionally, the difficulty is to uncover what this distinction consists in. His data suggests that moral considerations play an important role in folk psychology. People's intuitions about whether an outcome was intentionally produced seem to vary depending on the moral status of the outcome.

The verdict is still out whether Knobe's conclusion is correct; numerous studies have proposed alternative views, i.e., psychological bias (Malle 2006; Nadelhoffer 2004a, 2004b, forthcoming), individual difference and interpretive diversity (Nichols and Ulatowski 2007), or a Gricean account (Adams and Steadman 2004a, 2004b, ms).

While Knobe's work was at least partly an inspiration for my experiment, my work is distinct from Knobe's work because I am interested not in the folk concept of 'intentional' but in the folk concept of 'action'. I want to discover whether people intuitively believe that a multitude of descriptions refer to the same act or each

description designates distinct acts. I am not concerned with the intentionality of an agent's action.

In my experiment, I ask whether the agent, whose acts we can describe in different ways, performed one act or distinct acts. I remove the *intentionality* language. Just as in cases raised by Davidson, Goldman, Thomson, etc., e.g., does "Sirhan pulled the trigger" refer to the same action as "Sirhan killed Kennedy?" (Thomson 1977, 1), the question is whether two or more descriptions designate a single act or distinct acts.

The data I collected show an asymmetry in people's intuitions about act individuation. An asymmetry suggests that people sometimes accept that two action descriptions designate the same act. But, at other times, people reject that two action descriptions designate the same act. The latter suggest that people intuitively think that each action description designates a unique and distinct act. So, I believe that an explanation of people's intuitions about action individuation is that they somehow depend upon the valence of the consequences of the action.

Finally, I will try to predict the future of the debate and how my dissertation research will contribute to it. To my mind, the future of the debate includes further empirical research. Empirical research will inform what action description we ought to endorse. Supplemented by normative argument, the action individuation debate will contribute to action theory in particular and moral theory in general.

The aims of this dissertation

What I hope to show in the course of this dissertation is something modest. I will try to establish that there is a statistically significant variation in the subjects' responses. If there is statistically significant variation in the responses, then we will have to question the legitimacy of the armchair philosopher's appeal to intuitions and, more specifically, we will have reason to say that the currently accepted theories of action individuation have failed to account for individual differences in the population. When it comes to views on distinguishing actions, different people have different views about how to individuation action.

CHAPTER 1

WHY ACT INDIVIDUATION MATTERS[11]

Introduction

Sometimes the motivation for philosophical problems is absent or trivial. A good philosophy paper or project should be well motivated. The conventional motivation story of act individuation is somewhat unappealing. The problem of act individuation has been motivated by the nature of action, i.e., "what is action?" If the nature of action question is important, then individuating actions matters too. A part of moral philosophy consists in arguments about whether an agent's *action* is moral or immoral. Thus, what an action is is important for moral philosophy. Since it is important to know what an action is for problems in moral philosophy, a discussion about action individuation seems to be well motivated.

A better motivation may be derived from numerous sources. The aim of this chapter is to show that the problem of action individuation is more important than action theorists have considered in the conventional motivation story. My goal is not

---

[11] Thanks to Bob Barnard, Eric Hutton, Ron Mallon, Elijah Millgram, and Mark Timmons for conversations or correspondence leading to a draft of this chapter, and I am especially grateful to Anya Plutynski for alerting me to the inadequacies of the conventional motivation story.

necessarily to replace the conventional motivation story but to supplement it in a way that shows that the problem of act individuation is compelling. First, I will reconstruct the conventional motivation story action theorists have offered for action individuation. Then, I will contend that the conventional story can be supplemented by other motivations that overlap with concerns in other areas of philosophy, such as moral philosophy. Action individuation matters because: (1) the doctrine of double effect presupposes that we know how to distinguish between actions, (2) planning theories of practical reasoning could be improved if it were to commit itself to one or other account of action individuation, and (3) it could be analogous to other projects in closely related problem that arises in virtue ethics, the individuation of virtues (Cooper 1999). The list of motivations is not exhaustive, but it should provide sufficient evidence for the importance of the problem of action individuation. Finally, I will conclude that these alternative motivations seem to strengthen the reasons for exploring action individuation.

The "Conventional" Motivation Story

Historically, the problem of act individuation has had a single motivation story. The story is derived from the works of G.E.M. Anscombe and Donald Davidson.[12] In this section, I will reconstruct the conventional motivation story, and I

---

[12] Anscombe and Davidson's motivations for the problem of action individuation are slightly different. Since my primary concern in this chapter will be to clarify Davidson and his successor's motivations, I will limit discussion of Anscombe's motivations to this footnote. Anscombe is interested in discussing action individuation to elucidate how act descriptions may serve to specify intentions and also how

will provide some arguments that show why the conventional motivation story fails to generate much interest. The conventional motivation story is not necessarily wrong, but we should find some other motivations that improve our reasons for taking up the problem of action individuation. If there are stronger motivations available, then we should use them.

According to the conventional motivation story, the problem of act individuation is motivated by our interest in answering the question: "what is action?" because a solution to the problem of act individuation will lead to a solution to the nature of action. Whereas the problem of act individuation is concerned with distinguishing between actions, the nature of action is concerned with identifying actions and determining whether actions are events (or not), for example. A person coughs, sneezes, blinks, and blushes, and, in some sense, the person *has done* something. But sneezing, blinking, and blushing are different than walking across a room or operating a pump. Action theorists interested in the nature of action investigate the distinction between these two types of 'doings'.

Several action theorists have cited the nature of action as the reason for carrying out an investigation of action individuation. Alvin Goldman and Judith

---

intentions behind an act are related to one another. For example, when we ask the man why he moved his arm, he may respond, "I intended to operate the pump." If we ask him why he operated the pump, he may reply, "to replenish the water supply." And if we ask him why he replenished the water supply, he may respond, "to poison the inhabitants." Anscombe's position is that having "one action with four descriptions, each dependent on wider circumstances, and each related to the next as description of means to end… [permits us to] speak equally well of *four* corresponding intentions, or of *one* intention – the last term that we have brought in in the series" (Anscombe 2001, 46). It is such concern that motivates the problem of action individuation for Anscombe.

Jarvis Thomson have used the nature of action question to motivate their discussion of action individuation. Ultimately, a response to the problem of act individuation implies a solution to the nature of action question. For example, Alvin Goldman writes:

> What is an act? One of the problems concerning the nature of acts is the problem of individuation. Suppose that John does each of the following things (all at the same time): (1) he moves his hand, (2) he frightens away a fly, (3) he moves his queen to king-knight-seven, (4) he checkmates his opponent, (5) he gives his opponent a heart attack, and (6) he wins his first chess game ever. Has John here performed *six* acts? Or has he only performed *one* act, of which six different descriptions have been given? […] An answer to such questions will provide a partial answer to the question of the nature of acts. (Goldman 1970, 1)

A solution to the problem of action individuation, distinguishing between actions, presumably leads to an answer of the question of the nature of action. Goldman has asserted that the action individuation dispute should "lay bare the nature, or ontological status, of an act" (Goldman 1971, 768). An account of action individuation that fails to "lay bare" the nature of an act is not as comprehensive as one that does. According to Goldman, Davidson's account of action individuation has told us something about the ontological importance of action, i.e., actions are "terms in causal relations" (Goldman 1971, 769) but Davidson leaves out an analysis of ordinary action sentences. Since, according to Goldman, an analysis of ordinary action sentences will bring about an understanding of the nature of acts, an account of

action individuation that does not tell us about the nature of action – like Davidson's – is unsatisfactory.

Goldman's belief that the nature of action motivates a discussion of act identity carries over into the work of Judith Jarvis Thomson. Thomson has said that the identity or non-identity of actions, "will bring us to what seems to me a good place at which to begin an investigation into the nature of acts" (Thomson 1977, 14). So, on Thomson's view, interest in the ontology of action has motivated the problem of action individuation.

Proponents of a componential account of act individuation use a different motivation for the individuation project. Thalberg, Ginet, and McCann, componential theorists of action individuation, each have different motivations for their accounts. First, Hugh McCann has echoed and has somewhat extended Goldman and Thomson's motivation story. McCann has written "the problem of how action is individuated is of interest for general ontology, and for understanding the metaphysics of time and change" (McCann 1998, 1). At the very least, an acceptable ontology of action requires a principled response to the problem of action individuation. His explanation of a motivation for the problem of action individuation has included the metaphysical problem of time and ontology generally. So, McCann's motivation story involves the nature of action and other – more broadly construed – metaphysical problems, such as time and general ontology.

Second, Irving Thalberg writes

> I developed [an account of action individuation] not to uncover the essence of events generally, or action generally, but to reduce our bafflement over various kinds of occurrences which seem tied up with perception, emotion, and behavior. (Thalberg 1977, 9)

His account of action individuation is supposed to show how perception and emotion are bound up with what we do. Thalberg's component analysis of action

> permits us to say both that the agent performs only one deed, and that he carries out non-basic as well as basic actions (Thalberg 1977, 4)

without a paradox arising.

Finally, Ginet's componential account more or less has quashed the action individuation debate by insisting that it has nothing whatsoever to do with action theory generally. The problem might be interesting "in its own right," but it has no bearing on anything else in action theory (70f). The motivation for taking up the problem of act individuation was to fit it into his systematic theory of action generally.

The conventional motivation story has moved from act individuation to the nature of action, but the story has failed to explain how the two are related. Ginet raises this problem. The story has omitted a principled response to the following question: is the problem of action individuation related to the nature of action? If so, how are they related? Call a response to these two questions a "bridging argument." The omission of a bridging argument may make us skeptical of the story's

importance. If the problem of act individuation and the nature of action are not related (or not related in the right way), then we may doubt whether an investigation has been motivated adequately. So, it is important to have a clear understanding of how the two are related if we expect the conventional motivation story to be persuasive.

Davidson has supplied a bridging argument. His argument has attempted to "bridge" the gap between action individuation and the nature of action.[13] Goldman has pointed out that Davidson does not get us to fully satisfactory understanding of the nature of action. Nevertheless, his bridging argument is worthy of consideration.

In Davidson's causal theory of action, our actions are to be explained causally in terms of mental events or, more precisely, beliefs and desires. Since Davidson's causal theory of action is an attempt to work out a comprehensive view of action, Davidson seems obligated to tell us what actions are in terms of his causal account.

Davidson has derived what actions are from the logical form of action sentences. If there are entailment relations between different action descriptions, then what we have is a single action, which can be variously described. Davidson writes:

> But what is the relation between my pointing the gun and pulling the trigger, and my shooting the victim? The natural and, I think, correct answer is that the relation is that of identity. (Davidson 1968, 84)

---

[13] The causal theory of action was first introduced in Aristotle (see *Nichomachean Ethics* 1139a31-32). Since Aristotle's work is outside the scope of this dissertation, I will not address his theory of action.

Suppose that Plum shot and killed White. At least two action descriptions are possible for this event: "Plum killed White" and "Plum shot White." (There are, of course, many more possible descriptions of Plum's action, e.g., "Plum mortally wounded White" or "Plum murdered White.") The first and second descriptions have the same extensions. Any two descriptions with the same extensions refer to the same event. An action is an event. The two descriptions, therefore, refer to the same action.

Davidson's argument presents a logical form of action sentences to demonstrate that two or more descriptions refer to the same action. Suppose we have the following five sentences about Plum and White:

(1) Plum pointed the gun and pulled the trigger.
(2) Plum moved his right index finger.
(3) Plum pulled the trigger.
(4) Plum shot White.
(5) Plum obeyed an order.

An interest in presenting the logical form of sentences (1) to (5) and also

(6) Plum shot White with a gun.

forces us to admit some identities between acts can be differently described. Thus (6) entails (4); the logical form of these sentences makes this clear.

Davidson proposes the following logical form for (4) and (6):

(4') (∃*x*) (Shooting (Plum, White, *x*) and
(6') (∃*x*) (Shooting (Plum, White, *x*) & With (a gun, *x*).

The variable '*x*' ranges over events: (4') says there was an event that was a shooting of White by Plum, and (6') adds that this *same* event involved a gun. The entailment from (6') to (4') is straightforward. Likewise, the identity of Plum's shooting White with Plum's shooting him with a gun is straightforward because the two action descriptions are extensionally equivalent. Statements about actions and events have an underlying form expressible in quantifications over events.

Davidson's bridging argument shows that act individuation gives rise to a study of the nature of action. Just as we have learned from Quine that we should not tolerate entities unless we are prepared to make sense of sentences affirming and denying the identity of such sentences, we ought not to countenance actions without sentences that affirm and deny the identity of action-sentences (Quine 1969, 23). Thus, on Davidson's account, Quine's dictum, "no entity without identity" becomes "no action without identity."

Problems with the Conventional Motivation Story

I want to raise two problems for the bridging argument and the conventional motivation story. First, Davidson's bridging arguments depends on the idea that actions are events. If actions are not events as Kent Bach (1980) has argued, then Davidson's account is incorrect. Second, since Quine's dictum is controversial and

we have a reason for rejecting Quine's dictum, then, analogously, Davidson's modified version of Quine's dictum fails too. These problems are not insurmountable. My discussion will try to show how the conventional motivation story can be supplemented, not replaced.

First, Davidson's motivation story depends on a view of actions. According to Davidson, actions are a special subclass of events. Since they are a subclass of events and events are entities, an action is an entity as well. If actions are not events (or a special subclass of events), then an action is not an entity. We cannot speak of actions as entities, which is a presumption of Davidson's motivation story. So, we may reject the conventional motivation story because it assumes a specific view of action we need not endorse.

Kent Bach has argued that actions are not events. Bach has stated that an action is an "instance of the relation of bringing about between agents and events" (Bach 1980, 119). On this account, agents perform actions when they bring about a series of events, one after another. For example, Smith killed Jones by shooting the gun by pressing the trigger by moving his finger. Since, according to Bach, performing an action is equivalent to bringing about an event, it is not necessary to count actions but only the relations between agents and events. So, Bach has concluded that actions are not events.

Since actions are not events and Davidson's argument assumes actions are events, Bach argues against Davidson's account of act individuation. He writes:

> If actions are instances of a relation, we are not obliged to produce a theory of individuation of actions… Instances are not individual and not subject to quantification. Imagine the silliness of a debate over whether, in the case of a red, round ball, there is one instance or two. There is an instance of redness and an instance of roundness, but one and the same individual, the ball, is an instance of both. Rather than get mixed up in such a debate, we can simply say that one individual is an instance of both properties, that is, that the ball is red and round. Similarly, we need not worry about individuation in the case of actions, construed as instances of the relation of bringing about between agents and events. Since an action is performed if and only if someone has brought about an event, we need not count actions but only agents and events. (Bach 1980, 119)

Actions are said to take place at some point in time. They have relatively clear beginnings and endings but unclear spatial boundaries. This makes actions event-like. Similarly, actions and events appear to be homogenous in causal explanations. Actions can be causes of which events are effects (cf. Davidson 1967). Davidson believed that actions are events. Therefore, we can doubt that Bach's assessment has overturned Davidson's view that actions are events.

Next, Davidson's bridging argument presents a modified version of Quine's dictum. Just as Quine has said, "no entity without identity," Davidson says "no action without identity." P.F. Strawson has alerted us to alternative interpretations of Quine's dictum (Strawson 1997, 21-51). Just as Quine's dictum is ambiguous, so too is Davidson's formulation. Thus, we need not accept Davidson's bridging argument or the conventional motivation story.

On Strawson's view Quine's slogan "no entity without identity" can be interpreted in a wide sense and a narrow sense. First, according to the wide interpretation, everything that exists has identity conditions. If *x* exists, then, for any *y*, there is a determinate fact of the matter about whether *x* is identical to *y*. If we can make sense of sentences affirming or denying the identity of such entities, then the entity exists. On the wide interpretation, we do not need to be able to tell what the criteria for identity are. If interpreted liberally in this manner, then the slogan is *prima facie* acceptable.

The *prima facie* acceptability of the wide interpretation points to one of its deficiencies. If *x* and *y* are identical, then it is just that "*y*" is another name for "*x*"; "*x*" and "*y*" co-refer. In saying "*x* and *y* are identical," we merely say that *x* is identical to *x*. For example, when we say that Clark Kent and Superman are identical, we just mean Superman and Superman are identical. The statement is tautologous. Tautologous statements fail to give us new information. The lack of new information makes tautologous statements uninteresting. Therefore, a deficiency of the wide interpretation is that it is uninteresting.

Whereas the wide interpretation seems to be *prima facie* acceptable, the narrow interpretation severely limits the scope of the dictum. This narrow interpretation countenances that entities are identical if and only if they belong to a sort with clearly stateable general criteria of identity for *all* entities of a certain sort. For example, a criterion for material objects might hold that they are the same if and only if they occupy the same place at the same time. The identity criteria are fixed.

For every ontological category, we should be able to supply criteria of identity. This will tell us if *x* and *y* belong to the same category, and under what conditions *x* is identical with *y*.

The narrow interpretation would affirm that *x* and *y* are identical when *x* and *y* occupy the same place at the same time. To return to the Clark Kent / Superman example, we may say on the narrow interpretation that Clark Kent and Superman are identical because wherever (in space and time) Superman is Clark Kent is also.

When we discover that Superman has a property Clark Kent does not, the narrow interpretation is questionable. For example, Clark Kent is unable to fly, even though Superman is able to fly. The two might be identical but Superman possesses a quality Clark Kent lacks. Even if Superman and Clark Kent are identical in spatio-temporal terms, Superman's possession of a quality Clark Kent lacks may cause us to question the legitimacy of the narrow interpretation's call for criteria of identity based upon spatio-temporal boundaries.

The wide and narrow interpretations are different ways to understand Quine's slogan "no entity without identity." Both interpretations are inadequate. The inadequacies cause to doubt the narrow and wide interpretations of the slogan "no entity without identity." Davidson's modified version of Quine's dictum does not fare much better. So, there is no reason to accept Davidson's bridging argument. Ultimately, the bridging argument calls for further exploration. Hence, for Davidson's bridging argument and motivation story to be acceptable, we will have to how the bridging argument overcomes the problems of the narrow and the wide interpretation.

Davidson has not provided an argument to overcome the deficiencies of the two interpretations. Thus, we merely would be speculating how he would attempt to overcome the deficiencies of the bridging argument.

The conventional motivation story may not convince us of the importance of act individuation. In this section, I reviewed two worries about the conventional motivation story. Alternative motivation stories may provide us with further reasons for exploring the problem of action individuation.

The Doctrine of Double Effect and Action Individuation

The conventional motivation story has attempted to show that individuating actions play a role in discussions about the nature of action. In this section, I will offer an alternative reason for thinking the problem of act individuation is compelling. I will argue that the doctrine of double effect calls for an investigation of act individuation. First, I will provide a definition of the doctrine of double effect. In the course of the discussion, I will introduce an example. The example should help clarify why double effect compels us to explore act individuation.

The doctrine of double effect (hereafter "DDE") seeks to explain under what circumstances one may act when an action has both good and bad consequences. The

DDE is alternatively known as the Principle of Double Effect.[14] The DDE holds that there are stronger moral reasons against an agent intending harm as a means to achieving some benefit than there are against an agent merely foreseeing, but not intending, harm that comes about as a result of his actions to achieve some benefit. The moral permissibility of some action depends on whether the bad effect is intended, or merely foreseen or permitted to happen.

We should clarify the distinction between merely foreseeing harm and intending harm. A good explanation of the difference comes out of the work of St. Thomas Aquinas. Aquinas drew a distinction between merely foreseeing the bad effect and intending the bad effect in his discussion of whether it is morally permissible to kill someone in self-defense.[15] In his discussion, he asserted that there is nothing to prevent an act's having two effects. One effect may be intended, and the other merely foreseen. Aquinas writes:

> Nothing hinders one act from having two effects, only one of which is intended, while the other is besides the intention. Now moral acts take their species according to what is intended, and not according to what is beside the intention, since this is accidental… Accordingly the act of self-defense may have two effects, one is the saving of one's life, the other is the slaying of the aggressor. Therefore this act, since one's intention is to save one's own life, is not unlawful, seeing that it is natural to everything to keep itself in "being," as far as

---

[14] Two anthologies are good resources for articles on the doctrine (or principle) of double effect. See Woodward (2001) and Steinbock and Norcross (1994).
[15] Summa Theologica, II-II, Q. 64, Art. 7.

> possible. (Aquinas *Summa Theologica* II-II, Q. 64, Art. 7)

Suppose that Jones attacks Smith in such a way that he jeopardizes Smith's life. According to Aquinas, then, Smith may be morally justified in taking Jones's life if he kills Jones in self-defense. Smith's intent is to defend himself, not to kill. What is "beside the intention" in the example is Jones's death. The morality of an act is determined by what is intended, not by what is "beside the intention." Therefore, it may be morally permissible for Smith to ward off an attack, even if that involves killing Jones.

Aquinas never properly formulated the DDE, but discussed the importance of the distinction on which it turns. There are conditions or presuppositions that form the background of the DDE. The conditions or presuppositions that form the background of the DDE is suggested by Aquinas's work on the distinction between intending harm and merely foreseen harm. Besides the fact that the bad effect may not be directly intended, but only permitted to happen, it must also be the case that:[16]

1. the intended final end of the agent must be morally good or at least indifferent;
2. the intended means to the final end must be morally acceptable;
3. the foreseen bad effect that is expected to occur as a result of pursuing the good must not be intended; and
4. the beneficial final end must be important enough in order to justify bringing about the bad effect.

---

[16] These conditions are due in large part to Warren Quinn (1989).

Besides these four conditions, it is also important to remember (5.) the good effect must be produced directly (causally, not temporally) by the action and not by the bad effect. If the good effect were produced causally by the bad effect, then that would be using a bad means to achieve a good end, which is morally impermissible.

The DDE is typically invoked in situations that involve an agent in making a difficult moral decision.[17] Some situations are more frequently discussed than others. Perhaps the most often discussed situations invoking the DDE are the terrorist/tactical bomber and the mother/unborn fetus cases. [18]

In one version of the terror bomber case, a bomber has been instructed to bomb an airport full of civilians in order to demoralize the enemy. The bomber proceeds with the bombing. Subsequently, he kills the civilians in the airport, which maximizes his ability to demoralize the enemy. Proponents of the DDE would argue that the terror bomber intends to harm the civilians as a means to demoralize the enemy. Since the terror bomber deliberately involves the citizens in order to fulfill his purpose, the DDE would find the terror bomber's actions morally impermissible.

In one version of the tactical bomber case, a bombardier has been instructed to bomb an enemy forward observatory in order to destroy it. The bombardier intends to hit the observatory, but foresees that the bomb will very likely cause damage to

---

[17] Difficult moral decisions are not necessarily a moral dilemma; see Sinnott-Armstrong 1988 on the nature of moral dilemmas.

[18] The cases are discussed most often in the works of Philippa Foot, Warren Quinn, and Judith Jarvis Thomson.

nearby apartment buildings, possibly killing the civilians inside. Proponents of the DDE would say that the tactical bomber is merely foreseeing, but not intending to harm the civilians inside the apartment building. Since the tactical bomber merely foresees involving the citizens in order to fulfill his purpose, the DDE would deem the tactical bomber's actions morally permissible.

The distinction between the terror bomber and the tactical bomber case is that in one case the agent intends to harm the civilians and in the other the agent merely foresees the harm to civilians. The DDE provides us with a means of saying that the terror bomber's actions are morally impermissible because he intends to demoralize the enemy by bombing the civilians. The tactical bomber merely foresees the harm that will come to the civilians, and he does not intend to harm them. When an agent merely foresees the harm that will come about but does not intend the harm, it is morally permissible for the agent to perform the action, according to the DDE. So, according to the DDE, the terror bomber's actions are morally impermissible, whereas the tactical bomber's actions may be morally permissible (subject to the other conditions being met).

Two more cases will further clarify the DDE. Suppose that complications arise that threaten a pregnant woman's life. Doctors surmise that they must remove the fetus to save the mother's life. The only way to remove the fetus from the womb is to crush the fetus's skull. The doctors perform the craniotomy, and the fetus dies as a result of its head being crushed. In this case, proponents of the DDE would say that the doctors intended harm to the unborn fetus as a means to saving the mother's life.

Since they intended to harm the fetus, proponents of the DDE would find their actions morally impermissible.

In the fourth case, suppose that a pregnant woman's life is threatened by advanced uterine cancer. Doctors conclude that they must remove the uterus to save the woman's life. Since the uterus contains a fetus, the doctors would also remove it in the process. The fetus will surely die outside the womb. The doctors perform the hysterectomy, and the fetus dies as a result. Proponents of the DDE would say that the doctors merely foresee the harm that comes to the fetus, but they do not intend to harm it. Since they did not intend to harm the fetus, according to the DDE, the doctor's actions are morally permissible.

All of the cases seem to distinguish between those consequences that *intend harm* and the *merely foreseen* harmful consequences. A distinction between intended and merely foreseen consequences may affect how we distinguish between actions.

According to (5), one cannot act so that the bad effect is a means to the good effect. For this reason, proponents of the DDE have said that the craniotomy case is morally impermissible. Since the fetus's death is only an effect of the obstetrician's "wresting it from its mother," it is not itself an action. Since it seems that acts are the objects of moral assessment, the *means* to which (5) refers should be understood as an act. Evil means are evil acts. Bombing the civilians in terror bomber case would count as a means. But, in the craniotomy case, the fetus's death is an effect of the obstetrician's action, so (5) is not so strong as to rule out the craniotomy case for

proponents of the DDE. Both the DDE and especially condition (5) may be made less problematic as more sophisticated theories of act individuation develop.

The problem of act individuation seeks to distinguish between actions. According to Anscombe, it is of the essence that the DDE distinguish between the intended and merely foresee consequences of an action (Anscombe 1970, 42-53). The DDE contends that the boundary between intending harm and merely foreseeing it depends on the view that an action can have different consequences, i.e. a good or bad effect. If an action can be described in various ways, then there might be some ambiguity as to whether the consequences of some action were intentional or merely foreseen. This has prompted Jonathan Bennett to say, "There are various criteria for drawing the line between what someone did and the consequences of what he did; and there can be several proper ways of drawing it in a given case." (Bennett, 1966, 86) So, there are logical limits on what can be included in, or left out of, descriptions of intentional actions.

Proponents of the DDE would be wise to explore the dimensions of the act individuation debate. Perhaps on one account the craniotomy may be morally impermissible because it uses a bad effect as a *means* to a good effect. On another account the craniotomy may be morally permissible because the child's death is merely hastened by the doctor's actions.

The problem of act individuation and the DDE explore overlapping problems. Proponents of the DDE distinguish between foreseen bad side effects and intended harmful side effects. Proponents of action individuation make a distinction between

the actions that could result in harmful side effects. Therefore, an exploration of action individuation is a good idea would complement arguments for (or against) certain views of the DDE. Such an argument shows that the problem of action individuation is at least partially motivated by our concern over the DDE.

Plans, Practical Reason, and Action Individuation

Plans are helpful for the coordination of activities. A plan is a method of achieving one's goals that is developed to a sufficiently specified extent prior to the person's carrying out her action. Hasty decisions are avoidable if a person makes plans in advance. A plan is sufficiently specified if the plan's implementation will serve the person's interests better than any alternatives open to her, given the anticipated conditions in which the plan will be implemented. Successful coordination of a person's activities enable her to achieve complex goals she might not have been able to complete if she did not have a plan.

Two conceptions of plans are possible. First, according to Michael Bratman, plans are an "appropriate state of mind" (Bratman 1983, 272); for Bratman, then, a plan is a matter of "having a plan." "I have a plan to *A* only if it is true of me that I plan to *A*." (Bratman 1999b, 29) Bratman has argued that plans are conduct-controlling pro-attitudes. Although plans are in a similar category as beliefs and desires, they are distinct from them. They are distinct because intention should be categorized in terms of its function in rational action. Intention and related

phenomena are best understood as manifestations of the human capacity to plan and coordinate activities in advance. Bratman writes, "Intentions are… conduct-*controlling* pro-attitudes… The volitional dimension of the commitment involved in future-directed intention derives from the fact that intentions are conduct controllers. If my future-directed intention manages to survive until the time of action, and I see that that time has arrived and nothing interferes, it will control my action then." (Bratman 1999b, 16) Thus, a plan resembles beliefs and desires, in that it is a mental state of an agent, but plans are different than beliefs and desires since they exhibit a stability for which belief-desire models cannot account.

A plan does not have to be like Bratman's intention-based account. A plan could be an abstract structure represented by a set of instructions or a procedure.[19] For the purpose of this section, plans may be understood as either an appropriate state of mind or an abstract structure represented by a set of instructions.

We use plans to coordinate our activities in order to achieve our goals. On the "standard view" of planning, a rational agent will choose to perform some action if he is rational and the proposed course of action serves his interests better than any alternative open to him. For example, suppose Frank, an information technology director at XYZ Corporation, is interested in deploying new software, and he wants to deploy the software today. Frank adopts the plan to install the software on each of the company's computers himself. But, when the time for action arrives, Frank's

---

[19] Michael Bratman has discussed the distinction in his (1999a), p. 28f.

concerns would be served best if he did not install the software on each of the company's computers himself. Frank's concerns would be served better if he built a script that automatically installed the new software on each of the company's computers upon re-boot instead. The standard view says that Frank should abandon his intrapersonal plan to install the new software himself and he adopted the plan to write script that will automatically install the new software on re-boot.

Frank's original plan can be abandoned and a new one adopted without any problem arising. Changing plans are not always that simple. When we are confronted by cases in which the optimal course of action that would serve the agent's concerns best the standard view of planning may be undermined, or at least called into question.

The standard view has been challenged by cases in which it is not clear what course of action would serve the agent's concerns best. Cases of temptation, like Quinn's puzzle of the self-torturer (1993) or Kavka's toxin puzzle (1983), seem to conflict with the standard view. Quinn's self-torturer case is:

> Suppose there is a medical device that enables doctors to apply electric current to the boddy in increments so tiny that the patient cannot feel them. The device has 1001 settings: 0 (off) and 1… 1000. Suppose someone (call him the self-torturer) agrees to have the device, in some conveniently portable form, attached to him in return for the following conditions: The device is initially set at 0. At the start of each week he is allowed a period of free experimentation in which he may try out and compare different settings, after which the dial is returned to its previous position. At any other time,

> he has only two options – to stay put or to advance the dial one setting. But he may advance only one step each week, and he may *never* retreat. *At each advance he gets $10,000*.
>
> Since the self-torturer cannot feel any difference in comfort between adjacent settings, he appears to have a clear and repeatable reason to increase the voltage each week. The trouble is that there *are* noticeable differences in comfort between settings that are sufficiently far apart. Indeed, if he keeps advancing, he can see that he will eventually reach settings that will be so painful that he would then gladly relinquish his fortune and return to 0. (Quinn 1993, 198)

It seems unclear what it is rational to do in temptation cases.

Chrisoula Andreou (2006) has argued that recognizing actions as a part of larger actions or courses of action is important for practical deliberation in cases of temptation. According to Andreou, whether an action serves an agent's concerns well depends on what action or course of action it is a part of. She uses the following example:

> Tanya is having fun at a party. A friend has ordered another round of tequila. She has already had two generous shots but is tempted to have another despite her plan to stop at two; she realizes, however, that having another tequila will make all the difference in terms of how she will feel tomorrow. If she does not have another tequila, she will feel a bit groggy when she wakes up tomorrow, but will feel just fine after a few glasses of water. Though she is tempted to have another shot, Tanya does not consider the fun of doing one more round tonight to be worth the extra suffering she will experience tomorrow.

Andreou admits that the case is not straightforwardly one of temptation, but it serves to clarify what her argument is. Tanya does not want to be sick in the morning. We will call this her plan. According to Andreou, Tanya's drinking another shot serves her concerns well only if it is a part of an action or course of action that serves her concerns equally well. Since drinking another shot would cause her to be ill in the morning and cause her to be ill in the morning, Tanya's having a third tequila shot is not a part of her larger plan not to be sick in the morning. So, on Andreou's account, if Tanya is rational and it is clear at the time of action that having another tequila will not serve her concerns well, then she will not drink the third shot of tequila.

Andreou's proposal is consistent with the standard view, and I believe that her argument raises some interest in action individuation. Should I be correct, a problem in practical reasoning motivates our investigation of act individuation.

When Andreou suggests that an action, *x*, is a part of another action or course of action, *y*, we should attempt to understand what it means for *x* to be a part of *y*. The literature on act individuation's priority is to determine when *x* and *y* designate the same act, distinct acts, or arguing that *x* is a part of *y*. That some particular act serves one's cumulative plan may depend on one's view of *x*'s being a part of *y*. Consider the Tanya case again. We will say that Tanya's cumulative plan is *not to be sick in the morning*. If she drinks the third tequila, then it is likely that she will be sick in the morning. If she does not drink the third tequila, then it is likely that she will not be sick in the morning. So, it appears that Tanya's having the third drink will not fit her cumulative plan not to be sick in the morning because such an act will

result in her becoming sick. The conclusion assumes that Tanya's drinking another shot and her being sick in the morning are intimately related.

Two things have to be shown for Andreou's analysis to work. It has to be shown that Tanya's having another shot is *the same as* her being sick in the morning, or is *a part of* her being sick in the morning. If it turns out that Tanya's having another shot *refers to a completely distinct act from* her being sick in the morning, then the two are conceptually unrelated, even if her drinking another shot brings about her becoming sick.

We may not think Tanya's having another shot and her becoming sick are completely distinct because it is odd to say that having another shot did not bring about her sickness in the morning. We may contend that the two are related because the former brings about the latter, even if the action theorist contends act descriptions designate distinct actions.

The two descriptions refer to distinct acts when each description exemplify different properties. In Tanya's case, this seems like a semantic point, not a metaphysical one. The action theorist who contends that two or more act descriptions designate distinct acts because the two actions under consideration are distinct metaphysically speaking will argue that Tanya's case does not rule out potential counterexamples. For example, Tanya's drinking tequila last week lead her to sleep with Johnny. She is pregnant. So, her sickness in the morning is brought about by her pregnancy, not by having that last tequila. Or perhaps Tanya ate soft-shelled crab before her evening of tequila drinking. Her sickness could have been brought about

by her eating something that disagreed with her digestive system. Given that we cannot rule out these counterexamples, on the view that act descriptions designate completely distinct acts, Tanya's having a third drink is not necessarily a part of her becoming sick the next morning.

If we believe that Tanya's having the third tequila is the same as her being sick in the morning, then more perplexing temporal problem arises. Suppose Tanya performs one action, of which there can be a multitude of descriptions. If "Tanya's having the third drink" and "Tanya's being sick in the morning" refer to the same action, then Tanya is ill before she becomes sick in the morning. Similarly, Tanya's having that third drink continues to happen when she's becoming sick in the morning. This is odd for two reasons: (1) Tanya's drinking the tequila happens several hours before she's sick and (2) Tanya's illness occurs several hours after she stops drinking tequila.

Second, we can suppose that Tanya's having a drink and Tanya's becoming sick in the morning are distinct. If Tanya's being sick in the morning is distinct from her having the drink the night before, then her concern for not wanting to be sick in the morning has very little to do with having a drink the night before. This is something Andreou's account wants to avoid because she wants to show how Tanya's concerns are served best by the action(s) or course(s) of action that are part of her cumulative plan. Andreou may have shed more light on intrapersonal decision-making (and interpersonal decision-making), but her analysis calls for further analysis of how we distinguish between actions.

If Andreou's proposal is correct and there is a way of accommodating our considered judgments about the relevant cases of temptation without surrendering the standard view, then – according to the brief discussion above – it seems to promote further discussion of what it means for an action "to be a part of" another action or course of action. The investigation of what it means for one action "to be a part of" another action is at least one way of characterizing the problem of action individuation.

The Unity of Virtue and Action Individuation

The previous section advances a motivation for discussing act individuation based upon planning theories of practical reasoning. In this section, I will show why the unity of virtue is another motivation for the exploration of the problem of action individuation.

The discussion of the unity of virtue developed in response to the idea that "Virtue is one." Ancient Greek philosophers, such as Socrates, Plato, Aristotle, and the Stoics, subscribed to the unity of virtue. According to the unity of virtue, the possession of one virtue is necessarily related to the possession of all the others (Cf. Penner 1973).

There are two versions of the unity of virtue thesis. First, the *stronger* version of the thesis states that virtue is (*quite literally*) one. Our names for all the apparently different virtues refer to different aspects of the same single property. Accordingly,

any one of virtues is identical to all of the other virtues. "Bravery," "temperance," "courage," and "wisdom," for instance, are four different names of the same thing. Not only are there brave men but also there is such a thing as bravery.

Second, the *weaker* version of the thesis states that virtues are so integrated with one another that a person cannot have one without having all the others. One cannot be truly courageous unless one is also wise, as well as temperate, truthful, friendly, and so on.

The stronger version of the unity of virtues thesis entails, but is not entailed by, the weaker version of the thesis. That bravery *is* wisdom and wisdom *is* temperance, etc. (identity), entails that men are brave if and only if they are wise and men are brave if and only if they are just, and so on (equivalence). But the converse cannot be true since the identity statements carry with it certain ontological implications the equivalence statements do not. Call proponents of the weaker thesis *equivalence theorists*, and call proponents of the stronger thesis *identity theorists*.

One may deny the stronger thesis by arguing that the meaning of "bravery," for example, is not synonymous with the meaning of "wisdom." Since the two virtues do not have the same meaning, bravery and wisdom – along with the other virtues – cannot be identical.

Proponents of the weaker thesis have argued that what some virtue theorists, i.e., Aristotle, mean by the doctrine of the unity of the virtues must be equivalence. On the weaker account, when one asks of some virtue, *x*, "what is *x*?", the interpreters

have believed this to be a request for meanings of the term – an exercise in conceptual analysis.

The debate between the identity theorists and the equivalence theorists is about the parts of virtue. The proponents of each view have tried to determine how closely related is the possession of one virtue to the possession of another virtue. What I want to show now is that the two positions in the unity of virtues debate motivate an exploration of action individuation.

On Aristotle's account, having a virtue is a matter of having a character disposed to do the right thing in the right way, at the right time, to the right person, for the right reason. An individual can only have such a virtue if she has a sense of what the right action and the right reason are. (I should note that on one reading this may be too strong. For an individual to act on the right reason, the person may not have to know what the right reason is or how to act on it.) One needs to have the right values and the right priorities. Knowing what the right values and priorities are requires the knowledge of the importance of everything else against which it may in principle have to be balanced. If an action may be interpreted in different ways, then the person should have an explanation for why the action should be interpreted in a way that reflects the right values and the right priorities. Part of interpreting an action is being able to distinguish it from other – possibly, though not necessarily, distinct – actions.

A person should know how to distinguish between actions if she is to know what the right action is. Act descriptions are our best ways of interpreting whether an

action is virtuous. Any action could have a number of descriptions. The right action should have at least one description that designates it. For example, Eliza knows that her action of "laying down suppressive gunfire" to protect her comrades from hostile enemy fire exemplifies courage because she can distinguish her act from "running away from the battlefield" (cowardice) and from "charging the enemy position" (brashness). The act of cowardice and brashness are outside the scope of actions we would say are courageous. We understand this because we can distinguish between Eliza's actions of "laying down suppressive gunfire," "running away from the battlefield," and "charging the enemy position." So, if the possession of one virtue requires at least the knowledge for the possession of every other virtue, then the action under one virtue requires at least the knowledge of why the other actions under consideration fall outside that virtue.

Other Motivations

Problems of individuation often arise in technical and practical affairs. Sometimes these problems will persuade us to do or to refrain from doing something. That problems of individuation involve our deliberating about the optimal course of action may lend credence to the claim that the problem of action individuation is important for action theory. In this brief section, I will review two areas in which the problem of individuation comes about and suggest that these problems help motivate a discussion of action individuation.

First, biologists sometimes conduct experiments on animals within a particular habitat. The results of the biologists' experiment may be intended to apply to that habitat at other times than when the experiment first took place. The way that biologists describe the habitat may influence whether other biologists think that the experiment performed in one habitat applies to that habitat at a later time.

Whether an experiment is considered to be correct may depend on how the zoologist or biologist individuates habitats. Someone who thinks that all habitats are different and – subsequently – that each habitat description refers to a distinct habitat will probably believe that the research involving one habitat will not concern that habitat at a different time. Any two descriptions of a habitat may designate different habitats, even if their geographic location and ecological footprint is the same. Some biologists, however, may think that one habitat can have multiple descriptions. The biologist, thus, may be inclined to accept that those descriptions apply to the habitat at a different time. Biologists may accept or reject the results of an experiment based upon how they individuate habitats.

Biologists' views of how to individuate habitats may affect how they will describe others' research. Describing others' research will involve a description of their actions. So, their view of another's actions will be affected by the way they individuate environments.

Second, the problem of individuation also affects practical affairs. Suppose that Smith wants to go to Bi-Lo on Yadkin Road. He asks Jones for directions. Jones tells Smith he does not know of a Bi-Lo on Yadkin Road, but he knows of one on

U.S. 87. Jones gives Smith directions to the Bi-Lo on U.S. 87. As Smith makes his way to Bi-Lo, he discovers that U.S. 87 *is* Yadkin Road. U.S. 87 and Yadkin Road designate the same road. What Jones did not know was that U.S. 87 and Yadkin Road refer to the same street.

Jones did not know that U.S. 87 and Yadkin Road refer to the same street. Jones thinks that U.S. 87 is a highway, whereas Yadkin Road is not a highway. He may think that the two are distinct streets because one has a property the other does not. Jones would say of Smith that he's going to the Bi-Lo on U.S. 87, but he would not say of Smith that he's traveling to the Bi-Lo on Yadkin Road. Jones's description of Smith's action depends on whether he knows how to distinguish Yadkin Road from U.S. 87.

Jones will not give Smith bad directions because the two are not really distinct. In Jones's mind, however, the two are distinct because each has different properties. That he believes the two have distinct properties means he will say of Smith's "going to the Bi-Lo on U.S. 87" is true and he will say of Smith's "going to Bi-Lo on Yadkin Road" is false. The way that Jones divides the streets will affect how he individuates Smith's actions. So, in this way, something as simple as giving directions may motivate a discussion of how we individuate actions.

Conclusion

The beginning of the chapter outlined the conventional motivation story. The last-half of the chapter provided further reasons for thinking the problem of action individuation is important. The goal of this chapter will have been met if it has supplied further motivation for the problem of action individuation.

First, the conventional motivation story uses a Quinean idea that has its own problems. It could be interpreted in at least two ways. None of the action theorists have explicitly endorsed either interpretation. We should wonder which interpretation is correct. Second, the conventional motivation story depends on a controversial view of actions. According to Davidson, actions are events. Bach's argument against the broadly Davidsonian view that actions are events sheds light on a possible reason to reject Davidson's position. Finally, action theorists have appeared to take Davidson's motivation for pursuing the individuation of action as their own. We should dedicate some time to discover other motivations. If we find that there other motivations, then those motivations could help Davidson's conventional motivation story. So, on these three counts, we should look for other motivation stories that complement the conventional motivation story.

The last half of the chapter has been dedicated to supplying some compelling reasons for exploring action individuation. First, the DDE uses descriptions of an agent's actions in order to determine whether the action is morally permissible. Since an action may be described in myriad ways, it matters for DDE whether the description refers to one and the same act or distinct acts. A theory of action

individuation will enable proponents of the DDE to better understand whether some action is morally permissible or impermissible..

Second, the literature on the unity of virtues has discussed how to individuate virtues. Some accounts defend an identity view where all of the apparently different virtues refer to different aspects of the same single property, while others defend an equivalence view where virtues are so integrated with each other that a person cannot have one virtue without having all the others. Since having a virtue is a matter of having a disposition to act in accordance with virtue, the morally virtuous person will be disposed to choose one action over another. Agent's must be able to distinguish between actions in order to choose one over any of the others. Distinguishing between actions is possible if the agent can individuate actions. So, individuating virtues motivates a discussion of individuating action.

Third, plans assist us in coordinating our activities. Sometimes larger plans are composed of smaller plans. The smaller plans are a part of the larger plans. All of the plans include actions. When an action is a part of a plan, either we could mean that the action and the plan refer to the same thing or we could mean that the action is a part of the plan, where the action and the plan are distinct. The difference in the two positions is a matter of formulating a position in action individuation. Planning theories of practical reasoning seem to promote a discussion of action individuation.

This chapter has tried to provide a groundwork for the remainder of the dissertation. The groundwork of the dissertation is to develop a sufficient motivation for exploring the issues the rest of the work will undertake. Since the rest of my work

depends on implementing a special methodology, I will now turn to a defense of that strategy. The next chapter seeks to defend the social scientific methods experimental philosophy employs.

CHAPTER 2

WHAT IS EXPERIMENTAL PHILOSOPHY?[20]

Introduction

There is a movement in philosophy that is quickly becoming a way to open new avenues in long-standing philosophical disputes. Its proponents (as well as its detractors) call the movement *experimental philosophy*.[21] What is characteristic of the movement is the application of social scientific methods to philosophical problems.

The aim of this chapter is to outline experimental philosophy. First, I will discuss a few characteristics of experimental philosophy. Part of the discussion will include: (1) what an intuition is, (2) explore the characteristics of two prominent projects within experimental philosophy, and (3) offer a few advantages of experimental philosophy. Next, I will distinguish between at least two different

---

[20] Thanks to Eric Amsel, David Chalmers, Terry Horgan, Joshua Knobe, Clayton Littlejohn, Ron Mallon, Elijah Millgram, Matthew Mullins, Shaun Nichols, Ashlee Rogers, and, especially, Bob Barnard, Richard Greene, and Rachel Robison, as well as audiences at Weber State University and the University of Mississippi for their comments and questions on an earlier draft of this (and the next) chapter.

[21] Several players in the metaphilosophical debate about experimental philosophy will be mentioned in this chapter. Prominent among them are Joshua Alexander and Jonathan Weinberg (2006), advocates of experimental philosophy, and Antii Kaupinnen (2007), a critic of experimental philosophy.

branches of experimental philosophy. Following Joshua Alexander and Jonathan Weinberg (2006), I will call one branch the *proper foundational view* and the other the *restrictionist view*. The third part will concern the purpose of experimental philosophy. In the next chapter, I will defend experimental philosophy from its detractors by contending with some popular objections to it.

Some Characteristics of Experimental Philosophy

The aim of this section is to give a few characteristics of experimental philosophy. I will introduce experimental philosophy by discussing three things: (1) what an intuition is, (2) explore two projects within the experimental philosophy domain,[22] and (3) offer a few advantages of experimental work in philosophy. Together these components should provide an adequate introduction to experimental philosophy.

An introduction to experimental philosophy would not be complete if it did not include a brief discussion of its methodological predecessor. So, I want to discuss conceptual analysis, experimental philosophy's methodological predecessor, first.

Conceptual analysis is one of the methodological predecessors of experimental philosophy. Others are popular, such as reflective equilibrium, but

---

[22] Thomas Nadelhoffer and Eddy Nahmias (2007) have argued that experimental philosophy is composed of at least three projects. I agree with them, but I believe two projects are too similar to explore under different subheadings.

conceptual analysis is the hallmark method of philosophical analysis since the early twentieth century.

Conceptual analysis consists in providing a set of necessary and sufficient conditions for a concept's application. Since it is reasonable to say that it is the standard view, I will call it "standard philosophical practice."

Because experimental philosophy uses empirical data and *a priori* investigations seem resistent to the incorporation of empirical methods, it has been perceived as a threat to standard philosophical practice. Calling experimental philosophy a threat is to overstate its objective. Experimental philosophy is not meant to supplant standard philosophical practice, e.g., conceptual analysis or reflective equilibrium; it is meant to complement them. Thus, we might think of experimental philosophy as making up for some of conceptual analysis's deficiencies.

Proponents of "standard philosophical practice" do not necessarily use empirical data in their analyses. They call on intuitions produced in response to thought-experiments as evidence in favor of accepting or rejecting some philosophical claim. Alexander and Weinberg have a short explanation of standard philosophical practice:

> Going back arguably at least to Frege (and, in some sense all the way back to Socrates), it has been a standard practice in analytic philosophy to employ intuitions generated in response to thought-experiments as evidence in the evaluation of philosophical claims. A philosopher, wishing to either establish or prosecute some philosophical claim proposes a thought-

> experiment intended to generate an intuition relevant to evaluating the philosophical claim. According to standard philosophical practice, the generated intuition provides evidence for the acceptance or rejection of the philosophical claim: the philosophical claim is *prima facie* good to the extent that it accords with the generated intuition, *prima facie* bad to the extent that it fails to accord with the generated intuition. (Alexander and Weinberg 2006, 1)

This type of standard philosophical practice can be done "from the armchair" because practitioners believe their own intuitions about cases are typical ("ordinary"). Philosophers attribute these "typical" or "ordinary" intuitions to people.

Some philosophical practitioners do not gather empirical evidence, but they use it in their work. These philosophers do not necessarily ignore empirical data but they do not go about collecting it either. The work of Jerry Fodor is a good example of such a practitioner who uses but does not collect empirical data.

Jerry Fodor, in *The Language of Thought* (hereafter "LOT;" 1975), argues for an undiluted mentalist approach to psychological reality, a nativist thesis about mental content, and the defense of the view that the mind works along the lines of a computer program. Fodor did not propose LOT as a merely conceptual or speculative thesis. Fodor believes that LOT is compellingly entailed by the assumptions at work in *empirical* theories of certain central human cognitive abilities. Although Fodor did not employ empirical methods to gather evidence for his central claims, he – nevertheless – used empirical data to argue for LOT.

People like Fodor, and other prominent figures such as the early work of Stephen Stich or Fred Dretske, who use empirical data but do not collect it are not necessarily paradigmatic figures in standard philosophical practice because they are attentive to the empirical data. They use the empirical data to draw informed conclusions. Since they are attentive to the empirical data and attempt to formulate theories with an eye toward this data, Fodor and others like him are beyond the scope of what I have termed "standard philosophical practice."

The standard philosophical practitioner claims that "what we think…" is supposed to be indicative of what ordinary people believe is true. Jackson (1998) has defended this view. He writes

> I am sometimes asked… why… don't I advocate doing serious opinion polls on people's responses to various cases? […] often we know that our own case is typical and so can generalize from it to others" (Jackson 1998, 36f).

Not to corroborate evidence by empirical means is the paradigmatic component of what I call standard philosophical practice.

Some practitioners of standard philosophical practice use ordinary intuitions in their research, but the intuitions are generated from what *they think* ordinary people's intuitions are. I call on Frank Jackson's work again:

> How should we identify our ordinary conception? […] Intuitions about how various cases… are correctly

> described in terms of free action, determinism, and belief are precisely what reveal our ordinary conceptions of free action, determinism, and belief or… our folk theory. For what guides me in describing an action as free is revealed by my intuitions about whether various possible cases are or are not cases of free action. Thus my intuitions about possible cases reveal my theory of free action – they could hardly be supposed to reveal someone else's! Likewise, your intuitions reveal your theory. To the extent that our intuitions coincide, they reveal our shared theory. To the extent that our intuitions coincide with the folk, they reveal the folk theory. (Jackson 1998, 31f)

They attempt to formulate a folk theory of *x*, whatever kind of philosophical claim (or concept) *x* may be, based upon their view of folk intuitions.

Standard philosophical practice employs intuitions generated in response to thought experiments. The intuitions are evidence for the evaluation of philosophical claims. The researcher proposes a thought experiment meant to uncover an intuition relevant to the philosophical claim under consideration. They use intuitions garnered from the thought experiment to support or to criticize the philosophical claim under consideration. A philosophical claim is correct insofar as it agrees with the intuition and is incorrect insofar as it disagrees with the intuition.

The use of thought experiments to mine ordinary intuitions is common in the problems of personal identity, philosophy of mind, epistemology, and ethics. In ethics, for example, Judith Jarvis Thomson uses a hypothetical thought experiment to give substance to a general principle. She writes:

> But now let me ask you to imagine this. You wake up in the morning and find yourself back to back in bed with an unconscious violinist. A famous unconscious violinist. He has been found to have a fatal kidney ailment, and the Society of Music Lovers has canvassed all the available medical records and found that you alone have the right blood type to help. They have therefore kidnapped you, and last night the violinist's circulatory system was plugged into yours, so that your kidneys can be used to extract poisons from his blood as well as your own. The director of the hospital now tells you, 'Look, we're sorry the Society of Music Lovers did this to you – we would never have permitted it if we had known. But still, they did it, and the violinist now is plugged into you. To unplug you would be to kill him. But never mind, it's only for nine months. By then he will have recovered from his ailment, and can safely be unplugged from you.' Is it morally incumbent on you to accede to this situation? (Thomson 1971b, 123)

The example has characteristics and expectations that enable it to engage the readers and encourage them to actively participate in the issue and not just passively observe. Kathleen Wilkes is a critic of using thought experiments as a part of devising a solution to a deep philosophical problem. She has observed that thought experiments, like Thomson's, are

> forays of the imagination… for concluding that a philosophical thesis is plausible or implausible… [which] obey many of the constraints on experimentation. (Wilkes 1988, 2)

Standard philosophical practitioners suppose that one's own conclusions about the thought experiment apply to everyone who will read it.

Proponents of experimental philosophy believe that the standard philosophical practice assumes too much about ordinary intuitions. When a philosopher calls on "our" intuitions or says "we would all agree that…", people like Anthony Appiah protest rather forcefully: "What do you mean "we," Kemo Sabe?" (Appiah 2008, 80). The assumption that a philosopher is a reliable source of ordinary intuitions may be incorrect. A philosopher's own intuitions about shared intuitions with the folk maybe inaccurate and not the most reliable evidence.

Joshua Alexander and Jonathan Weinberg (2007) have addressed the concern over whose intuitions are most reliable. They have divided the landscape into three convenient spaces: *intuition solipsism*, *intuition elitism*, and *intuition populism* (Alexander and Weinberg 2007, pp. 2-5). Intuition solipsism, on one end of the spectrum, seems not to care about anyone's intuitions besides the philosopher's own. Intuition elitism seems to prefer a philosopher's own intuitions over any ordinary person's views. Intuition populism, on the other end of the spectrum, attempts to incorporate people's intuitions. Each of these positions deserves some consideration.

First, intuition solipsism is the view that a philosopher relies on her own intuitions when a philosopher depends on intuitions as evidence. Since philosophers often speak of intuitions in an impersonal voice, e.g., "we have the intuition that…" or "our intuitions are…", such appeals to intuition are beyond the scope of one's own intuitions. Intuition solipsism is too narrow to be philosophically profitable. When philosophers speak of *ordinary intuitions* or a *folk theory*, they want what ordinary people think is the case, not what they think is the case.

Second, according to intuition elitism, when a philosopher appeals to "our" intuitions, she takes her own intuitions to be representative of the intuitions of other professional philosophers. If a philosopher's own intuitions are representative of the intuitions of other professional philosophers, then that would limit our understanding to some technical sense of the concept under consideration. Epistemologists, for example, are interested in the concept of knowledge as it is ordinarily understood outside of strictly philosophical discourse. Intuition elitism constrains us to the intuitions of the few. A philosopher's intuition about some concept, such as "knowledge," is fascinating. But, if epistemologists seek "knowledge" generally and if there is a chance that a philosopher's own intuitions are inconsistent with an ordinary person's intuitions, then the intuitions of the few fail to capture intuitions about "knowledge" generally. I will suggest later that the incorporation of sophisticated forms of polling and statistical analyses will complement the elitist's methods.

Finally, intuition populism is the view where the philosophers own intuitions are representative of folk intuitions. The ability to have the same intuition as a representative sample of the folk about some philosophical claim may be hard to come by. We may err toward our own intuitions. So, some other means of discovering folk intuitions may be helpful if we seek ordinary intuitions about some philosophical claim. Why not go directly to the source of the intuitions? We may perform some experiments to discover what the folk intuitions are.

So, this discussion implies that experimental philosophy is a class of philosophical research methods used in order to examine intuitions by using systematic experimentation and statistically analysis on human subjects. Systematic experimentation and statistical analysis are associated with experimental psychology or cognitive science. What makes experimental philosophy 'experimental' is that proponents of it actually run studies to get the data they need about ordinary intuitions. What makes experimental philosophy 'philosophical' is that proponents of it discuss the implications the accumulated data have for deep philosophical problems. Experimental philosophers design surveys to test laypersons' intuitions about some philosophical claim.

The definition of experimental philosophy, however, seems to be insufficient because it does not tell us *who* the targets of the empirical research are, *what* ordinary intuitions are, *how* systematic experimentation leads to philosophically interesting conclusions, or *why* naturalistic and empirically informed philosophy is any different than experimental philosophy. I would now like to discuss *what* an intuition is.

What are intuitions?

Intuitions have been the subject of serious philosophical controversy (see e.g., Ramsey and DePaul 1998 or Pust 1999). So, nothing I say here may settle the debate over the nature of intuitions. My focus is much more narrowly conceived. I want to discuss what experimental philosophers have said about intuitions in their work and the problems that arise from them.

Let me begin by surveying what experimental philosophers have said an intuition is. Jonathan Weinberg has been most vocal in his characterizing intuitions, within the experimental philosophy literature. So, much of what I say here has to do with his work. Weinberg has given various explanations of intuitions. An intuition is:

1. Intellectual happenings in which it seems to us that something is the case without arising from our inferring it from any reasons that is so, or our sensorily perceiving that it is so, or our having a sense of remembering that it is so. (Weinberg 2007, p. 318)

2. An intellectual seeming of opaque origin. (Alexander and Weinberg 2006, p.1n1)

3. A spontaneous judgment about the epistemic properties of some specific case – a judgment for which the person making the judgment may be able to offer no plausible justification. (Weinberg, Nichols, and Stich 2001, p. 432)

Lets presume that these three stipulations are a good sample of what an intuition is. Some experimental philosophers may disagree with Weinberg's characterization of intuitions. But, for the moment, these are the ways we will characterize intuition until more experimental philosophers devise a well staked-out view of intuitions.

The views are fair accounts of what we mean by "intuition." But (1) suggests that intuitions arise *ex nihilo* because they seem to be completely unrelated to anything else in our cognitive repertoire. (2) suffers from problems of ambiguity because there is nothing that mediates between the view that something is the case and one's belief that that is so. Finally, (3) seemingly over-emphasizes the

spontaneity of an intuition. In what follows, I want to offer an account that captures the merits of (1), (2), and (3) while leaving some of their problems behind.

(1) suggests that intuitions arise out of nothing, *ex nihilo*. On this view, intuitions are completely unrelated to anything. Intuitions are not derived from reason(s), sensory perception, or memory. They *just are*. Perhaps this is true of intuitions. It is possible that people's intuitions arise because of some cultural or technological influence upon a person. Although Weinberg's notion of intuition does not rule out the influence upon people's intuitions, it does not forthrightly accommodate these influences either. People's intuitions may be derived from some overwhelming yet undisclosed or still unknown cultural influence. We may agree that intuitions are non-inferential, but they are not completely unrelated from all other influences, even if undetected. A view of intuitions should have to account for these influences.

Second, what an intellectual seeming is is ambiguous. Such ambiguity may result in some problems. Alexander and Weinberg's understanding of an intuition seems to follow Bealer's (1993, 1998) and Pust's (2000) definition of intuition. Bealer and Pust believe that intuitions are "intellectual seemings." Designating intuitions as "intellectual seemings" is ambiguous. Intuitions could be *perceptual* intellectual seemings. Perceptual intellectual seemings could be types of sensory experiential states. The problem with this view of intuitions is that no sensory experience mediates between the fact that 7 and 5 yield 12 and my belief that this is so. Nothing like sensory experience plays that role, either. When someone thinks

about the equation, they do not find any distinct intellectual seeming. The person just finds his belief that 7 and 5 yield 12. Thus, what I will call the Alexander/Weinberg view of intuition should be revised in order to avoid the ambiguity that arises from perceptual intellectual seemings.

Finally, intuitions do not have to be had spontaneously, as (3) suggests. Weinberg, Nichols, and Stich's view of intuition seem to suggest that intuitions are time-sensitive. For them, it appears *as if* only spontaneous judgments are intuitions. The problem with this view is that intuitions do not necessarily depend on the amount of time that passes between the time one is presented with a scenario and the time at which one has formulated a response to that scenario. An important point Weinberg, Nichols, and Stich's view of intuition raises is that intuitions are non-inferential. Intuitions do not develop out of other beliefs. But non-inferential beliefs need not be spontaneous; in fact, some anecdotal evidence suggests that respondents may not have intuitions immediately after reading the vignettes we philosophers distribute to subjects.[23]

All of this is a bit of conceptual analysis about the nature of intuition. And I have tried to point out a few problems about some of the common views about the nature of intuitions in experimental philosophy. Given that these views of the nature of intuition seem problematic, I want to offer an account of intuition that seems to capture what experimental these two attempts have sought.

---

[23] After a recent experiment, a subject sent a brief email to me apologizing for not answering the question following the vignette. He cited "not having any initial reaction" to the thought experiment for his .

The account of intuition roughly follows Sosa (2006, 2007a, 2007b). On this view, to intuit that *p* is to be consciously attracted to assenting to *p*. Intuitions must satisfy at least four conditions: (a) we must understand *p* well enough, (b) *p* is modally strong or self-presenting, (c) our attraction to judge that *p* does not derive from any of the usual sources of evidence: introspection, perception, memory, testimony, or inference, and (d) our attraction to judging that *p* is virtuously based (that is, an ability to discriminate truths from falsehoods).

Sosa's view of intuitions seems to capture all the benefits of (1), (2), and (3) while leaving their deficiencies behind. First, unlike Alexander and Weinberg, it does not equate intuitions with "intellectual seemings." On Sosa's view, one is "consciously attracted to assenting" to some proposition. So an intuition is more like a feeling toward *p* than having or possessing intellectual concept *p*. Second, it seems to endorse the view that intuitions are conscious states. Although intuitions are a kind of conscious state, they are not the types of conscious states associated with sense perception. Third, it permits that intuitions are fallible.[24] Fourth, it explains how intuitions can have probative force. Intuitions aim at truth, though they sometimes miss their target. If we assume – along with Sosa – that intuitions are virtuously based in reliable cognitive abilities, they still can be said to "enjoy *prima facie* epistemic justification" (Sosa 2006, p. XXX). Finally, the account is externalist

---

[24] Two points need to be discussed. First, Weatherson (2003) has an excellent discussion about what the fallibility of intuitions should teach us about their role in philosophical theory. Second, none of the other views would deny the fallibility of intuitions, so this point is not necessarily a criticism of these other views but raising an important point about intuitions.

because a person need not have any beliefs, justified or not, about one's intuitions in order for them to have probative force. It is enough on this account that they are virtuously grounded.

Types of Experimental Philosophy

Experimental philosophy investigates ordinary intuitions. We have explored intuitions. I now want to turn to two different types of experimental philosophy. These two types of experimental philosophy use intuitions differently. One abandons them altogether and the other explores them empirically to devise more grounded philosophical theories. The two types of experimental philosophy I will discuss here include, what Alexander and Weinberg (2007) have called, the proper foundational view and the restrictionist view.

The proper foundational view uses the results of experimental work to provide a proper evidentiary foundation for certain philosophical claims and projects. The standard philosophical practice incorporates an appeal to intuitions as evidence for or against some philosophical claim. When we assert something about the general population's intuitions, it is a claim about the distribution of intuitions. Claims about the distribution of intuitions are empirical claims. Empirical claims about the distribution of intuitions are testable predictions about how people will respond when presented with a thought experiment. So, we should be concerned with conducting experimental research to determine what are the intuitions held by philosophers and non-philosophers alike.

If we seek the folk account of some philosophical claim and philosophical intuitions need to be representative of ordinary intuitions, then the proper foundational view seems to be a more powerful way of going about doing philosophical research. These empirical results can deliver the intuitions that may serve as evidence for or against philosophical claims. Therefore, the proper foundational view supplements standard philosophical practice by providing a proper evidentiary foundation for certain philosophical claims and projects.

The second position is the restrictionist view. On the restrictionist view, the results of experimental studies should figure in the radical restriction of the employment of intuitions as evidence. Weinberg, Nichols, and Stich (2001) have shown that socio-economic status plays a "major" role in producing subjects' epistemic intuitions. The role is substantial enough for the three to conclude somewhat hesistantly:

> the data we've reported look to be yet another serious embarrassment for the advocates of IDR [intuition-driven romanticism]. As in the case of cultural difference, they must either argue that these intuitive differences, when plugged into an IDR black box, would not lead to different normative conclusions, or they must bite the bullet and argue that diverging normative claims are genuinely normative, and thus that the sorts of doxastic states that ought to be pursued by relatively rich and well-educated people are significantly different from the sorts of doxastic states that poor and less well educated folks should seek… we certainly don't envy the predicament of the IDR advocate who has to opt for one or the other. (Weinberg, Nichols, and Stich 2001, 447f)

For restrictionists, the experimental evidence seems to point to the unsuitability of intuitions to serve as evidence at all. Their arguments are against not only the standard philosophical practice but also the proper foundational view. Experimental evidence that has shown cultural differences or socio-economic diversity seem to question the correctness of utilizing intuitions as evidence.

The restrictionist view is too conservative. People have been raised in different cultures, and they are bound to have different intuitions about philosophical thought experiments. In spite of the diversity of people's intuitions at the micro-level, some people's intuitions will overlap with other people's intuitions at the macro-level. Some overlap in intuitions warrants the use of people's intuitions in partial support of some theoretical view. So, the suggestion that we eliminate the use of intuitions altogether might be too demanding because of our limited perspective on people's intuitions.

The proper foundational view and the restrictionist view represent two types of experimental philosophy. They do not exhaust the different approaches to experimental philosophy, but a description of these two projects suffices for the purposes of this project.

Some Advantages of Experimental Philosophy

Applying social scientific methods to philosophical problems is advantageous. First, introducing innovative methods into philosophy may result in greater creative

activity among its practitioners. Philosophers have generally accepted conceptual analysis as *the* method with which to do philosophy. Frank Jackson writes, "[Conceptual analysis] is what philosophers have traditionally spent a good deal of time doing" (Jackson 1998, vii). For too long, it has seemed that the commonly held belief is that if conceptual analysis is the widely accepted method to use, we should go on using it.

Experimental philosophy has stirred up some controversy. In fact, a cottage industry of articles has been generated by the metaphilosophical question whether experimental philosophy should be an accepted methodology of philosophy (e.g., Kauppinen 2007, Nadelhoffer and Nahmias 2007, Knobe 2007). The introduction of a methodological discussion is advantageous for philosophy generally. The discussion is reminiscent of what Thomas Kuhn called a *paradigm shift* (Kuhn 1961).

Kuhn said that scientific progress occurs in violent revolutions where one conceptual worldview is replaced by another. Such progress is good for science, even if the shift is a violent one. Kuhn's discussion of the incommensurability of competing paradigms reminds us philosophers the violent conceptual shift that seems to be underway in the discipline. He wrote:

> Practicing in different worlds… two groups of scientists see different things when they look from the same point in the same direction. Again, that is not to say that they can see anything they please. Both are looking at the world, and what they look at has not changed. But in some areas they see different things, and they see them in different relations one to the other. That is why a

> law that cannot even be demonstrated to one group of scientists may occasionally seem intuitively obvious to another. Equally, it is why, before they can hope to communicate fully, one group or the other must experience the conversion that we have been calling a paradigm shift. Just because it is a transition between incommensurables, the transition between competing paradigms cannot be made a step at a time, forced by logic and neutral experience. Like the gestalt switch, it must occur all at once (though not necessarily in an instant) or not at all. (Kuhn 1961, 150)

Experimental philosophy has generated a methodological paradigm shift in philosophy. Its employment of social scientific methods has shifted focus from purely *a priori* methods to *empirical* or *a posteriori* methods.[25]

A second advantage of experimental philosophy is the challenge it poses for the methodological assumptions of conceptual analysis. Manuel Vargas writes, "Experimental work raises important questions about the methodological assumptions that go undiscussed in a good deal of the philosophical literature" (Vargas 2006, ??). An important feature of the philosophical endeavor is the metaphilosophical question whether conceptual analysis or reflective equilibrium, for example, is adequate. If a method is inadequate, then either a new method should be devised or the old one

---

[25] Critics of Kuhn, such as Imre Lakatos, have berated his notion of theory choice based on the *paradigm shift* as a "matter of mob psychology" (Lakatos 1970, 178). Moreover, Dudley Shapere has written that Kuhn believes "the decision of a scientific group to adopt a new paradigm cannot be based on good reasons of any kind, factual or otherwise" (Shapere 1966, 67). Interestingly enough, staunch supporters of the traditional *a priori* methods of philosophy have raised similar criticisms against experimental philosophy. Critics have found experimental philosophy unsettling because it derives philosophical theories directly from the folks' intuitions.

revised. Regardless of the outcome, experimental philosophy has forced us to revisit issues about the adequacy of the methods we employ.

Finally, experimental work helps us become better acquainted with people. Philosophy arose out of simple though vitally important questions. Thales – the so-called first western philosopher – asked, "what is the basic stuff of the universe?" People – untrained and unspecialized people – have wondered about these questions too. Similarly, the foundation of the Platonic dialectic relies upon engaging in a question-and-answer session with ordinary people on the street. The historical foundations of philosophical inquiry are a testament to the importance of ordinary people in the philosophical dialogue. As long as we have questions about the universe that affect people and as long as people show a passing interest in these concepts, we ought to listen to what they have to say.

Sometimes we have to set aside our years of training and recognize the importance of people's intuitions about deep and interesting philosophical topics.

Given that some philosophers believe their own intuitions represent people's intuitions fairly accurately (Jackson 1998, Goldman 2001), the empirical data show that the philosophers' intuitions and people's intuitions do not always coincide (Weinberg, Nichols and Stich 2001). Thus, philosophers have sometimes mistakenly attributed their own intuitions to ordinary people. The purpose of experimental philosophy is to show where these philosophers have gone wrong (and – perhaps in some cases – show them how they can improve their philosophical claims by acquainting themselves with the folk).

CHAPTER 3

DEFENDING EXPERIMENTAL PHILOSOPHY

Introduction

In the previous chapter, I discussed experimental philosophy. Experimental philosophy, broadly speaking, is the systematic investigation of people's intuitions. Philosophers have not attended to people's *actual* intuitions because they have not believed it is important to discover what they are (Jackson 1998, 36f). Frank Jackson writes, "we know that our own [intuitions] are typical and so can generalize from it to others." So, the problem is whether the sort of conceptual analysis that Jackson (1998) defends may argue from his own intuitions to people's *actual* intuitions.

The aim of this chapter is to contend with objections to experimental philosophy. I will address one of experimental philosophy's most vocal critics: Antti Kauppinen. Once I have addressed his objections, I will show how experimental philosophy applies to action theory. Experimental philosophy is a valuable contribution to philosophy, so we cannot discount the methods experimental philosophy employs if our goal is to elucidate a folk conception of action individuation.

Some Challenges to Commonly Employed Philosophical Methods

The previous section defined experimental philosophy – including a discussion of some of its divisions – and explained some of its advantages. We still have to provide some reasonably persuasive arguments against the traditional methods philosophy employs: *conceptual analysis* and *reflective equilibrium*. This section will provide some *pro tanto* reasons for rejecting these methods.

The following analysis does not presume to reject conceptual analysis *simpliciter*. Such an endeavor is beyond the scope of this project. I want to challenge those projects that employ folk intuitions without checking if people actually have these intuitions. Since both conceptual analysis and reflective equilibrium assume to have access to ordinary intuitions and there is reason to be skeptical over whether they do get at ordinary intuitions, I will argue that we reject these methodologies unless they adopt experimental methods to confirm (or disconfirm as the case may be) philosophical claims. Only after we have done some empirical leg work will the claims about the folk gain the proper foundation.

Many philosophical problems or projects find their source in intuitions.[26] For example, "what is free will?", "what is knowledge?", or "what is it for an action to be intentional?" are questions of deep philosophical significance that find their origin in intuitions. We consult our intuitions about various cases to develop an account of

[26] When I say "many," I do not mean all philosophical problems or projects find their source in intuitions. There may be methods philosophers use which do not depend at all on intuitions. These methods are outside the scope of my dissertation, so I will not address them here.

these philosophically important notions. Therefore, some philosophical accounts depend on intuitions.

Two dominant methodological traditions, conceptual analysis and reflective equilibrium, sometimes depend on intuitions. Both show that philosophical inquiry is driven by intuitive judgments, by what "we would say" where "we" stands in for not just the specialist but people everywhere. I will now outline each of the methodologies.

Conceptual Analysis

Conceptual analysis refers to an activity, and it refers to a movement. The movement of conceptual analysis had its heyday in Britain in the 1950s and is now more or less defunct (cf. Hanna 2000). Since this is not an argument about an historical time-period, I will have fewer things to say about conceptual analysis as a movement. Conceptual analysis as an activity, which of course goes back much further than mid-twentieth-century Britain, is still frequently practiced and paradigmatic of philosophical endeavors.

We can see the origin of conceptual analysis in the works of ancient Greek philosophy and the early modern period. If we read a text such as Plato's *Republic*, we will find the characters, particularly Socrates, using an ancient form of conceptual analysis. For instance, Socrates encounters someone who claims to have figured out the true essence of some abstract notion. The person posits a definition or analysis of the notion in the form of necessary and sufficient conditions that are thought to

capture all and only instances of the concept in question. Socrates then refutes his interlocutor's definition of the concept by pointing out various counterexamples. The dialogue eventually yields a definition of the concept that may or may not conflict with our own intuitive judgments. So, the proposed analysis is rejected because it fails to capture our intuitive judgments about the concept in question.

Similarly, we see an origin of conceptual analysis in modern philosophy too. The early modern philosophers, particularly Descartes, used procedural rules and tests in their method. For example, Descartes's method of hyperbolic doubt was an "excessive desire to distinguish between the true and the false," in which he sought "the true Method of arriving at all things of which [his] mind was capable." (Descartes 1641/1984, 15f) The method of hyperbolic doubt was an early modern precursor of conceptual analysis seeking absolute certainty.

The Socratic Method and Descartes's method of hyperbolic doubt are complementary methods and represent the origins of conceptual analysis. There are updated forms of the analytic methods still in use today.

Conceptual analysis can have different purposes. First, it could be merely descriptive. One can use conceptual analysis in psychology, sociology, and anthropology. Philosophy is a normative enterprise. More often than not, philosophers are interested in the *explicatory* purpose of conceptual analysis. Carnap used the term "explication" to stand for "the transformation of an inexact, prescientific concept, the *explicandum,* into a new exact concept, the *explicatum*" (Carnap 1950, 3). The basic idea is that the explicatum provides the necessary and

sufficient conditions that the explicandum does not possess. The explicatum is an improvement in comparison to the explicandum, even though the explicatum uses the explicandum to come up with the necessary and sufficient conditions for the concept.

Frank Jackson (1998) has defended conceptual analysis. In his work, he has defined conceptual analysis more explicitly than Carnap. Jackson describes conceptual analysis as "the very business of addressing when and whether a story told in one vocabulary is made true by one told in some allegedly more fundamental vocabulary" (Jackson 1998, 28).[27] Conceptual analysis attempts to clarify our pretheoretic ideas through the careful analysis of a concept. My concern focuses on philosophical analyses that use intuitions the philosopher ascribes to ordinary people.

The properties of the term that interest us are those that are understood by ordinary people. C.D. Broad has said:

> the most fundamental task of Philosophy is to take the concepts that we daily use in common life and science, to analyse them, and thus to determine their precise meanings and their mutual relations. (Broad 1927, 16)

---

[27] Audi has presented a competing definition: "Let us simply construe it as an attempt to provide an illuminating set of necessary and sufficient conditions for the (correct) application of a concept" (1983, 90).

The analysis of concepts or the analysis of the meanings of expressions is a technique practiced nowadays. What we seek, according to Broad and others, is a folk theory – a theory that ordinary people would adopt (or already have adopted).

In contemporary analytic philosophy, it is quite explicit that the goal is to give an analysis of *folk* concepts (Gibbard 1990, Jackson 1998, Lewis 1972). For example, Jackson writes:

> Thus my intuitions about possible cases reveal my theory… they could hardly be supposed to reveal someone else's! Likewise, your intuitions reveal your theory. To the extent that our intuitions coincide, they reveal our shared theory. To the extent that our intuitions coincide with those of the folk, they reveal the folk theory (1998, 32).

By folk theory, Jackson seems to mean a set or collection of generally accepted platitudes. According to Jackson, the agreement of my intuitions with ordinary intuitions leads to a folk theory of some philosophically important topic (Jackson 1998, 30f).

If the goal is to analyze folk concepts, one might expect philosophers to ask the folk their opinions. But Jackson writes:

> I am sometimes asked... why, if conceptual analysis is concerned to elucidate what governs our classificatory practice, don't I advocate doing serious opinion polls on people's responses to various cases? My answer is that I do - when it is necessary (Jackson 1998, 36f).

Jackson, a staunch defender of conceptual analysis, claims that such "polls" are appropriate. But, on his view, such polls are unnecessary because "often we know that our own case is typical and so can generalize from it to others" (Jackson 1998, 37). I imagine that he means: if our own case is typical, then we need not query ordinary people about their own intuitions. People's intuitions are our own intuitions. Therefore, we can generalize from our own intuitions to people's intuitions.

A philosopher moves from his own intuitions to people's intuitions using nothing but an armchair. Here is how the story typically go: Richard sets out to analyze the concept of *knowledge*. Richard sits back in his armchair and thinks hard about knowledge. In particular, his interest is in the "what we would say when" of term knowledge. Richard tentatively proposes a set of necessary and sufficient conditions for his *correct* use of the term knowledge, which he either accepts or rejects. After some deliberation, he feels he has succeeded, writes the paper, and publishes it. Rachel reads Richard's paper about knowledge. She sits back in her armchair and tries to devise a counterexample to Richard's proposed analysis. She feels she has succeeded, writes the response paper, and publishes it. This is a paradigmatic case of conceptual analysis, or "counterexample philosophy" as Michael Bishop (1992) has named it.

To my mind, Jackson's belief that we may generalize from our own intuitions to those of the folk suffers from at least two problems. First, it seems that practitioners of conceptual analysis have believed that we can move from our own intuitions to ordinary intuitions very easily. Philosophers, so to speak, have a special

competency in accessing folk intuitions. Experimental studies, however, seem to undermine the special competency claim. The empirical data have shown philosophers' own intuitions do not track ordinary intuitions very well. If philosophers' own intuitions do not track ordinary intuitions consistently very well, then their intuitions have very little evidential value. One should not utilize a method without evidential value.

Next, if Jackson's claim may be construed as an identity statement, then saying that we can generalize about ordinary intuitions is controversial. When he claims that "our own case is typical" and we can "generalize from it to others," one may interpret that as: our own intuitions *just are* ordinary intuitions. Two facts play against the intuitions being identical. On the one hand, suppose that *x* is Frank's intuition and he believes we can generalize it to a group of people, call the intuition of the group: *y*. The two intuitions are identical only if all the properties of *x* and *y* are the same. But Frank's intuition and the group's intuition may differ. In fact, all members of the group may have generally the same type of intuition without also having the same particular intuition. For example, I may believe generally that "the sky is blue," and everyone else could have the same general intuition that "the sky is blue." When we examine our individual intuitions, we find that my intuition that the sky is aquamarine fails to jibe well with another's intuition that the sky has more periwinkle in it than I intuit. So, we cannot say that our own intuitions just are folk intuitions.

On the other hand, we might mean that since we are one member of the group, we share the same intuition as the members in the group. The problem with this view is that our disposition may prevent us from sharing intuitions with other people. We may not share our intuitions with others because our specialized training may alert us to subtleties of language that a non-specialist would not consider very important. Our specialization seems to forbid us from saying that we share others' intuitions. Thus, we cannot conclude that we share anything – including intuitions – with other members of the group.

Reflective Equilibrium

A large group of philosophers hold a more restricted view of the reliability of intuitions as a source of philosophical truth. The more qualified view is reflective equilibrium.[28]

Reflective equilibrium originated in the work of Nelson Goodman; John Rawls capitalized on Goodman's innovation. Goodman writes, "*a rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend*" (Goodman 1983, 64; emphasis is Goodman's). According to Goodman, the method of reflective equilibrium is a matter of bringing into accord judgments about particular inferences and about general principles of general inference. Among these judgments Goodman includes the rejection of the

---

[28] John Rawls, in *A Theory of Justice* (1999, 20ff), dubbed the conservative method *reflective equilibrium*.

premise of an inference. The accord gained is the only justification possible for the inferential principles that emerged.

John Rawls's reflective equilibrium is a modified version of Goodman's procedure. Rawls's version is a process of justifying moral principles and moral judgments. Rawls writes:

> In searching for the most favored description of the situation we work from both ends. We begin by describing it so that it represents generally shared and preferably weak conditions. We then see if these conditions are strong enough to yield a significant set of principles. If not, we look for further premises equally reasonable. But if so, and these principles match our considered convictions of justice, then so far well and good. But presumably there will be discrepancies. In this case we have a choice. We can either modify judgments, for even the judgments we take provisionally as fixed points are liable to revision. By going back and forth, sometimes altering the conditions of the contractual circumstances, at others withdrawing our judgments and conforming them to principle, I assume that eventually we shall find a description of the initial situation that both expresses reasonable conditions and yields principles which match our considered judgments duly pruned and adjusted (Rawls 1999, 20).

For Rawls, moral principles and moral judgments are justified on the grounds that they accord well with our intuitive judgments concerning particular cases. Intuitions are defeasible. Even if our intuitive judgments are confused or inconsistent, the process of reflective equilibrium advances principles that seem to accord well with most of our intuitions.

Reflective equilibrium tests theories against the "data" of what ordinary people say about real and hypothetical examples. For instance, if we observe the native speakers of a language and infer from their use of language what counts as grammatical or ungrammatical, then we can judge whether other language user's assertions are grammatical. Similarly, in practical reasoning, reflective equilibrium is a method that has been used to exclude a handful of exceptions to means-end inference (Fehige 2001) and to satisficing (Slote 1989).[29] Reflective equilibrium tests intuitive judgments that a theory proposes. Intuitive judgments are revisable in the face of contradictory evidence and in the face of well-established theories.

Philosophers must seek to create a philosophical analysis, theory, or account of the target subject matter together with the totality of their relevant intuitive judgments. The analysis must bring together background philosophical, scientific, political, and theological beliefs into a coherent whole. To do this, in some cases, philosophical definitions or theories will have to be revised when they conflict with intuitive judgments. In other cases, it will be the intuitive judgments that will have to be revised when they conflict not only with a developing philosophical view, but with various stable background beliefs as well.

---

[29] Both essays by Fehige and by Slote are reprinted in Elijah Millgram's anthology, *Varieties of Practical Reasoning* (2001, 49-76 and 221-235 respectively). In his Introduction, Millgram includes a discussion of reflective equilibrium in Fehige and Slote. Millgram takes a stand against the use of reflective equilibrium in practical reasoning by offering his own method in his latest book *Ethics Done Right: Practical Reasoning as a Foundation for Moral Theory* (2005). He ingeniously calls it the "Method of Practical Reasoning" (Millgram 2005, 1-32).

There are fewer criticisms of reflective equilibrium than conceptual analysis primarily because it recognizes the defeasibility of intuitions. So, what I will do here is argue that adopting the methods experimental philosophy utilizes may assist reflective equilibrium theorists. After all, the philosopher who uses reflective equilibrium believes that it advances principles that seem to accord well with most of our intuitions, even if our intuitive judgments turn out to be confused or inconsistent.

Reflective equilibrium acknowledges that philosophical theories sometimes have to be revised because they conflict with intuitive judgments *and* sometimes intuitive judgments will have to be revised when they conflict with a new philosophical view. Berys Gaut's (2002) view of moral pluralism is a good example of how to use reflective equilibrium.

Gaut has outlined a sketch of how the use of reflective equilibrium as a justificatory method will favor moral pluralism (Gaut 2002, 146ff). According to Gaut, common-sense morality has tolerated a plurality of principles, giving reasons not to lie, not to steal, or not to kill. Reflective equilibrium has sought to clarify these moral commitments. The clarification of moral commitments has accomplished two goals: (1) it makes moral principles and judgments about particular cases consistent and (2) it renders moral commitments in general consistent with non-moral beliefs, i.e., the nature of persons and of society. Moreover, according to Gaut, a reason the process of reasoned improvement would not result in one fundamental moral principle from which all our judgments could be derived is that agent-relative reasons, e.g., I have a reason to help *my* children or friends because of the relationship

I have with these people, account for our moral commitments. Agent-neutral reasons cannot account for moral commitments and lead to moral dilemmas. As Gaut points out, one could always ignore or alter one's considered moral judgments when they clash with one's favored moral system. For these reasons, the pluralist seems to have good reason to be hopeful that he should win the battle against monistic views.

Gaut's use of reflective equilibrium informs the debate on moral pluralism. His analysis has not ruled out intuitions driving theory or theory driving intuitions. Perhaps what would be beneficial for Gaut position's is to present empirical evidence supporting the idea that we question the applicability of one of the options. If we can raise skeptical worries about the accuracy of ordinary intuitions, then we ought to seek alternative means of improving Gaut's position. The empirical data could cause us to revise or to endorse Gaut's moral pluralism if its account of ordinary intuitions is true.

Recent experimental work has challenged traditional methods, but I do not want to challenge reflective equilibrium here. The process of reflective equilibrium recognizes the importance and flexibility of ordinary intuitions. The experimental work has asked the folk for their opinion. By asking the folk for their intuitions, we would improve the accuracy of Gaut's position. Some recent evidence has shown that there are individual differences about intentional action (Nichols and Ulatowski 2007), that is, people interpret 'intentional' differently. These results have been employed to pose a problem for monistic views of morality. If we could devise a theory in which it is shown that people's intuitions are flexible and if that experiment

could test whether one's considered moral judgments are altered or ignored when they clash with a favored moral judgment, then we may have further reason to endorse Gaut's pluralism.

The overall goal of this section was to challenge the commonly held belief that philosophy from the armchair captures ordinary intuitions. But, in some cases, e.g., reflective equilibrium, the use of experimental methods may assist traditional philosophical analysis in devising the *best* theory or account. We have little reason to think that a philosopher's intuitions carry any normative weight (Weinberg, Nichols, and Stich 2004). A philosopher's intuitions about folk intuitions *are not necessarily* an accurate portrayal of folk intuitions. If philosophers who use the traditional methods of philosophy want to call on ordinary intuitions, then gathering empirical data to support what they say about the folk would lend some credence to their claims. We could trust what the philosopher says about ordinary intuitions and support what they want to say about some concept or support their folk theory of *x*. The way to ordinary intuitions is through experimental philosophy.

Kauppinen's Criticisms

In this section, I will explore one argument against experimental philosophy that shows robust intuitions are out of reach when we use the experimental method. In the next section, I will investigate another argument against experimental work in

philosophy that claims the philosophers own intuitions should still play an indispensable role in philosophical methodology.

Antti Kauppinen (2007) has argued that the subjects' intuitions studied by experimental philosophy may not reflect the relevant concepts. The variation of intuitions between subjects and philosophers or between different groups of subjects, e.g., cultural diversity, can be dismissed because the subjects are employing different concepts. Since the variation of intuitions is the result of multiple concepts in play, the variation does not challenge standard philosophical practice.

Kauppinen tells us that philosophers make three characteristic assumptions about the responses that count as revealing people's concepts. These assumptions include *the competence of the speaker*, *sufficiently ideal conditions*, and *basis in semantic rather than pragmatic considerations*.

First, the speaker's competence matters because an incompetent speaker would tell us nothing about the relevant concept. Second, Kauppinen believes certain *ideal* conditions need to be met for judgments to avoid conceptual mistakes. Ideal conditions are those in which there are no perturbing, warping, or distorting factors or limits of information, access or ability (Pettit 1999, 32). Finally, Kauppinen says:

> even if we limited ourselves to responses by competent speakers in ideal conditions, what they would say about particular cases would not necessarily reveal to us what we are interested in, namely the *semantic* contours of the concept at hand or the contribution it makes to the truth conditions of sentences in which it is used. (Kauppinen 2007, CITE)

Because of these assumptions, the intuitions gleaned by experimental philosophy may not reveal the relevant concepts. The multiple variations in people's intuitions fails to challenge standard philosophical practice.

First, Kauppinen's worry that variation of intuitions among ordinary people reveals that they are not using concepts univocally should carry over to philosophers too. Any substantive variation of intuitions between philosophers can be dismissed as resulting from the fact that different philosophers are utilizing different concepts. Failing to use concepts univocally shows that the relevant concepts have not been targeted. So, the intuitions of philosophers may not reflect the relevant concepts.

Second, Kauppinen seems to privilege philosophers' intuitions, and the argument makes sense only if we do privilege philosophers' intuitions. The argument begs the question against the experimental philosopher by flat-out rejecting the relevance of folk intuitions to standard philosophical practice. Kauppinen owes us a non question-begging account for why philosophers' intuitions are more reliable than the folk's intuitions. Moreover, not only is this objection mistaken about privileging philosophers' intuitions but also it is mistaken because philosophers assume they have *special privileged* access to folk intuitions. Philosophers rarely have the same intuitions among themselves. Why should we think that they get the folks' intuitions right? So, there is no reason for us to think that philosophers *know* what the folk's intuitions are.

Third, philosophers seek a folk account. If we did not seek a folk account (or one closely approximating what ordinary people think), then a philosophical matter would be unworthy of investigation. For example, Jackson writes of a folk theory of rightness:

> If we wish to address the concerns of our fellows when we discuss [a philosophical] matter – and if we don't, we will not have much of an audience – we had better mean what they mean. We had better, that is, identify our subject via the *folk* theory of rightness, wrongness, goodness, badness, and so on. We need to identify rightness as the property that satisfies, or near enough satisfies, the folk theory of rightness – and likewise for the other moral properties. It is, thus, folk theory that will be our guide in identifying rightness, goodness, and so on. (Jackson 1998, 118)

Even Kauppinen agrees with Jackson's sentiment when he writes, "why should anybody care about what philosophers do if they just argued about their own inventions?" (Kauppinen 2007, CITE) A folk account seeks more than just an expert's view. Thus, there is no reason to exclude the folk from empirical investigation if we take philosophy to be searching for how the folk use a concept.

Fourth, a somewhat complicit way of responding to Kauppinen's criticism is to concede that some folk intuitions are not philosophically interesting. Some philosophical topics will not be as accommodating as other topics are. For example, some areas of philosophy, such as the philosophy of mathematics or logic, are so obtuse that we may not care what people's intuitions about them are. That some

intuitions may not be relevant to philosophy does not entail (or even imply) that we rid philosophy of its use of folk intuitions. Folk intuitions ought to at least serve as constraints on philosophical investigations.

Finally, the appropriateness of what we say depends on pragmatic factors that are not a part of the meaning or semantic content of expressions. In fact, there are certain ordinary contexts in which it would be inappropriate for speakers to use the semantically appropriate expression. Kauppinen asserts the surveys of lay persons makes it practically impossible to "separate the contribution of semantic and pragmatic considerations to what people say (and what it is *proper* to say)" (Kauppinen, CITE). But pragmatics is important for experimental philosophy. Pragmatics is the study of language which focuses attention on the users and the context of language use rather than reference, truth, or meaning. In the Knobe effect literature, for instance, Fred Adams and Annie Steadman (2003a; 2003b) have staked out a position suggesting that the subjects' responses were being driven by pragmatic considerations having to do with the relationship between intentional action and moral responsibility. If they are correct, these pragmatic considerations outweighed the semantic considerations having to do with the folk concept of intentional action. So, it seems plausible that philosophers are interested in both semantics and pragmatics. I find it hard to believe that we philosophers drop pragmatics altogether.

Each of the characteristic assumptions philosophers make about the responses that count as revealing people's concepts have been overturned. What is left for Kauppinen is a framework argument that has its weaknesses too.

Kauppinen argues that when a philosopher claims that according to an intuition, e.g., Gettier cases are not knowledge, they are narrowly and typically making a claim of how competent users of the concept of knowledge would pre-theoretically classify the case in suitable ideal conditions (and without being influenced by irrelevant factors) (Kauppinen, CITE). Call these narrow intuitions *robust intuitions*. He doubts whether experimental studies are capable of shedding any light on robust folk intuitions. Philosopher's claims, therefore, are out of reach for those who use the survey model, such as experimental philosophers. According to Kauppinen, philosophically interesting folk intuitions can only be uncovered by a process of *dialogue and reflection* (CITE).

Two issues deserve consideration: Kauppinen's defense of dialogue and reflection model and his related claim that philosophers are only interested in *robust intuitions*. According to Kauppinen's dialogue and reflection model, philosophers ought to be concerned only with the intuitions competent language users would express under ideal conditions. Nadelhoffer and Nahmias (2007) have shown why Kauppinen's dialogue and reflection model is flawed. On one interpretation they call the *informal dialogue and reflection model*, "it suffices for the purposes of doing conceptual analysis that we engage in discussion with our colleagues, students, friends, and family" (Nadelhoffer and Nahmias 2007, CITE). But the informal model suffers from several problems. First, since a philosopher's intuition will be shaped by her social environment, her educational background, and her cultural heritage, "two different philosophers may understandably end up with two very different ideas with

respect to what most people actually think about a particular topic" (Nadelhoffer and Nahmias 2007, CITE). Second, a philosopher's intuition about ordinary folk intuitions and concepts have been influenced by the theories they have adopted during their extensive training. Consequently, philosophers may favor the ordinary intuitions that support their views.

On a second *experimental* interpretation of Kauppinen's dialogue and reflection model, a controlled and systematic experiment ought to be devised to test what people's reflective intuitions, judgments, and beliefs about a given philosophical claim are. Implementing an experimental version of Kauppinen's model has three important lessons: (1) "experimental philosophers have resources to get at what Kauppinen takes to be the philosophically interesting intuitions"; (2) the experimental model makes clear that any information a philosopher can access using informal conversation can be accessed using controlled experiments; and, (3) the experimental model shows that experimental philosophers are interested in more than just surface intuitions (Nadelhoffer and Nahmias 2007, CITE). The employment of systematic and controlled experimental studies, therefore, seem to help Kauppinen's special brand of conceptual analysis, the dialogue and reflection model.

Kauppinen's second problem stems from his concern that experimental philosophy does not get at robust intuitions. His argument that philosophers should not be interested in any old folk intuitions rests on the mistaken belief that robust intuitions cannot be examined using experimental methods. The experimental methods used are statistical methods. Statistical methods allow the researcher to

examine correlations between manipulated factors, even in light of expected effects of the "noise" of other factors. When the experimenter uses a sufficiently large sample size, he can show that the probability is extremely low that the relevant results were obtained because of the irrelevant factors. Experimental studies do not rule out a subject's response due to inattention, confusion, or a desire to mess up the experiment, but they do indicate that it is highly unlikely that most subjects were inattentive, confused, or mischievous.

Experimental philosophers make an effort to assure subjects avoid irrelevant factors. For instance, they instruct subjects to read the scenarios and questions carefully before they respond. Moreover, experimental philosophers give subjects ample time to complete the surveys. Since the experimental method is statistical analysis, experimental philosophers *control* for whether participants are following the instructions by using tests that check whether subjects have understood the scenarios they have read. If subjects miss these questions, then they are excluded from the analysis.

Experimental philosophers are very aware of the methodological difficulties the survey method may present. Nevertheless, statistical analysis enables experimental philosophers to avoid these problems to the best of their ability and allows them to get at the robust intuitions that Kauppinen seeks.

The people matter for a folk account. Therefore, the empirical investigations experimental philosophers use provide reasonably reliable evidence for claims about ordinary folk concepts.

Potential Objections to Experimental Philosophy

The previous section handled a formal argument against experimental philosophy by Antti Kauppinen. Since experimental philosophy is relatively new to the philosophical enterprise, very little written is written about it.[30] Nevertheless, the following arguments represent some of the conversations and correspondences I have had with others who are critical of the new methodology.

Critics of experimental philosophy target the empirical character of the methods it employs. In this section, I will deal with each of the following complaints: (1) experimental philosophy is not philosophy because it fails to give us normative conclusions and (2) the empirical assessments fail to offer any substantive arguments.

Each of these complaints tries to alienate experimental philosophy from standard philosophical practice and fails to account for the strengths of experimental philosophy. In conversation, I am often confronted with the question: "is experimental philosophy *really* philosophy?" Apparently, some philosophers think it is worthwhile to demarcate philosophy from other disciplines. I am not convinced that such an endeavor *is* worthwhile, but I cannot ignore their arguments because *I* think their position lacks a motivation. I would like to show that given most

---

[30] Only a handful of publications are available addressing the new methodology employed by experimental philosophers. But a number of edited anthologies have included articles whose authors have employed methods used by experimental philosophers (see Carruthers et al. 2007a, 2007b, 2007c; Sinnott-Armstrong 2008a, 2008b, 2008c). A new anthology, edited by Joshua Knobe (2008), on experimental philosophy will be available in June 2008.

arguments in philosophy have failed to resolve problems by using one of the traditional methods I have discussed above, it seems to be a good time to accept a new method that promises a possible resolution of the debate.

The Demarcation Problem

Philosophical arguments have a special normative characteristic. They are supposed to tell us, for instance, what we ought to do or ought not to do, or what we ought to believe or ought not to believe. According to some critics, the best outcome we can achieve in experimental philosophy is that either the world is the way we think it is or the world is not the way we think it is. Nothing normative hinges on this argument because it does not tell us what we should or should not believe or do. If an argument is missing the special normative characteristic, such as in the case of experimental philosophy, then it is not a philosophical argument. The normative property demarcates philosophy from non-philosophy. Therefore, experimental philosophy is not (and should not be a part of) philosophy.

Normative conclusions are distinctive characteristics of philosophical argumentation. But the point of experimental philosophy is not to rewrite the way that we do philosophy but to improve our assumptions from which we draw normative conclusions. Our assumptions about ordinary intuitions have been incorrect. So, there is at least one normative conclusion we can draw from experimental philosophy: *it tells us what we should believe about intuitions*.

If we want a folk theory, then we should take the folks' intuitions seriously. So, we should use empirical assessments to test whether our assumptions about the folk are correct. If they are not, then we should adjust the way that we argue for normative conclusions so that they account for ordinary intuitions. If they are correct, then we have no need to worry that our argument uses ordinary intuitions inappropriately. Thus, the worry that experimental philosophy is inconsistent with a crucial aim of philosophy generally is unfounded given that it's point is only to expand the methodological arsenal we may use to draw philosophically important conclusions.

Mere Verbal Disagreement

Finally, a criticism often leveled against experimental philosophy is that it can only point to verbal disagreements without being able to settle anything substantive. When an empirical assessment is distributed to respondents, their responses might be due to certain words appearing in the vignette. Sosa writes:

> Verbal disagreement *need* not reveal any substantive, real disagreement, if ambiguity and context might account for the verbal divergence. If today I say "Mary went to the bank yesterday" and tomorrow you say "Mary did not go to the bank yesterday" we need not disagree, given ambiguity and contextual variation. The experimentalists have not yet done enough to show that they have crossed the gaps created by such potential differences in meaning and context, so as to show that supposedly commonsense intuitive belief is

> really not as widely shared as philosophers have assumed it to be (Sosa 2007, CITE).

Sosa's criticism has argued that philosophers can go on assuming their intuitions reflect ordinary intuitions because problems of contextual variation and ambiguity arise often enough that an empirical assessment may not be able to overcome them.

Experimental philosophy has serious work to do if it is to overcome Sosa's objection. No empirical assessment, as they are currently conceived, completely disqualifies contextual variation or ambiguity among respondents. When we encounter a verbal disagreement in applied ethics, it usually leads to substantive disagreement. For instance, in debates about euthanasia, two people could quarrel about death criteria, a merely verbal dispute. The underlying or substantive disagreement, however, between the two people could be accounted for in their conception of what a person is or what life is. If our verbal disagreements lead to more substantive problems more often than not, then perhaps we should not be so quick to judge that experimental philosophers need to show "beyond reasonable doubt" that there are philosophically important disagreements found in cultural or socio-economic differences (Sosa 2007, CITE).

Experimental Philosophy, Folk Intuitions, and the Problem of Action Individuation

The previous section has provided some speculative critical arguments some philosophers may have against experimental philosophy. Before I conclude the

chapter, I would like to provide some evidence showing that action individuation theorists have tried to produce a folk account. Experimental philosophers, thus, are justified in testing assumptions about ordinary intuitions and action individuation.

Action individuation theorists have asserted that their accounts should be consistent with ordinary intuitions. First, the aim of Alvin Goldman's account of human action is to explain what any ordinary person would think about human action. He writes:

> One of the purposes here is to explicate certain aspects of our common sense conceptual scheme. Thus, my analysis of action is intended to capture, as closely as possible, our pretheoretic conception of an "act" or an "action" (Goldman 1970, vi).

If an account of action fails to explain what ordinary people understand by "act" or "action," then presumably we can dispose of it. For Goldman, an expectation of action theory is that it be consistent with ordinary people's intuitions.

Second, Irving Thalberg argues that if we fail to consider ordinary intuitions about action individuation, then our theories will be consumed by either vicious circularity or mere stipulation. Thalberg oscillates between these two complaints when he is confronted by something like the following: someone saying "Smith drove his car" is to assert that Smith exemplified the property of driving his car. What is it like for a person, Smith, to go around "exemplifying" properties? Either we merely

repeat the analysis given or stipulate what it is for a person to exemplify some property. Both options is are not good, according to Thalberg. Thus he writes:

> Perhaps we should not attach any philosophical significance to crude, ordinary speech. Instead we should think only of the axiomatic theories of events and actions in which these locutions are meant to be deployed. We forget ordinary meanings of the terms, and look exclusively at the formal definitions, formation rules, substitution rules, postulates, and rules of inference. But then… we will not be able to interpret its key terms at all, and we can hardly expect it to "lay bare the nature, or ontological status" [quoting Goldman 1971] of events and deeds (Thalberg 1977, 8f).

By foregoing ordinary intuitions about action, according to Thalberg, we also have to forgo the ultimate aim and motivation for exploring action theory: an explanation of the nature of action. We also forgo the ultimate aim of conceptual analysis if the analyst has interest in exposing a folk account (Cf. Jackson 1998; and what I have said above).

Finally, Ginet has believed that we can set the action individuation debate aside if each of the accounts can be supported by ordinary intuitions. He writes, "each [of the accounts of action individuation] can find *some* support in our ordinary talk about actions" (Ginet 1991, 71). His concrete account is "better supported [because] it refines and regiments what is meant by an action beyond anything demonstrable from our ordinary talk" (ibid.). Since the other accounts are not refined enough because they depend on ordinary intuitions, Ginet has argued that we dismiss

the other accounts. The problem is that his view may be the one that a lay person supports. His assumption that his concrete account is better than the others because it does not rely on ordinary intuitions can easily be undermined if we find at least one person who upholds the same position Ginet does. His position – though hostile toward folk intuitions – seems to beg for more careful consideration of them before we set the debate aside.

Since ordinary intuitions are important for action theory and few theorists have bothered asking people for their intuitions, it is up to us to ask people for their intuitions. Then, we can build an informed account of action based in ordinary intuitions.

Conclusion

Experimental philosophy challenges the traditional view of philosophical inquiry. For this reason, it has received tremendous resistance from people who are not ready to revise philosophy's long-standing method.[31] Some have already dismissed experimental philosophy as a passing fad.[32] I do not expect to convince

---

[31] See, for example, J. David Velleman's blog, *Left2Right*, for some pointed comments about the nature of experimental philosophy. Though, it is clear from the dialogue Velleman has with Weinberg and Vargas that he has no clear understanding of experimental philosophy or its procedures. Experimental philosophy is not going to replace philosophical inquiry altogether. Its purpose is merely to help it along.

[32] Many philosophers have said so in conversation with me, including David Chalmers, Terry Horgan, Bob Barnard, and Michael Tye. Barnard thinks that experimental philosophy is nothing but logical positivism returned from the philosophical dead. We have had many quarrels over the relationship

them that experimental philosophy is a worthwhile endeavor. But I do want conceptual analysts, like Frank Jackson, to stop talking about the folk if they care not to ask them for their intuitions.

In this chapter, I have argued that Kauppinen's criticisms of experimental philosophy fail for several reasons. We must revise philosophical methodology if we want to formulate a folk conception of a philosophical topic. I want to suggest that the methodology include asking people for their intuitions.[33]

The first two chapters have outlined a motivation for interest in the problem of action individuation and argued for incorporating experimental philosophy methods into the debate. Chapter 4 will summarize the major views in the action individuation debate. In particular, it will situate the contribution that this dissertation will make to the field, more generally.

---

between logical positivism and experimental philosophy. Recently, however, Barnard has embraced experimental philosophy and its methods (in spite of his reservations).

[33] With their intuitions in hand, we may begin to draw some normative conclusions about folk concepts. A substantial part of this chapter will involve an argument showing that we can draw normative conclusions from descriptive information. In fact, I will show that standard epistemological analysis does so already.

CHAPTER 4

ACT INDIVIDUATION: THE CURRENT STATE OF PLAY

Introduction

The previous chapter outlined a defense of experimental philosophy. The aim of this chapter is to provide an overview of the literature on the problem of act individuation. This chapter describes the state of the debate from its historical precursors in the works of H.L.A. Hart and Georg von Wright to the componential view of Irving Thalberg. By the end of this chapter, we will have reviewed the state of the debate on the problem of action individuation. Chapter 5 will discuss the results of an experimental study on act individuation and how that study affects two common invariant accounts of act individuation.

The action individuation debate focuses on the question: what distinguishes one action from other actions? An example serves as an illustration of the focus of the debate. Suppose that Fletch does each of the following things at the same time: (1) he stands up, (2) he startles the watchman, (3) he says, “I didn't want to do this, but I'm going to have to pull rank on you. I'm with the mattress police. There’s no tag on this mattress,” (4) he lies about his identity, and (5) he stares down the barrel of a

shotgun. How many actions has Fletch performed? People interested in this question have engaged in a debate over individuating action.

In the treatment of action individuation which follows I begin with a brief historical look at the problem and how the problem evolved. Next, I offer a description and summary of the three major positions in the debate. I will call the position that argues each action description designates one action as the *minimizing* view, the position that argues different action designators always refer to distinct actions as the *maximizing* view, and the intermediate accounts I will call the *componential* view.[34] Sections 3, 4, and 5 will give a summary of the positions, respectively.

The Early Beginnings of the Problem of Action Individuation

Two theses have been ignored by those working on the action individuation debate. But these views are nevertheless important for an understanding of action theory around the time the problem of action individuation becomes a popular avenue of research. Von Wright's and Hart's approaches can be reinterpreted as proposals of criteria for distinguishing those properties that do and those that do not individuate

---

[34] Various authors have used different terms for the positions in the debate. For example, Thalberg calls the *minimizing* view the unifier approach and the *maximizing* view the multiplier approach (Thalberg 1971, 781). Goldman refers to the *minimizing* view as the identity thesis (Goldman 1970, 2). Davis has different names for each position: the *minimizing* view is the austere theory, the *maximizing* view is the prolific theory, and the *componential* view is the moderate theory (Davis 1979, 27-41). Moreover, there is very little consensus among the proponents of the *componential* view. What I have done in this essay is summarize the more popular *componential* views.

action. For von Wright, it was the distinction between a result and a consequence, and for Hart it was ascription of responsibility.

Von Wright's Intention-based Theory of Action Individuation

One of von Wright's contributions to action theory is the distinction between a result and a consequence. According to von Wright, an action is the intentional production of change in the world. By "change," von Wright means, "the transition from one state of affairs to another" (von Wright 1963, 36). Hence, a description of an action contains three elements: (i) the state-of-affairs of the world at the beginning of the action, (ii) the state-of-affairs of the world after the action has been completed, and (iii) the state-of-affairs in which the world would be if the agent had not interfered with it.

The three elements of an action description tell us about the change that takes place in the world. By (i), we are alerted to the way the world is at the beginning of the action. Suppose, for example, the description is "Boswell's opening the window." The description implies, by (i), that the window was shut moments before Boswell changed the window's position. (ii) tells us about the state of affairs of the world after the action has been completed is what von Wright calls the "result." Since the window was shut before Boswell changed the position of the window, the result of Boswell's action is that "the window is open." Finally, (iii) has to do with the consequence of the action. The consequence of action is brought about by the "result" of performing some action. In Boswell's case, had she not opened the

window the room would have been very hot. (iii) is a description of the way the world would have been if the action had not taken place. (iii) tells us the reason for the action.

The "result," the state of the world after an action has been completed, plays an important role in von Wright's action theory. He writes:

> By the *result* of an act we can understand either the change corresponding to this act or, alternatively, the end-state […] of this change. Thus, by the result of the act of opening a certain window we can understand either the fact that the window is opening (changes from closed to open) or the fact that it is open. (von Wright 1963, 39)

There are three things we ought to notice about von Wright's "result." First, if the final state of affairs is not brought about, then the action has not taken place. An action cannot be described as an act of an individual if the result corresponding to the respective action is not produced. Von Wright's definition of action requires that a result be produced by some action because an action, on his view, is an intentional production of change. If the world were not changed, then the action would have no result. The relationship between an action and its result, i.e., the change that has taken place, is a conceptual one.

Second, the result and consequences of an action are differently related to the action. The result of an action may give rise to other transformations in the world. In this case, the result of an action is the cause of these transformations. Von Wright

defines the "consequences" of action as changes that are causally produced by the result of an action. Whereas the relationship between an action and its result is of a conceptual nature, the relationship between an action and its consequences is of a causal nature. An action necessarily has a result, but it does not necessarily have consequences.

Finally, the distinction between results and consequences of an action depends on which action is under consideration. For example, suppose that Harrison starts the engine by turning the key. When we say that Harrison (a) started the engine by (b) turning the key, the change in position of the key is the result of (b) and the change in the state of the engine is the consequence of (b). Also, the change in the state of the engine is a result of (a).

The agent's intention plays a pivotal role in the distinction between a result and a consequence. Since an action may consist in several causally related phases, we can single out one of them as the object of the agent's intention. It is the thing which the agent intends to do. What the agent intends to do is the result of his action.

Some actions have unintended results. For example, if Bulger intended to teach Flynn a lesson and three people die because Bulger's teaching Flynn a lesson even though he did not intend for them to die, then Bulger's killing the three people was not a part of the result of Bulger's action. If the agent did not intend to do something, then that is not a result of his action.

The agent may intend to do something but fail to achieve the intended results. Since an action is the intentional production of change in the world, the failure to

achieve the intended result does not mean it is not the result of the agent's action. For whatever reason, the world did not come out the way the agent intended it to be, but the state of affairs of the world was brought about by the agent's intention.

According to von Wright, if we want to know what an agent has done we must find out what his intention was. An example to illustrate von Wright's point is the story of *Oedipus Rex*. An oracle predicted that the King of Thebes and Jocasta would bear a son who would kill his father and marry his mother. When Oedipus was born of the King of Thebes and Jocasta, they abandoned him and he was later adopted by a couple from Corinth. Another oracle repeated the prophecy to Oedipus. Oedipus fled the city so that the prophecy did not come true, but he did not know that the couple in Corinth had adopted him. On his way out of Corinth, Oedipus met the King of Thebes on the road. The King refused to let Oedipus pass, so Oedipus killed him. When he reached Thebes, Oedipus solved the riddle of the Sphinx and as a reward was married to Jocasta, not knowing she was his mother.

We can describe Oedipus's actions in at least four ways: (A) Oedipus killed a traveler; (B) Oedipus killed his father; (C) Oedipus married Jocasta; and (D) Oedipus married his mother. It follows from Von Wright's theory that only (A) and (C) describe actions performed by Oedipus because only they correspond with Oedipus's intentions. For von Wright, only movements that conform to the intention of an agent are actions. One may notice that (A) and (B) include the same movements. So, it appears that (B) is an act of Oedipus's. But, since (B) does not conform to Oedipus's intention, (B) is not an action of Oedipus's. It seems that we cannot say that (B) and

(D) were actions of Oedipus. According to von Wright, the decisive property for the individuation of actions is intention and the change represented in the intention.

Von Wright's position has a serious disadvantage because it only takes into account the viewpoint of the agent. Von Wright seems to believe that actions are individuated by the agent's intention. From the agent's point of view, the most natural interpretation of the result of the agent's action is one that accounts for the agent's intention. Von Wright summarizes his position:

> When the outer aspect of an action consists of several causally related phases, it is normally correct to single out *one* of them as the *object* of the agent's intention. It is the thing which the agent *intends to do*. This is the result of his action. (von Wright 1971, 89)

This does not mean there is not a fact of the matter about the agent's intention. There is a fact of the matter, even though observers have no access to the agent's intention. An observer's interpretation of the agent's intention might differ significantly from that of the agent. These other descriptions might be relevant to the agent's intended action because they are consistent with the transitional change that took place. But, since the interpretation was not the same as the agent's intention, the observer cannot ascribe the action to the agent. Moreover, since von Wright fails to explain whether unintentional actions have results and what they are, it seems that we cannot even say in the case above that (B) and (D) are consequences of Oedipus's actions, but not themselves actions of his. Ultimately, according to von Wright, the property that is

decisive for the individuation of actions is agent's intention and the change represented in that intention.

H.L.A. Hart's Ascriptive Theory of Action Individuation

Whereas von Wright's view accounts for actions solely from the point of view of the agent, Hart's proposal ignores the agent's intention in favor of an attributive view of action. Hart's proposal cites ascription of responsibility for choosing which properties among all those exemplified by a sequence of bodily movements or by an agent that permit the individuation of actions. He writes:

> The philosophical analysis of the concept of human action has been inadequate and confused, at least in part because sentences of the form 'He did it' have been traditionally regarded as primarily descriptive, whereas their principal function is what I venture to call ascriptive, being quite literally to ascribe responsibility for actions much as the principal function of sentences of the form 'This is his' is to ascribe right in property. (Hart 1948/49, 145)

According to Hart, actions and legal concepts are related because they are revocable or defeasible. The definition of a legal concept includes certain conditions of exception, or negative circumstances, which exclude the application of the concept or stipulate that it be applied in a more moderate or partial form. Just as legal concepts cannot be given precision by specifying a number of conditions that are necessary and jointly sufficient, the same applies to action. He writes:

> Our concept of 'action', like our concept of property, is a social concept and logically dependent on accepted rules of conduct. It is fundamentally not descriptive, but ascriptive in character; and it is a defeasible concept to be defined through exceptions and not by a set of necessary and sufficient conditions whether physical or psychological. (Hart 1948/49, 161)

According to Hart, statements of the kind 'S did *x*', where S stands for an agent and *x* refers to an action performed by S, are not descriptive, but ascriptive or attributive statements. Their primary function is to ascribe responsibility to an agent.

'Ascription' and 'responsibility' can be understood in a conceptually precise way. For ascription, we can distinguish between a weak and a strong sense of the term. In the weak sense, 'ascribing' is equivalent to answering the question "Who did *x*?" Whereas the weak sense of ascribing requires an answer to the question "Who did *x*?," a description requires an answer to the question, "What did S do?" In the strong sense, ascription implies a degree of discretion, is relative to a context, and it is defeasible. Ascription in the strong sense is a matter of decision, and it is not a matter of discovery. To ascribe one action or other to some agent, it depends on the context.

Joel Feinberg has extended Hart's ascriptivist theory of action. Feinberg provided some further explication of responsibility. He distinguishes between five senses of the expression "ascription of responsibility" (Cf. Feinberg 1968, 116). According to Feinberg, the five senses of the expression "ascription of responsibility" include: ascription of causal relationship, ascription of causal agency, ascription of

simple agency, allegation of fault, and ascription of strict liability. (1) *Ascription of causal relationship*: when we say "S was responsible for *x*," we may mean that *x* was the causal result of S's action, without implying any judgment of the action. (2) *Ascriptions of causal agency*: responsibility in the causal sense can be replaced by ascriptions of agency. For example, on this view, "Harrison caused the engine to start" and "Harrison turned on the engine" show that Harrison is causally responsible for the engine's turning on. (3) *Ascription of simple agency*: a basic action is one that is not performed by way of another action. For example, Harrison smiling and Harrison lifting his arm are examples of an agent performing a basic action. (4) *Allegation of fault*: some ascriptions may find an agent blameworthy. (5) *Ascription of strict liability*: some ascriptions of an agent's actions may justify punishing the agent.

Feinberg's extended analysis and his introduction of five senses of "ascriptions of responsibility" are consistent with Hart's ascriptive theory of action. Hart's ascriptive theory of action is a social concept. The criteria for determining what action someone has performed depend on social, moral, and legal conventions and rules. According to Hart, all cases of individuating action are instances of ascription in the strong sense. Similarly, we can use the five senses of Feinberg's ascriptions of responsibility to individuate action.

The problem for the ascriptivist account contrasts with that of von Wright. Since the ascriptive theory of action depends on social, moral, legal conventions and rules, it fails to account for the point of view of the agent. When an observer ascribes

an act to an agent, she should be careful to ascribe that act correctly. If an agent were ascribing an act to herself, then she would not have to be as careful because she is the author of the act. An agent has special access to her intention because she is well acquainted with what she is doing, has done, or will do. She is well acquainted with the act. What she says about her own action is presumably correct. When an observer ascribes an act to an agent, she may account for contextual features that affect the agent. The agent performing the action may not be aware of these features. So, ascribing an act to another may incorrectly interpret the agent's action.

Von Wright cannot adequately account for non-intentional action, and Hart cannot adequately account for intentional actions, unless we would be willing to accept that we ascribe in the strong sense intentions to ourselves.

The Minimizing View

Neither von Wright's theory nor Hart's theory seem to have influenced the three popular views of act individuation, i.e., the minimizing view, the maximizing view, or the componential view, directly. But the three popular views may have been indirectly motivated by von Wright's and Hart's theories. Moreover, von Wright and Hart's accounts are interesting in their own right and deserve a place in the history of the problem of act individuation. Von Wright emphasized the agent's intention for divvying up action, which is a popular view still held by action theorists today, e.g., Alfred Mele, and Hart's account of act individuation focused on ascriptions. The

three popular views deviate somewhat from von Wright and Hart's theories. They suggest that the concern of act individuation is the notion of identity.

Advocates of the minimizing view argue that an agent performs one action, which can have many descriptions. The source of the minimizing view is in the works of Anscombe and Davidson. This section will summarize the arguments of the minimizing view. At the beginning of the next section, I will review the arguments for rejecting the minimizing account.

Discussion of the problem of action individuation was initiated by Anscombe's question:

> Are we to say that the man who (intentionally) moves his arm, operates the pump, replenishes the water supply, poisons the inhabitants, is performing *four* actions? Or only one? (Anscombe 1957/2000, 45)

Anscombe has given us four action descriptions that contain cognates of the action verbs "move," "operate," "replenish," and "poison." Some action must answer to each of these descriptions. So, the question is whether the same action answers to all of these descriptions or whether each description has a different action answering to it.

Anscombe argues that a variety of action descriptions refer to the same action. She writes (with reference to the man operating the pump):

> For moving his arm up and down with his fingers round the pump handle *is*, in these circumstances, operating the pump; and, in these circumstances, it *is* replenishing the house water-supply; and, in these circumstances, it *is* poisoning the household. (Anscombe 1957/2000, 46)

Her answer serves the purpose of having a clear understanding of whether a number of different descriptions designate the same action; it also is meant to elucidate how action descriptions may serve to specify intentions and how intentions behind an act are related to one another. If we ask the man why he moved his arm up and down, then he may answer that he intended to operate the pump. Since his intention was to operate the pump and moving his arm up and down was a part of operating the pump, we can say that a characteristic of his moving his arm up and down was intentional. Similarly, when we ask him why he operated the pump, he may answer that he intended to operate the pump. His intending to operate the pump was the intention with which the act in each of its other descriptions was done. Anscombe writes:

> One action with four descriptions, each dependent on wider circumstances, and each related to the next as description of means to end… speak equally well of *four* corresponding intentions, or of *one* intention – the last term that we have brought in in the series. By making it the last term so far brought in, we have given it the character of being the intention… *with* which the act in its other descriptions was done. (Anscombe 1957/2000, 46)

When Ansombe talks about four intentions, she is speaking of "the character of being intentional" that belongs to the act in each of the four descriptions. When Anscombe

speaks of *one* intention, she is talking about *the intention* with which the act in each of its other descriptions was done.

A corresponding minimizing account of action individuation appears in Davidson's work on action and event individuation. Davidson's view on the question of action individuation is *prima facie* the same as Anscombe's.[35] He too endorses the idea that a variety of action descriptions can designate the same action. He endorses the minimizing account in various places:

> I flip the switch, turn on the light, and illuminate the room. Unbeknownst to me I also alert a prowler to the fact that I am home. Here I do not do four things, but only one, of which four descriptions have been given. (Davidson 2001, 4)

> That the bullet pierced the victim was a consequence of my pointing the gun and pulling the trigger. It is clear that there are two different events, since one began slightly after the other. But what is the relation between my pointing the gun and pulling the trigger, and my shooting the victim? The natural and, I think, correct answer is that the relation is that of identity. (Davidson 2001, 109)

---

[35] Julia Annas (1976) has argued that Anscombe and Davidson may argue for the same conclusion, but they do not have the same argument for that conclusion. According to Annas, Anscombe is committed to something much weaker "involving essential reference to means-end chains" than what Davidson argues in his minimizing view. Suppose that the man is moving his arm up and down (A), operating the pump (B), replenishing the water supply (C), and poisoning the inhabitants (D). Anscombe's point is that we can say that we have one action under different descriptions if the descriptions are related as descriptions of means to descriptions of ends (Anscombe 2000, p. 46). Annas, consequently, claims that, on Anscombe's view, "we have to be able to say only that B is the same A as C and C is the same B as D, not that A,B,C, and D are all the same F." Annas contends that Davidson is not interested in means-end reasoning but is interested in the proper canonical notation for action sentences.

Despite Annas's effort to show how Davidson and Anscombe argue differently for the same conclusion, Anscombe later admitted that her view was the same as Davidson's. So, it is hard to reconcile Annas's argument with what Anscombe said herself in a later article (Cf. Anscombe 1979).

The descriptions "Donald's pulling the trigger" and "Donald's shooting the victim" designate the same action.

Davidson's identity theory or minimizing account of action individuation depends on two things: his view of event individuation and what counts as an action.

First, Davidson provides a criterion for the identity between events. He writes, "events are identical if and only if they have exactly the same causes and the same effects" (Davidson 2001, 179). More formally, Davidson writes $x = y$ iff ((z) (if z caused x, then z caused y) and (z) (if x caused z, then y caused z)) (ibid.).[36] On Davidson's account of event individuation, to describe an event as a killing where the killing is caused by a shooting or by a person's pulling a trigger *is* to describe an event as causing a death (Davidson 2001, 177). This is the first step that implies different descriptions designate the same action.

Second, we need to know that actions only consist in bodily movements. Davidson argues this way:

> This welter of related descriptions corresponds to a single *descriptum*… When we infer that he stopped his car from the fact that by pressing the pedal a man caused his automobile to come to a stop, we do not transfer agency from one event to another, or infer that the man was agent not only of one action but of two. We may indeed extend responsibility or liability for an

[36] Ernest Lepore (1985, 156f) has mentioned that Davidson changed the formal definition later, though it retained some of the formal features of the original definition, e.g., causal features.

> action to responsibility or liability for its consequences, but this we do, not by saddling the agent with a new action, but by pointing out that his original action had those results. (Davidson 2001, 59)

What we may conclude from this is that mere bodily movement is all that there is to an action. Accordingly, we "never do more than move our bodies: the rest is up to nature" (Davidson 2001, 59).

According to the minimizing view of action individuation, a bodily movement gives rise to a chain of events. Each event allows a different description of action. Thus, for each bodily movement, we have a number of descriptions that satisfactorily describe that single event.

Objections to the *Minimizing* View

The previous section outlined the minimizing account of Anscombe and Davidson. Since the driving force behind the maximizing account is the objections to the minimizing account, I begin this section with an outline of three major objections to the minimizing account. Once I provide the three standard arguments against the minimizing account, I will outline the positive argument for the maximizing account. At the end of the section, I will provide a few objections to the maximizing account. The next section will summarize a moderate theory of action individuation, what I term the *componential* view.

Goldman has proposed three objections to the minimizing view. The three objections are: the *causal explanatory relations* problem, the problem of the *by-locution*, and the *temporal order relations* problem.[37] I will reconstruct an argument for each of these objections.

First, the argument from causal explanatory relations shows that identity claims are at odds with the minimizing view's notion that a variety of descriptions can refer to the same action. Suppose we have two descriptions: "Smith killed Jones" and "Smith fired the gun." Something enters in the causal explanation of "Smith killed Jones" that does not enter into the causal explanation of "Smith fired the gun." "Smith fired the gun" *because* the movement of the gun's trigger causes the gun to fire and Smith's finger moved in such a way that he pulled the trigger. Smith killed Jones *because* Jones had insulted him earlier. But Smith's killing Jones does not enter into the causal explanation of Smith's firing the gun. Goldman writes, "If we find… that [two descriptions] have somewhat different sets of causes or causal factors, that would give us reason to conclude that [the two actions] are not the same after all" (Goldman 1970, 3; Cf. Goldman 1971, 767). This implies that Smith's firing the gun refers to a different action than the one that describes Smith's killing Jones.

A primary reason for thinking that Smith's killing Jones did not cause the gun to fire is that it sounds odd or seems false. That a statement seems to be false or that

---

[37] These labels are Carl Ginet's (1990).

it sounds odd is not a good reason to think there is something wrong with it. Claims that "it sounds odd" or "it seems false" have to be bolstered by the fact that if it is true to say that Smith's killing Jones caused the gun to fire, then that would commit us to holding that Smith killed Jones before the gun fired (Goldman 1970, 2). Such an explanation cannot be true because the bullet fired from the gun killed Jones; it is Smith's pulling the trigger that caused the gun to fire, not Smith's killing Jones.

A minimizer might deny that the description is not extensional in the way the objection contends it is. Defenders of the minimizing view may assert that the relation of causation holds between two events only under certain descriptions. For example, in the Oedipus story, minimizers may say that "Oedipus wanted to marry *x*" holds true under some descriptions, e.g., Oedipus wanted to marry the Queen of Thebes or Jocasta, and not true under other descriptions, i.e., Oedipus wanted to marry his mother. Therefore, on the minimizing account the causal relationship is intensional.

The primary problem with the minimizer's response to the argument from causal explanatory relations is that it makes causation language dependent. Oedipus wanted to marry *x*, but a description of *x* is the Queen of Thebes, Jocasta, and his mother. Regardless, each of these descriptions designates *x*. If the above argument were true, then Oedipus's desire to marry *x* has more to do with the semantic description of *x* than *x* herself. Thus, we cannot say that the relation of causation depends on language.

The second argument against the minimizing view argues that when two action descriptions are connected using the word "by," the two descriptions do not refer to the same action. Since we often speak of one act being done or performed *by* doing another, we might be tempted to think that the two descriptions denote the same action. We think that the two descriptions denote the same action because the first is somehow generated by the second. For example, I entered freeway traffic by stepping on the accelerator. One can show that two action descriptions do not designate the same action if the relationship between the descriptions is asymmetrical and extensional. If two or more action descriptions are asymmetrical and extensional, then that provides some evidence supporting the idea that two or more action descriptions designate distinct actions.

When we say that "Smith killed Jones by firing the gun" and "Smith fired the gun by pulling the trigger" and "Smith pulled the trigger by moving his finger," the relationship between these descriptions is not arbitrary. The relationship between them is one of means and ends, not of identity. While we say that "Smith killed Jones by firing the gun," we cannot say with the same sense that "Smith fired the gun by killing Jones." The relationship between the two descriptions is asymmetric. Moreover, the BY-relation is extensional. This means that the relationship between action descriptions is irreflexive (Goldman 1970, 5). For example, we would not say for pragmatic reasons that Smith killed Jones by killing Jones, or that Doe turned on the light by turning on the light. We cannot explain how Smith killed Jones by indicating that he killed Jones, and we cannot explain how Doe turned on the light by

indicating that she turned on the light. Goldman has seized upon the non-symmetry of the relevant properties (Castaneda 1979, 242f).

According to the minimizing view, any two actions are identical only if the relations that hold between them are equivalence relations. What we have shown through the discussion of asymmetry and irreflexivity is that the by-relation is not an equivalence relation. Thus, no acts standing in the by-relation to one another are identical.

Finally, the argument from temporal order relations shows that two action descriptions fail to identify the same action when the two descriptions designate distinct moments in time (Goldman 1970, 5; Thomson 1971). Suppose that Jones dies of a gunshot wound hours (or perhaps days or months) following Smith's shooting him. Smith shoots Jones at t_1 and Jones dies at t_2. So, Smith's shooting Jones occurred before Jones died, even though Jones's death is hastened because Smith shot him. While it is true that Jones death occurs after Smith shot him, it is false to say that Jones death occurs after Smith kills him. Thus the shooting and the killing must be distinct since one event seems to have a property the other lacks.

These three arguments seem to have shown that the minimizing view for individuating actions is inadequate.[38] To remove the minimizing account is not enough to solve the problem of action individuation. It is just that we have one less option to consider. We need a positive account that distinguishes actions to secure an

---

[38] Castaneda seems to think that at least the second argument from by-relations fails to undermine the minimizing view, and I tend to agree with him. He argues that…

alternative position in the debate. What follows is a summary of the positive account of the maximizing view. After I outline the maximizing account, I will show why the maximizing account does not run into the same three problems the minimizing account does. Finally, I will provide some objections to the maximizing account.

The Maximizing View

The maximizing view has an underlying rationale. According to the maximizing view, actions appear to involve and can be characterized as the loss, retention, acquisition, or having of properties by an object at a time. This implies that actions are exemplifications of properties.

Goldman's proposed solution to the problem of action individuation begins with a distinction between what he calls act-types and act-tokens. Act-types are a kind of property. He writes:

> An act-type, as I construe it, is simply an act property, something that an agent exemplifies. When we say, "John weighed 170 pounds" or "John was bald," we ascribe to John the property of weighing 170 pounds or the property of being bald. Similarly, I suggest that when we say "John signaled for a turn" or "John killed George," we ascribe act properties or act types to John: the property of signaling for a turn or the property of killing George. (Goldman 1971, 769)

An act-type is a property of the agent. The action itself is the agent's exemplifying that property at that time. If an agent exemplifies two properties at the same time, then there are two exemplifyings, two such properties at the same time. Thus, there are two distinct actions.

An act-token is an exemplifying of an act-type by a particular agent at a particular point in time. Goldman's account of act-tokens is:

> To perform an act, then is to exemplify a property. To perform the act of giving a lecture is to exemplify the property of giving a lecture. A particular act, then consists in the exemplifying of an act-property by an agent at a particular time. I shall call such particular acts: "act-tokens." An act-token is not itself a property. It is the exemplifying of a property by an agent. (Goldman 1970, 10)

The terminological distinction between act-types and act-tokens provides a criterion for the identity of an action. The criterion of identity is:

> Since an act-token is the exemplifying of a property by an agent at a time, it is natural so to individuate act-tokens that *two act-tokens are identical if and only if they involve the same agent, the same property, and the same time*. (ibid.)

Suppose there is an action x that has the property P at time t_1. According to Goldman's maximizing account, if there is some action y that has property Q at time t^*, then the identity condition for the two actions x and y is that they are the same if

and only if $x = y$, $P = Q$, and $t_1 = t^*$. Otherwise, x and y are distinct act-tokens that exemplify distinct actions.

The distinction between act-types and act-tokens lead to the criterion that shows an agent performs distinct actions. For example, if we apply Goldman's criterion to the Smith case, moving his right index finger, pulling the trigger, firing the gun, shooting Jones, and killing Jones are different actions because each one exemplifies a different property. They are distinct actions, but Goldman contends that there is a certain relationship between them. The relationship between them is not one of identity. If it were identity, then Goldman's account would have to overcome the same problems faced by the minimizing account. According to Goldman, there are generational relations between the actions produced by one and the same basic action.

According to Goldman, there are four kinds of generational relations: causal, conventional, simple, and augmentation. First, there is causal generation. Causal generation is the relationship between actions that depends on relationships of causality. Goldman writes:

> *S*'s act-token *A* has a certain effect, *E*, and because it has this effect, *S* may be credited with performing act *A'*. For example, *S*'s flipping the switch has the effect of the light's going on. And in virtue of this, *S* may be credited with the act of turning on the light. That is, we may say that *S* exemplified the property of turning on the light… To generalize: *Act-token* A *of agent* S *causally generates act-token* A′ *of agent* S *only if* (a) A

> *causes* E, and (b) A′ *consists in* S's *causing* E. (Goldman 1970, 22f)

Thus, Smith's pulling the trigger causally generates Smith's firing the gun since Smith causes the gun to fire, and firing the gun consists in Smith's pulling the trigger.

Second, an action performed in certain circumstances produces another action because there is a rule to that effect. This is what Goldman calls conventional generation. *S*'s performance of *A* justifies the ascription of *A′* when and only when some rule *R* exists as a justification for that ascription. For example, suppose that Carne, a train conductor, extends his arm horizontally at shoulder height. Given that extending one's arm horizontally at shoulder length along a train track counts as signaling for the fireman to reduce speed, Carne has signaled for the fireman to reduce speed. Such examples may be expressed by the following condition:

> *Act-token* A *of agent* S *conventionally generates act-token* A′ *of agent* S *only if the performance of* A *in circumstances* C *(possibly null), together with a rule* R *saying that* A *done in* C *counts as* A′*, guarantees the performance of* A′. (Goldman 1970, 26, his italics)

Third, simple generation relation occurs when an action is a sufficient condition for the production of another action. But the explanation of the simultaneous action is neither a case of causal nor of conventional generation. According to Goldman, simple generation is "the existence of certain circumstances, conjoined with the performance of *A*, ensures that the agent has performed *A″*" and

"simple generation is like conventional generation minus the rules" (ibid.). For example, suppose that John Kerry asserts that we should not have gone to war with Iraq. Together with the fact that Kerry supported the war with Iraq at an earlier time, this implies that Kerry contradicted his earlier statement. The relevant circumstance is Kerry's claim that we ought not have gone to war (call this *C*) and Kerry's earlier support of the war (call this *A*) imply that Kerry contradicted his earlier statement (call this *A'*). Simple generation can be schematized as *A* and *C* imply *A'*, while conventional generation also requires a rule *R* jointly implies *A'*. Kerry's "flip-flop" on the Iraq issue is an example of simple generation in Goldman's account.

Finally, when an action is modified as to manner and circumstance, Goldman contends that an augmentation generation is present. For example, Carl Lewis's running 100 meters, if done in the appropriate manner, will generate Lewis's running 100 meters swiftly.

> What is distinctive in these cases is the fact that the performance of the generated act, *A',* entails the performance of the generating act, *A*. On the other hand, the performance of the generating act does not entail the performance of the generated act. The generated act is formed by "augmenting" the generating act with some relevant fact or circumstance. The fact that the generating act is not merely performed, but performed *in a certain manner,or in certain circumstances,* entails that the generated act is also performed. (Goldman 1970, 28)

The generated act of Lewis's running 100 meters swiftly entails the generating act of

Lewis's running 100 meters. The speed of Lewis's running 100 meters augments his running 100 meters. In contrast to simple generation, there is no relationship of implication for augmentation generation.

The parameters of Goldman's maximizing theory, particularly his introduction of the generation relation solves the problems encountered by Davidson's minimizing account. I will now review how Goldman's theory may accommodate the problems.

The first problem Davidson's minimizing account faced was the argument from causal explanatory relations. On the minimizing account, the action descriptions identify the same action and that was a problem because it was Smith's pulling the trigger and not Smith's killing Jones that caused the gun to fire, which caused Jones's death. The trigger pulling and the killing are not unrelated. The principle of causal generation says that act-token *A* of agent *S* causally generates act-token *A′* of agent *S* only if (a) *A* causes *E* and (b) *A′* consists in *S*'s causing *E*. For the Smith example, then, *A* is Smith's pulling the trigger, *E* is Jones's death, and *A′* is Smith's killing Jones. Thus, on Goldman's account, causal generation is met because Smith's pulling the trigger causes Jones's death and Smith's killing Jones consists in Smith's pulling the trigger causing Jones's death. It is the case that one's pulling the gun's trigger, shootings, and other actions typically generate killings.

The second problem is the argument from BY-relations. Since any acts *A*, *A′* are identical only if the equivalence relation holds between them and since the by-relation is not an equivalence relation, no acts standing in the by-relation designate the same action. Goldman asserts that when we speak of Smith's killing Jones by

shooting him, we express a certain relationship between the two acts that is asymmetric, irreflexive, and transitive. The four generation relations are built to handle the relationship expressed in the by-relation. So, the maximizing view does not fail in the way that the assumption of identity for acts fails because the generation relations do not presume an equivalence relation between the two acts.

Finally, the argument from temporal order has shown us that since Jones's death occurs well after Smith shot him, the shooting of Jones is not identical to the killing of Jones. What makes the temporal order objection more problematic is that at the time of Jones's death, which occurs several hours, days, or months after Smith shot him, the minimizing account would have us believe that Smith killed Jones at *that* time and not when Smith shot him at the earlier time.

Goldman's maximizing account can handle the temporal order objection. One of the generation relations should permit us to say that Smith's shooting Jones and Smith's killing Jones refer to distinct acts, even though one, Jones's death, is the result of the other, Smith's shooting Jones. Goldman provides an alternative example to show that the maximizing theory, along with generation relations, can accommodate the relationship between acts and their temporal parts. He writes:

> Consider, for example, *S*'s act of driving a nail into the wall. Suppose this was accomplished by *S* striking the nail four times with a hammer… There are four relevant basic acts performed during the period in question, at time t_1, t_2, t_3, and t_4 respectively… Each of these basic acts is an act of *S*'s swinging his hand, each of which generates an act of *S*'s swinging the hammer,

> which in turn generates an act of driving the nail a little way into the wall. Thus, *S*'s swinging his hand at t_1 generates *S*'s swinging the hammer at t_1 which generates *S*'s driving the nail a little way into the wall at t_1. The sequence of these four basic acts constitutes a larger act, viz., *S*'s swinging his hand four times (between t_1 and t_4). This larger act generates *S*'s act of swinging the hammer four times (between t_1 and t_4), which in turn generates *S*'s act of driving the nail into the wall (between t_1 and t_4). None of the larger acts is generationally related to any of the smaller acts, but there are generational relationships among the three larger acts. (Goldman 1970, 35f)

The account distinguishes between the acts that take place at t_1 and at t_4 so that it is not the case that different act descriptions refer to one "single" action. We can then say that Smith's pulling the trigger, moving his right index finger, and so on generates Smith's killing Jones.

For these reasons, the maximizing account overcomes the three common objections to the minimizing view. There are problems that arise for the maximizing account, and I will now review some of these problems.

Objections to the Maximizing Account

First, Goldman introduces entities, i.e., basic acts, act-tokens and act-types, and act-trees, beyond necessity. The maximizing account increases the number of actions performed by an agent on a given occasion. Since an account of individuation

should be parsimonious, we ought to reject these entities if they are not really necessary.[39]

We might think that the maximizing account and the minimizing account differ with respect to the number of entities they introduce, but they do not. If we analyze the two positions, we will find that they are not very distantly related in terms of ontological requirements. Just as the minimizing theory recognizes an action, a description of that action, and a set of these descriptions, so too does the maximizing account. The difference is that the maximizing account uses basic action, action (act-token and act-type), and act-tree, respectively, to distinguish between actions. The debate then seems to be a verbal dispute, upon which nothing else hinges – others have emphasized this exact point (Castenada 1979 and Ginet 1990, 70).

There is a second more serious problem with the maximizing view. Since there are as many act-tokens as there are act-properties, it follows that the number of actions an agent can perform at any one time is infinitely large. On this view, the actions an agent performs at one moment in time is impossible to count. But an agent's actions are limited to the space and time it occupies, so the number of actions should be countable. Thus, the maximizing view is mistaken in asserting that there is no limit to the number of distinct actions an agent performs at any one time.

---

[39] It has been pointed out to me that this objection may be begging-the-question against Goldman. I agree, but it is one of the most common problems raised against Goldman's account of action individuation. So, I chose to include it among the objections I recited in this review of literature.

The Componential View

The maximizing and minimizing view of action individuation represent two-thirds of the positions in the debate. They are the diametrically opposed views in the debate. The source of the disagreement is how many distinct particular actions an agent performs when she moves her body. For the maximizer, there are as many actions as there are action descriptions, and, for the minimizer, there is one action for which we can have a variety of descriptions.

The middle position of the debate is unsatisfied with both the maximizing account and the minimizing account. The middle position is more difficult to define than the other positions in the debate. Thalberg has asserted that the middle position (particularly that view offered by Lawrence Davis) "wants to hold down the birth rate of act tokens" (Thalberg 1971, 786). According to the middle position, when we have a number of different descriptions, e.g., "the Prime Minister's moving his finger," "his pressing a button," or "his destroying an ancient capital city," there are not separate distinct actions here but events in the series have components in common. An event added on to the action of the Prime Minister's moving his finger yields the action of his pressing the button (Hornsby 1979, 195f; Hornsby 1981, 18-21; Weil and Thalberg 1981, 12ff). I will call the middle path the *componential view*.[40]

---

[40] Irving Thalberg has termed his own view the non-reductive unifying account (Thalberg 1977, 109).

A brief summary of the three positions may be useful here. The debate on action individuation has been outlined in numerous places. Thomson's summary is probably the clearest. She writes:

> There is, first, what we might call the "relaxed view," according to which we are to identify the replenishing with the pumping, a killing with a shooting with a pressing of a trigger, a flipping of a switch with an alerting of a prowler. It is this view which I think Davidson's causal criterion for act identity was meant to express. There is, third, Goldman's "extreme view" [i.e., no replenishing of the water supply by operating the pump by Sebastian is identical with any replenishing of the water supply by Sebastian]. But there is, between them, a "middle ground," according to which we may not identify a replenishing with a pumping, but may, and indeed should, identify a replenishing with a replenishing by pumping, and that with a replenishing with a pump; according to which we may not identify a killing with a shooting with a pressing of a trigger, but may, and indeed should, identify a killing with a killing by shooting, and that with a killing with a gun; and so on. (Thomson 1971, 780)

Other accounts of the middle (componential) view do not necessarily entirely agree with Thomson's assessment. What we should glean from Thomson's summary is that a third position may be staked out.

The aim of this section is to summarize a few arguments for and against the componential view of action individuation. In this section, I will review two of the most prominent componential views. First, I will summarize Irving Thalberg's arguments for a componential view. On Thalberg's account, since we may

distinguish between basic and non-basic actions and since some basic and non-basic actions overlap with one another in a special way, we may conclude that a single act can be of more than one type (Thalberg 1977, 85-129). Then, I will provide a summary of Ginet's "concrete" account of act individuation (Ginet 1990, 45-71). On Ginet's "concrete" account, he explains the necessary and sufficient conditions for it to be the case that two different action descriptions designate one and the same concrete action.

Irving Thalberg outlines a componential account of action individuation in a series of papers (Weil and Thalberg 1974, Thalberg 1971) and a chapter in his book *Perception, Emotion, and Action* (Thalberg 1977). According to Thalberg's componential view, the problems facing the *maximizing* view and the *minimizing* view can be convincingly handled by making use of the notion of components (or parts) of actions.

Thalberg's componential account largely consists in two theses. First, the description of an action entails something about the agent's state of mind. A person cannot be said to have acted if he did not have the appropriate mental states. Second, an event that a person brings about can be described by some kind of causal analysis in terms of these appropriate mental states. The relevant mental states cause the

bodily movement. Thus, both the mental states and the bodily movement of an agent are components of an action.[41]

An example will help clarify Thalberg's account. Suppose that the Prime Minister of Gulbrin and his National Security Board hold an emergency meeting. They are convinced, erroneously perhaps, that the country of Bringaray, Gulbrin's sworn-enemy, is about to launch a nuclear attack. The Prime Minister has a red button on his desk that if pushed triggers the launch of a dozen nuclear missiles which are all aimed at key targets in Bringaray. The missiles will launch from their silos within two hours after the button is activated. These missiles will strike their intended target four hours after the button is pushed by the Prime Minister.

An inventory of the Prime Minister's will help us understand the componential account:

> (RFA): The Prime Minister decides to press the red button;
> (BP): Events in the Prime Minister's cerebral cortex occur that set off happenings in his body;
> (FP): Neural feedback processes travel through his nervous system;
> (BM): Finally, his left index finger moves downward.[42]

---

[41] One might believe that not much distinguishes Thalberg's two theses from Davidson and Anscombe. We must keep in mind that Thalberg's componential view falls somewhere in between the minimizing account and the maximizing account. Thalberg has attempted to construct a view of act individuation that incorporates as much of the two opposing views as possible while leaving aside the problems that these two positions encounter. So, if the reader believes that Thalberg's two theses reflect Davidson and Anscombe's view, then the reader has an appropriate understanding of the two theses.

[42] An interesting development has occurred recently. Some data, following on studies by Benjamin Libet (2004), have surfaced as to whether these four descriptions could be true.

The basic action of the Prime Minister crooking his left index finger consists of RFA, BP, FP, and BM. Each of RFA, BP, FP, and BM are components of the Prime Minister's basic action.

The analysis works just as well for non-basic actions.[43] Non-basic actions are those actions which the Prime Minister inaugurates in his basic actions that occur beyond his epidermis (Thalberg 1977, 89; Weil and Thalberg 1974, 115f). The Prime Minister's non-basic deeds consists of elements RFA through BM, plus one or more of the following "extra-bodily" events:

> (RC) The Prime Minister's left index finger reaches the red button.
> (E_1) The red button moves.
> (E_2) Two hours from the time the button is depressed, the missiles blast off.
> (E_3) Four hours from the time the button is depressed, the missiles raze targets in Bringaray.
> (CC) There occurs a violation of the law which allows the Prime Minister of Gulbrin to use that button only with parliamentary approval.

Just as RFA, BP, FP, and BM are components of the basic action of the Prime Minister moving his finger, so too they are components of those more complex actions represented by RC, E_1, E_2, E_3, and CC. To provide a comprehensive understanding of non-basic actions we simply add to those components associated with the basic action the happenings outside the agent's body.

---

[43] There are other ways of drawing a basic/non-basic distinction. I outline here Thalberg and his collaborators way of distinguishing basic and non-basic actions.

According to Thalberg's account, each item on this list does not refer to the same action, and each item on this list does not refer to distinct actions. It is just that RFA, BP, FP, and BM are in a sense contained by RC, E_1, E_2, E_3, and CC. So, "all we should concede… is that… [the Prime Minister] accomplished more – not that he does something extra" (Thalberg 1977, 106). What does he accomplish? When the Prime Minister bends his finger, the basic action results in the destruction of targets in Bringaray.

Time is a problem for Thalberg's componential account. In many ways, the problem resembles the kind of problem encountered by the minimizing account, i.e., the problem of temporal order relations (Cf. p. 126f). The minimizing account makes actions end sooner than they actually seem to end. For example, Smith's killing Jones ends before Jones has died, and, in fact, Smith's killing ends even before the gun has gone off. Somewhat analogously on Thalberg's account, actions go on too long. Suppose that the Prime Minister were to play golf between the time he depresses the button and the time at which the missiles hit their targets. When the Prime Minister is putting for birdie on the third hole, Thalberg's componential account entails that his putting for birdie is a part of his destroying selective Bringarayian targets. Similarly, if while the Prime Minister is playing golf a thunderstorm suddenly erupts and lightning strikes and kills the nefarious head of state, the Prime Minister is in the process of killing innocent Bringarayians after he has died.

Thalberg's view represents an alternative to the maximizing and minimizing account of action individuation. His account has shown that non-basic actions are parts of basic actions. Carl Ginet has offered an alternative to Thalberg's componential account. His view is founded on what he terms "concrete actions." His view teeters between the maximizing and minimizing view primarily because he finds "the disagreement between the extreme minimizing position and [Thalberg's account] on that question is a disagreement about how certain action-designators work to pick out particular things, not necessarily about what ontological category the things they pick out belong to." (Ginet 1990, 49) On Ginet's view, either an action is volition or an action is composed of a core action plus a layer of consequence or circumstance, or an action is a conjunction of actions.

According to Ginet, actions are events that possess a "layered structure" (Ginet 1990, 50f). At the core of an action is volition. Volition is a mental action of the agent. Parts are then added to the core action to make a larger action and parts are added to those larger actions to make larger actions, but these larger actions are typically not actions of that agent. The core action is the action of the agent, while larger and larger actions form parts added to the core action of the agent.

Core actions are not concrete actions. According to Ginet, "in order to establish that there [are concrete actions, he needs]… to give a criterion of individuation for concrete actions" (Ginet 1990, 65). He argues that the criterion must be very specific. He needs to explain, "what is necessary and sufficient for it to be the case that two different canonical action-designators designate one and the same

concrete action" (ibid.). Canonical action-designators "pick out a particular personal event or state *uniquely*" (Ginet 1990, 1). But since canonical action-designators, such as "S's making a movement with her arm just now of the same sort that caused her such pain a few minutes ago" (Ginet 1990, 45 [his example]), "fail to make fully explicit the type of the action designated" (Ginet 1990, 45f), there must be a criterion of individuation where the layered structure of actions yield that distinct actions have the same parts.

The layered structure of actions yields distinct actions having the same parts. For example, suppose that Mancuso fetches some bread because Terrell told him to do so. Mancuso's fetching some bread because Terrell told him to do so is distinct from Terrell's making Mancuso fetch some bread by telling him to do so. But the two actions have the same parts. The action has parts: (a) *S* fetching some bread, (b) *R*'s command, and (c) the causal relation that obtains between them, Terrell's commanding Mancuso to fetch some bread causes Mancuso to fetch some bread, or vice versa as the case may be. These are not distinct actions by distinct actors (like what we would expect Goldman to argue) because "actions are not… [an] abstract sort of particular that Goldman says they are but something more concrete." (Ginet 1990, 48) What matters on Ginet's account is the volition. If the core volition is the same for two or more action descriptions, then the action descriptions designate the same action. If, on the other hand, the core volitions differ for two or more action descriptions, then the action descriptions designate distinct actions. According to

Ginet, the action-designators describe an action and "pick out" the further components beyond the basic action. He writes:

> I find it natural to think of what is denoted by, for example, <<*S*'s raising her hand at t_1>> as including among its parts not only the volition denoted by [<<*S*'s willing (volition) to exert force upward with her arm and her hand at t_1>>], but also the result of that volition implied by [<<*S*'s raising her hand at t_1>>], namely *S*'s hand's rising; or, for another example, to think of what is designated by <<*S*'s offending *R* at t_3>> as including among its parts not only the action designated by <<*S*'s voting against a proposal at t_2>>, but also the result of that action implied by… *R*'s being offended by *S*'s action. I also find it natural (and here some middlers might part company with me) to think of the action designated by… <<*S*'s raising her hand just after she has heard the chair of the meeting say "And those opposed?" at t_1>> as including among its parts not only the action designated by… [<<*S*'s raising her hand at t_1>>], but also the *circumstance* of that action implied by [<<*S*'s raising her hand just after she has heard the chair of the meeting say "And those opposed?" at t_1>>]. (Ginet 1990, 49f)

This argument alone is insufficient to show that actions have layered structures, so Ginet goes on to explain the necessary and sufficient conditions for two action descriptions (in his terms canonical action-designators, see e.g., Ginet 1990, 18f) to refer to one and same concrete action.

Ginet's criterion of action individuation establishes that actions are layered structures. First, he outlines some conditions. He writes:

> Two canonical action-designators designate the same action only if the agents they refer to are the same and the times they refer to overlap in such a way that the action designated by each designator occurred during the period of overlap. Let us call two canonical action-designators related in this way *co-agential* and *co-temporal*. (Ginet 1990, 65)

Given that *X* and *Y* are co-agential and co-temporal canonical action-designators, any pair of canonical action-designators that intuitively designate the same concrete action will fall into one or another of three classes. These three are: (i) *equivalence*; (ii) *"consists in"*, and (iii) *"generation"* classes (Ginet 1990, 66f).

First, according to Ginet:

> <<S's *V*-ing at *t*>> is equivalent to <<S's *U*-ing at *t**>> if and only if it is necessarily true that S *V*-ed at *t* if and only if S *U*-ed at *t**". So, for example, <<*S*'s speaking a Finnish sentence at *t*>> is equivalent to <<*S*'s uttering a sentence of Finnish at *t*>>. (Ginet 1990, 66)

Similarly, Smith's just now rubbing her toe and at the same time closing the microwave door *is equivalent to* Smith's just now closing the microwave door and simultaneously rubbing her toe. Ginet calls the pair: *equivalent designators* (ibid.). The designators <<*S V*-ing at *t*>> and <<*S*'s *U*-ing at *t**>> are equivalent if and only if it is necessarily true that *S V*-ed at *t* if and only if *S U*-ed at *t**.

A second class of co-designating pairs designates the same action when one of the descriptions further explains an aggregate of actions and the other description refers to an action that *consisted in* that aggregate. Ginet defines the "consisted in"

relation this way: "Suarez's making a phone call at *t*" consisted in "Suarez's turning the rotary dial at *t*" if and only if Suarez's making a phone call at *t* consisted in Suarez's turning the rotary dial at *t* (ibid.). The by-relation holds for the consisted in relation. If *X* consisted in *Y*, then *X* by *Y*. Ginet writes, "If *X* CONSISTED IN *Y*, then *X* BY *Y*: If *S*'s typing *by* consisted in *S*'s first typing *b* and then typing *y*, then just because of that, *S* typed *by* by first typing *b* and then typing *y*." (ibid.) So, in the Suarez example, if Suarez's dialing 762-2239 (assume that Suarez has a rotary dial telephone) consisted in Suarez's first turning the dial from 7 and then turning the dial from 6, and so on, then, just because of that, Suarez dialed 762-2239 by first turning the dial from 7 and then turning the dial from 6, and so on.

Finally, when there are two action descriptions where one of them contains more information than the other and the information provided in the second description includes different properties than the first, Ginet notes that "both intuitively designate the same concrete action" (Ginet 1990, 66f). Suppose we have the following pair of descriptions: "Palmer's lifting his leg at *t*" and "Palmer's lifting his leg slowly at *t*." (Palmer is not a dog.) The second description gives us more information about the intrinsic features of Palmer's action, though the two seemingly designate the same concrete action. The second description of Palmer's action gives us more information about the manner of Palmer's bodily movements. Nevertheless, the two denote the same concrete action, even though one contains more or different information about the intrinsic features of the action, i.e., the manner of the action.

Ginet spends the greatest amount of time defending the third condition of his criterion for act individuation (Ginet 1990, 66ff). Judging from the amount of space and time he spends defending the third condition, he does not believe that the first or the second condition need much defense. I think the second condition is very much a problem for his view because Ginet's account may find that an aggregate action consisted in an action that actually should be distinct from the aggregate action.

The second condition states that a whole action is composed of a variety of events. When Donald turns on a light, the action he performs includes as a part the event that is the illumination of the bulb. Moreover, the movement Donald makes causes the illumination. The whole action is the aggregate of these (and other) events, and they are related in a particular kind of causal (or perhaps, more generally, instrumental) structure.

We should say that an aggregate action has a plurality of parts that are themselves actions. But does this suffice? On Ginet's nested view of action individuation, Donald's turning on the light contains as a part the act that is the flipping of the switch, and the flipping in turn contains an act that is the movement of his finger. So it turns out that the turning on of a light can be thought of as an aggregate that includes as parts (at least) three overlapping acts.

That Donald's turning on a light has three aggregate parts, namely the turning on of a light, the flipping of a switch, and the movement of his finger, is an example of one way of looking at an aggregate action, but there could be a problem for this account of aggregate action. We perform some actions where the aggregate has more

than one non-overlapping action performed by a single agent. For example, suppose that Simpson prepares dinner on Monday night, and she attends a conference two weeks later. These two acts are discrete; that is, neither is intended as a means to or a condition for the other and that together do not constitute an intended means toward or condition of some further end of the agent. So, there is no reason to say that Simpson's attending the conference consisted in Simpson's preparing dinner on Monday. On at least one reading of Ginet's account, the "consisted in" relation obtains for these two actions. For example, in Ginet's view, Simpson's attending the conference two weeks in the future *consisted in* Simpson's preparing dinner on Monday if and only if it is true that Simpson's attending the conference two weeks in the future consisted in Simpson's preparing dinner on Monday. On Ginet's view, if Simpson had not prepared dinner on Monday, then Simpson would not attend the conference two weeks later.

Ginet believes the concrete action an agent's volition picks out is the one the designator is supposed to describe. So, two action designators co-refer when and only when the agent's volition picks out an action and this action can be referred to in more than one way by two or more action-designators. In the Simpson case, the volition would be <<*S*'s willing to exert force downward with his arm and hand at t_1>>. Such a volition could be applied to both dinner preparation and to attending the conference. As long as these two action-designators could be construed as a result of the volition, that there are non-overlapping actions performed by a single action is a problem for Ginet's account.

Ginet might contend that the example includes an action that is temporally scattered. He may assert that his conception of concrete action does not share this feature. So, if aggregate action is a typical sort of concrete action, then it cannot be scattered. Suppose that when Simpson is preparing dinner, she suddenly recalls that she is late for a meeting and rushes out of the kitchen. Simpson's rushing out of the kitchen and Simpson's preparing dinner are temporally connected in a way that her preparing dinner and her going to a conference two weeks later were not. Thus, on Ginet's account, Simpson's preparing dinner and her rushing out constitute an aggregate action. It follows that specifying the condition under which various actions of an agent constitute parts of a larger action is the fundamental issue concerning aggregate action and it appears that Ginet's account suffers from serious deficiencies because it fails to accommodate these counterexamples.

Conclusion

This chapter has provided an overview of the literature on the problem of action individuation. It is not, by any means, comprehensive. The chapter's central aim has been met if the reader has a better understanding of what the problem of action individuation is. In the following chapter, I will discuss an experimental study on action individuation.

CHAPTER 5

EXPERIMENTAL RESULTS[44]

Introduction

The aim of this chapter is to report the results of an experiment on act individuation. Using these results I will discuss what future projects should address about people's intuitions. People's intuitions about act individuation are different depending on the case they are given. The data seem to point out that people may be paying attention to an agent's reasons for action when they respond to questions about action individuation.

First, I will provide some background information about the two predominant accounts of action individuation. Second, if ordinary intuitions are supposed to inform Goldman's *maximizing* theory of action individuation and if he hasn't asked people for their intuitions, then we might conclude that his account of action individuation hasn't captured folk intuitions. I provide some empirical evidence to support this claim in the second part of this paper. Finally, I will argue that Goldman's account is wrong on two counts. On the one hand, ordinary intuitions do

---

[44] I would like to thank Eric Amsel, Ron Mallon, Elijah Millgram, and Walter Sinnott-Armstrong for their helpful comments.

not coincide with his *maximizing* account. On the other hand, people's intuitions shouldn't coincide with Goldman's account since the folk aren't aware of the proper conceptual tools Goldman cites. So, if people's intuitions correspond with Goldman's account, then his analysis is unnecessary. Thus, the empirical evidence provides partial support for the argument that we abandon Goldman's account of action individuation and endorse the *minimizing* account of action individuation.

The problem of act individuation explores the identity conditions of human action. Suppose, for example, that Tiffany moves her arm, depresses the lever, lifts the weight, operates the Nautilus machine, and scares the man on the rowing machine all at the same time. Do the descriptions 'Tiffany's moving her arm' and 'Tiffany's depressing the lever' refer to two distinct actions? Or do the descriptions designate the same action? Our intuitions tell us that actions have boundaries, though the boundaries between them may not be clear.

At least two accounts have seemed intuitively plausible responses to these questions. On Donald Davidson (1963) and G.E.M. Anscombe's (1957) account, Tiffany performs one action for which there are many descriptions. On Goldman's (1970) account, each description refers to a distinct act. Numerous others have sided with Goldman or Davidson/Anscombe, or have staked out a middle view. Work on the debate has been all but abandoned because each of the positions has seemed to be intuitively plausible and equally coherent (Ginet 1990, 70).

Some Background Information

First, Donald Davidson and G.E.M. Anscombe, advocates of the *minimizing* view, argue that an agent performs one action for which there are many descriptions (Davidson 2001; Anscombe 1979). Advocates of the minimizing account suggest that an agent performs one action, which can have many descriptions. Anscombe writes, "Are we to say that the man who (intentionally) moves his arm, operates the pump, replenishes the water supply, poisons the inhabitants, is performing *four* actions? Or only one?" (Anscombe 1957/2000, 45) Anscombe has given us four descriptions that contain cognates of the action verbs "move," "operate," "replenish," and "poison." Some action must answer to each of these descriptions. So, the question is whether the same action answers to all of these descriptions or whether each description has a different action answering to it.

Anscombe argues that a variety of descriptions refer to one and the same action. She writes:

> In short, the only distinct action of his that is in question is this one, A. For moving his arm up and down with his fingers round the pump handle *is*, in these circumstances, operating the pump; and, in these circumstances, it *is* replenishing the house water-supply; and, in these circumstances, it *is* poisoning the household. (Anscombe 1957/2000, 46)

Her answer serves the purpose of having a clear understanding of whether a number of different descriptions designate the same action; it also is meant to elucidate how action descriptions may serve to specify intentions and how intentions behind an act

are related to one another. If we ask the man why he moved his arm up and down, then he may answer that he intended to operate the pump. Moreover, when we ask him why he operated the pump, he may answer that he intended to replenish the water supply.

Donald Davidson offers a similar account. His view depends on a theory of event individuation and the fact that all actions are bodily movements. First, he writes, "events are identical if and only if they have exactly the same causes and the same effects" (Davidson 2001, 179). On Davidson's account, to describe an event as a killing where the killing is caused by a shooting or by a person's pulling a trigger *is* to describe an event as causing a death (Davidson 2001, 177). Second, Davidson believes that the "welter of related descriptions corresponds to a single *descriptum*" (Davidson 2001, 59). Mere bodily movement is all that there is to an action. Accordingly, we "never do more than move our bodies: the rest is up to nature" (Davidson 2001, 59).

While Davidson and Anscombe do not explicitly address the relationship of their account to everyday practices of action individuation, their views are meant to capture how ordinary people distinguish between actions. Davidson used an example where he moves his finger, thereby flicking the switch, turning on the light, illuminating the room, and alerting the prowler outside (Davidson 1963, 686). (Anscombe used a different example in her work to explain a similar effect.) Any description of Davidson's action designates the same action because: (i) an action cannot be going on unless the agent is doing something and (ii) we don't want to say

that Davidson's moving his finger and his flicking the switch are different acts since they have the same effect. Davidson and Anscombe assume that everyone would agree that Davidson's illuminating the room is not distinct from moving his finger or flicking the switch because nothing beyond his finger movement or flicking the switch is required. Thus, although Davidson and Anscombe's theory renders a folk account of action individuation, they are less explicit about doing so than Goldman.

Second, Goldman's *maximizing* account argues that each action description designates a distinct act because each action exemplifies a unique property (Goldman 1970, 1971).[45] The maximizing account has an underlying rationale. According to the maximizing view, actions appear to involve and can be characterized as the loss, retention, acquisition, or having of properties by an object at a time. This implies that actions are exemplifications of properties.

Goldman's proposed solution to the problem of action individuation begins with a distinction between what he calls act-types and act-tokens. Act-types are a kind of property, a property of the agent (Goldman 1971, 769). For example, when we say "Ringo signaled for a turn," we ascribe act properties or act types to Ringo: the property of signaling for a turn. If an agent exemplifies two properties at the same time, then there are two exemplifyings, two such properties at the same time. Thus, there are two distinct actions.

---

[45] There is a third position in the action individuation debate. The third position, known as the *componential* view, will not be reviewed in this paper. Proponents of the *componential* view argue that an action consists in one summative action, which has many parts.

An act-token is an exemplifying of an act-type by a particular agent at a particular point in time. So, for example, Joan's moving her hand at time t_1 is an act token of the type moving one's hand. So, for Goldman:

> Since an act-token is the exemplifying of a property by an agent at a time, it is natural so to individuate act-tokens that *two act-tokens are identical if and only if they involve the same agent, the same property, and the same time*. (Goldman 1970, 10)

Suppose that there is an action x that has the property P at time t_1. According to Goldman's maximizing account, if there is some action y that has property Q at time t^*, the identity condition for the two actions is that action x and action y are identical if and only if $x = y$, $P = Q$, and $t_1 = t^*$. Otherwise, x and y are distinct act-tokens that represent distinct actions.

The minimizing and the maximizing account are the two predominant views in the action individuation literature. According to the minimizing view of action individuation, a bodily movement gives rise to a chain of events. Each event allows for a different description of action. Thus, for each bodily movement, we have a number of descriptions that satisfactorily describe the event. According to the maximizing view of action individuation, actions are particulars exemplifying a

property at a specific time. On this view, identity requires identity of the constitutive properties, physical objects, and time.[46]

There are two important characteristics of Goldman's account. The purpose of his project is (1) "to explicate certain aspects of our common sense conceptual scheme" and (2) "to develop a set of sharp conceptual tools that will be useful for studying action in a systematic way" (Goldman 1970, vi).

First, Goldman believes his explanation will be associated with a common sense conceptual scheme. He writes, "one of the purposes here is to explicate certain aspects of our common sense conceptual scheme. Thus, my analysis of action is intended to capture, as closely as possible, our pretheoretic conception of an "act" or an "action" " (ibid). Presumably, by "our," Goldman means not only philosophers but non-philosophers with little or no training in philosophy. Thus, his account attempts to capture folk intuitions.

Goldman has tried to state the implicit theory of act individuation held by the folk that produces our intuitions about particular cases. But his account seems vulnerable to the criticism that it's merely a reflection of what he – a specialist – thinks is the folk theory because he hasn't asked people for their pre-theoretic conceptions or intuitions. A better way of discovering what the folk think is to ask them directly. If we know the folk's pre-theoretic conceptions or intuitions, then we will be in a better position to summarize the folk theory of action.

---

[46] One should note that Goldman's maximizing account overlaps with Jagewon Kim's account of event individuation (Cf. Kim 1966, 1976).

Second, Goldman acknowledges that a common sense conceptual scheme may fail to exhibit the clarity a more thorough analysis may provide. Besides the delivery of an account closely associated with a commonsense conceptual scheme, Goldman says, "it is also my purpose to develop a set of sharp conceptual tools that will be useful for studying action in a systematic way" (ibid.). Systematicity cannot be achieved by ordinary intuitions alone. Therefore, ordinary intuitions must be supplemented by carefully constructed arguments supporting the introduction of novel concepts.

Ordinary pre-theoretic intuitions and conceptual analysis play a crucial role in Goldman's account of action individuation. These characteristics are components of Goldman's rejection of Davidson and Anscombe's *minimizing* account of action individuation (Goldman 1970, 1-7).

The Experiment

In developing a folk account, both Goldman and Anscombe/Davidson assume that an account of act individuation applies invariantly to different sorts of cases. Call this the assumption of invariant individuation.

Elsewhere in action theory, Joshua Knobe (2003) has been exploring intuitions with surprising results. When Knobe surveyed people's intuitions, he discovered that a majority of subjects judged that an agent intentionally brought about a side effect when the side effect was bad. An asymmetry occurs because subjects

didn't believe a side effect was brought about intentionally when that side effect was good. Knobe's methodology gives us a way of exploring intuitions about action individuation as well as examining the assumption of invariant individuation, and his findings suggest a hypothesis: action individuation will be sensitive to the valence of the consequences of the action (just as judgments of "intentionally" are).

Using the Knobe experiment as a prototype, I designed a pair of vignettes that would yield an asymmetry similar to what Knobe found in his experiment. My hypothesis is that subjects' views of the valence of the consequences of S's actions will influence whether they believe two or more descriptions designate the same or distinct acts in the manner we are considering. Consequently, we may want to say that the search for an account of action individuation that applies invariantly in different cases may well fail.

To test this hypothesis, I used a vignette to educe ordinary intuitions about action individuation. I used the Knobe experiment as a prototype for the vignette's design, and the cases are based on a classic example in Anscombe. Subjects were 75 undergraduates in introductory philosophy classes at Weber State University. Each subject was randomly assigned either to the 'poison condition' or to the 'savior condition'. Subjects in the poison condition read the following vignette:

> Smith's job is to pump water into the cistern which supplies the water of a house.
>
> One day Smith operates the pump and replenishes the house's water-supply. The occupants of the house are healthy and have no health

> problems. Jones tells Smith that someone has found a way of systematically contaminating the water's source with a deadly cumulative poison whose effects are unnoticeable until they can no longer be cured.
>
> Smith says, 'I don't care about contaminating the water's source; I just want to earn my pay.'
>
> The occupants of the house drink the water. Sure enough, they are poisoned and die.

Subjects receiving the poison condition were then asked: 'Was Smith's operating the pump the same thing as his poisoning the house's inhabitants or were they distinct?'[47]

Subjects in the savior condition received a vignette that was almost exactly the same, except that the inhabitants had a severe infection and the water had a cumulative antibiotic in it that would save them:

> Smith's job is to pump water into the cistern which supplies the water of a house.
>
> One day Smith operates the pump and replenishes the house's water-supply. The occupants of the house are sick and have severe infections. Jones tells Smith that someone has found a way of systematically purifying the water's source with a cumulative antibiotic whose effects are unnoticeable until they cure someone who has a severe infection.
>
> Smith says, 'I don't care about purifying the water's source; I just want to earn my pay.'

---

[47] Someone might wonder why I used 'thing' rather than 'act.' 'Thing' is less loaded than 'act.' Many action theorists, such as Goldman, have discussed the distinction between 'act' and 'action.' I didn't want to discuss the distinction between 'act' and 'action' or be compelled to discuss the distinction with subjects, so I decided to use 'thing' instead.

> The occupants of the house drink the water. Sure enough, they are saved and live.

Subjects receiving the savior condition were then asked: 'Was Smith's operating the pump the same thing as his saving the house's inhabitants or were they distinct?'

Each condition elicited different patterns of responses. In the poison condition, most subjects (61%) said that Smith's operating the pump and his poisoning the occupants were the same thing, whereas in the savior condition, most subjects (87%) said that Smith's operating the pump was distinct from his saving the inhabitants. The difference was highly significant, $\chi^2 (1, N = 75) = 17.613, p < .001$.[48]

Explanation

An asymmetric pattern of responses occurs in the results of the two cases, even though very little distinguishes them. It seems that in each case some feature of Smith's action affects people's intuitions. The question is: why would people respond to the two vignettes about action individuation differently?

In the poison condition, people are more prone to say that each description designates the same act. People seem to believe that Smith's operating the pump is the same thing as his poisoning the house's inhabitants. It appears that people believe the two descriptions designate the same act when a negative consequence is brought

---

[48] Joshua Knobe has recently turned to action individuation in a series of experiments exploring relevant relations such as the 'by-locution' (Knobe ms).

about by Smith's action. Although Smith is not operating the pump in order to harm the house's inhabitants, people's intuitions support the view that the two descriptions of Smith's activities designate the same action.

In the savior condition, people are far less inclined to say that each description designates the same act. When Smith's operating the pump brings about a positive consequence, people are reluctant to say that the two descriptions designate the same act. Smith is not operating the pump in order to save the house's inhabitants. The results are asymmetric because, in the savior condition, a majority of people believe each description designates distinct acts, while, in the poison condition, a majority of people believe that both descriptions refer to the same act.

What these results suggest is that the valence of the consequences of an action plays a role in how we distinguish between actions. The assumption of invariant individuation among action theorists is that a unified account of action individuation will apply to different sorts of cases. For example, Goldman has shown how his simple, unified view of action individuation applies to many different cases (Goldman 1970, 5-10; Goldman 1971, 487ff). But, as the data seem to suggest, it might be that people individuate actions in different ways in different cases. People's responses might cluster around a central core or theory, but – then again – they might not. The experimental results seemingly overturn the assumption of invariant individuation.[49]

---

[49] In fact, I believe that peoples' intuitions on individuating action may be *flexible* in the way that Nichols and Ulatowski had discovered 'intentionally' to be flexible in experiments on the Knobe effect (Nichols and Ulatowski 2007).

Of course, none of what I've said here is the last word on the folk account of action individuation. More empirical work should be completed to shed light on the folk account of action individuation. Other empirical questions may also be important for further inquiry. That there is an asymmetry present in the data goes to show only that how people distinguish between actions might depend on the valence of the consequences of an action.

CHAPTER 6

A FUTURE FOR ACT INDIVIDUATION?

Several action theorists have written off the problem of act individuation. Most recently, Andrei Buckareff has written a blog post about whether action theorists should care about the problem of act individuation (http://gfp.typepad.com/the_garden_of_forking_pat/2006/08/should_we_care_.html). In it, he expressed some views about how individuating action matters for an account of moral responsibility. But many action theorists who responded to his initial post were pessimistic. For example, Manuel Vargas was not convinced that individuating action has an impact on other problems in action theory or philosophy generally.

Buckareff is among the few action theorists who believe we should care about action individuation (or at least the few who has provided an argument supporting why we should care about action individuation), whereas Vargas is the norm. Action theorists have claimed that accounts of individuating action are uninformative. I already have pointed out Ginet's concern over the importance of action individuation (Ginet 1990, 70; above). Since action theorists believe that the accounts are uninformative and serve no purpose in any of the other (more prevalent) debates in

action theory, action theorists have argued that the exploration of the problem of act individuation should be abandoned. Investigators explore the problem of act individuation for no other reason than they have an interest in it.

The first chapter attempted to motivate an exploration of the problem of act individuation. Recent developments in philosophy provoked me to devise an experiment addressing the problem of act individuation. Chapter five showed how an empirical investigation of people's intuitions shows that invariantist accounts of act individuation have failed. The experimental results support the claim that invariant accounts of act individuation have failed to consider people's intuitions about individuating action. If these data are accurate, then the future of the problem of act individuation lies in experimental work. Thus, we can suppose that this project has revealed a new way of attending to the problem of act individuation.

Experimental work will serve action theory and the problem of act individuation well. If action theorists accept that we can re-open the case about the importance of the problem of act individuation and if experimenting on people's intuitions is one way that we can address the problem of act individuation, then there is a future for research in the problem space.

In this section, I would like to discuss two future projects. The first project is an empirical investigation of people's intuitions about temporal problems and individuating actions. Advocates of the componential account used the problem of temporal order relations to show that the minimizing and maximizing views were untenable. For this project, I hypothesize that disagreement over the time of an action

depends on people's views of the consequences of the action under consideration. If my hypothesis is correct, then the componentialists' rejection of the minimizing and maximizing view can be called into question. The second project is an empirical investigation that aims to show people's intuitions about act individuation are flexible. By furthering the study I performed for this dissertation project, I believe that an explanation of the asymmetry is due in part to the stable individual differences in how people interpret an action.

Persisting Acts

The componentialists' dissatisfaction with the minimizing account and the maximizing view arises largely from a disagreement over the time of an action. Suppose that Smith shoots Jones and Jones dies six months later from the gunshot wound she sustained. First, on the minimizing account, since the two descriptions "Smith's shooting Jones" (call it "*A*") and "Smith's killing Jones" (call it "*B*") designate the same act because Jones's death was hastened by Smith's shooting him, Smith killed Jones several months before Jones died. "Killing" entails that a person has died at the hands of another. So, if *B* applies to the same act as *A*, then Jones has been killed long before he is dead, which is inconsistent with our understanding of "killing."

Second, on the maximizing view, *A* and *B* are distinct because they exemplify different properties. The maximizing view seems more tenable than the minimizing

view until we find out that Smith died three months after shooting Jones. If Smith dies before Jones dies, then Smith acted after he is already dead. So, it seems either we accept the maximizing view of a *zombified* Smith who continues to act after he is dead or we agree that Smith can kill someone without the person he killed is dead.

In this project I subjected the componentialists objection to experimentation. Preliminary data indicate that even if people are prompted to notice the temporal distinction, people tend to judge the identity of an action based upon the valence of the consequences.

The experiment used the following two vignettes. For the offending condition (similar to the poison condition above), I asked:

> Smith is a member of the City Council. Part of his obligation is to vote on laws that affect millions of people. Smith's vote typically determines whether or not a proposal passes into law. Citizens know of Smith's crucial role, and they urge him to vote against a proposal permitting the sale of handguns to minors.
>
> On Tuesday night, the proposal is presented to a closed session of the City Council. Since the session is closed, the public will not know of Smith's vote for a week. City Council members are asked to vote on the proposal. Smith raises his hand and votes for the measure. As a result, the proposal passes into law.
>
> A week later citizens learn of Smith's vote and become offended.

Subjects were asked how much they agreed or disagreed with the following statement: *Even though citizens learn of Smith's vote a week after it has been cast,*

Smith offends them when he votes for the proposal. Respondents believed that Smith offended them when he cast his vote (mean = 5.0).

The excitement condition was similar to the offending condition, but it replaced Smith's voting against the measure with his voting for the measure, thereby passing the proposal the citizens wanted passed.

> Smith is a member of the City Council. Part of his obligation is to vote on laws that affect millions of people. Smith's vote typically determines whether or not a proposal passes into law. Citizens know of Smith's crucial role, and they urge him to vote against a proposal permitting the sale of handguns to minors.
>
> On Tuesday night, the proposal is presented to a closed session of the City Council. Since the session is closed, the public will not know of Smith's vote for a week. City Council members are asked to vote on the proposal. Smith raises his hand and votes against the measure. As a result, the proposal does not pass into law.
>
> A week later citizens learn of Smith's vote and become excited.

Subjects were asked how much they agreed or disagreed with the following statement: *Even though citizens learn of Smith's vote a week after it has been cast, Smith excites them when he votes for the proposal.* Subjects' responses suggest that citizens did not become excited when he voted for the proposal (mean = 2.75).

The preliminary data suggest that people's views on the timing of an action (and whether an act can persist through time) has much more to do with the

consequences of an act than the componentialist has supposed. Further analysis is required for my hypothesis to be empirically confirmed.

Individual Differences and Act Individuation

In a paper I co-authored with Shaun Nichols (2007), we argued that interpretive diversity is present in how people interpret the term "intentional." The term "intentional" admits of different interpretations, and such diversity can be seen in the minority responses of Knobe experiment cases. According to an interpretation of the data we collected, the minority responses are adopting a consistent (but systematically different) pattern of responses for the term "intentional."

The problem of "intentional" action presumes that we know something about action and that we can distinguish one action from another. The present project has attempted to show that an asymmetry arises in subjects' responses to vignettes about "action" in a similar way an asymmetry arises in experiments about "intentional" action. To say that the two projects are the same would be wrong.[50] One area in which our work might overlap is derived from the interpretive diversity thesis me and Nichols (2007) suggested in response to the Knobe effect.

Further examination of people's intuitions is required. If I were to find that the minority responses adopt a consistent yet systematically different pattern of

---

[50] In fact, several bloggers have commented that *my* experiment is just Knobe's experiment. They claim there is no difference between my experiment and Knobe's. But that is to miss the point of my work. My experiment concerns "action," not "intentional" action.

responses for the individuation of action, then interpretive diversity applies not only to the term “intentional” but also to action, as well.

These are just two ways in which my project on act individuation may be extended. Given the results of the experiment on action individuation I found leading up to the writing of this dissertation, I believe much more work can be (and should be) done in this area of action theory which many people had written off more than a decade ago.

REFERENCES

Adams, Frederick and Annie Steadman (2004a). "Intentional Action in Ordinary Language: Core Concept or Pragmatic Understanding?" *Analysis* 64: 173-181.

Adams, Frederick and Annie Steadman (2004b). "Intentional Action and Moral Considerations: Still Pragmatic." *Analysis* 64: 264-267.

Adams, Frederick and Annie Steadman (2007). "Folk Concepts, Surveys, and Intentional Action." in Christoph Lumer and Sandro Nannini (eds.) *Intentionality, Deliberation, and Autonomy: The Action-Theoretic Basis of Practical Philosophy*. Aldershot: Ashgate Publishers: 17-35.

Alexander, Joshua and Jonathan Weinberg (2007). "Analytic Epistemology and Experimental Philosophy." *Philosophy Compass* 2.1: 56-80.

Andreou, Chrisoula (2006). "Temptation and Deliberation." *Philosophical Studies* 131: 583-606.

Annas, Julia (1976). "Davidson and Anscombe on 'the Same Action'." *Mind* 75: 251-257.

Anscombe, G.E.M. (1970). "War and Murder" in Richard Wasserstrom (ed.). *War and Morality*. Belmont, CA: Wadsworth: 42-53.

Anscombe, G.E.M. (1979). "Under a Description." *Nous* 13: 219-233.

Anscombe, G.E.M. (1957/2000). *Intention*. Second edition. Cambridge, MA: Harvard University Press.

Appiah, Kwame Anthony (2008). *Experiments in Ethics*. Cambridge, MA: Harvard University Press.

Audi, Robert (1983). "The Applications of Conceptual Analysis." *Metaphilosophy* 14,2: 87-106.

Bach, Kent (1980). "Actions are not Events" *Mind* (89): 114-120.

Bennett, Jonathan (1966). "Whatever the Consequences." *Analysis* 26.3: 83-102.

Bennett, Jonathan (1988). *Events and Their Names*. Indianapolis, IN: Hackett Publishing Company.

Bennett, Jonathan (1995). *The Act Itself*. New York: Oxford University Press.

Bishop, Michael A. (1992). "The Possibility of Conceptual Clarity in Philosophy." *American Philosophical Quarterly* 29.3: 267-277.

Bishop, Michael and Dominic Murphy, eds. (forthcoming). *Stich and His Critics*. London: Blackwell Publishers.

Bratman, Michael (1983). "Taking Plans Seriously." *Social Theory and Practice* 9: 271-288.

Bratman, Michael (1999a). *Faces of Intention*. Cambridge, UK: Cambridge University Press.

Bratman, Michael (1999b). *Intention, Plans, and Practical Reason*. Palo Alto, CA:

CSLI Publications.

Broad, C.D. (1927). *Scientific Thought*. Patterson, N.J.: Littlefield Adams.

Buckareff, Andrei (2007). "Should we care about action individuation?" http://gfp.typepad.com/the_garden_of_forking_pat/2006/08/should_we_care_.html.

Carnap, Rudolf (1950). *Logical Foundations of Probability*. London: Routledge and Kegan Paul.

Carruthers, Peter, Stephen Laurence, and Stephen Stich (eds.) (2007a). *The Innate Mind: Volume 1: Structure and Contents*. Oxford: Oxford University Press.

Carruthers, Peter, Stephen Laurence, and Stephen Stich (eds.) (2007b). *The Innate Mind: Volume 2: Culture and Cognition*. Oxford: Oxford University Press.

Carruthers, Peter, Stephen Laurence, and Stephen Stich (eds.) (2007c). *The Innate Mind: Volume 3: Foundations and the Future*. Oxford: Oxford University Press.

Castaneda, Hector Neri (1979). "Intensionality and Identity in Human Action and Philosophical Method." *Nous* (13): 235-260.

Collingwood, Robin George (1938). *The Principles of Art*. Oxford, UK: Clarendon Press.

Cooper, John M. (1999). *Reason and Emotion: Essays on Ancient Moral Psychology and Ethical Theory*. Princeton, N.J.: Princeton University Press.

D'Arcy, Eric (1963). *Human Acts*. Oxford: Oxford University Press.

Davidson, Donald (1963). "Action, Reasons, and Causes." Reprinted in Davidson 2000: 3-19.

Davidson, Donald (1967). "The Logical Form of Action Sentences" in Davidson 2001: 105-148.

Davidson, Donald (1969). "The Individuation of Events" in Davidson 2001: 163-180.

Davidson, Donald (1971). "Agency" in Davidson 2001: 43-61.

Davidson, Donald (2001). *Essays on Actions and Events*. Oxford: Clarendon Press.

Davis, Lawrence (1970). "Individuation of Actions." *Journal of Philosophy* 67: 520-530.

Davis, Lawrence (1979). *Theory of Action*. Englewood Cliffs, N.J.: Prentice Hall.

DePaul, Michael and William Ramsey eds. (1998). *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*. Lanham, MD: Rowman and Littlefield.

Descartes, Rene (1641/1984). *The Philosophical Writings of Descartes: Volume II*. Translated by John Cottingham, Robert Stoothoff, and Dugald Murdoch. Cambridge: Cambridge University Press.

Evans, Jonathan St. B.T. (2002). "Logic and Human Reasoning: An Assessment of the Deduction Paradigm." *Psychological Bulletin* 128: 978-996.

Evans, Jonathan St. B.T. and David E. Over (1996). *Rationality and Reasoning*. East Sussex, UK: Psychology Press.

Fehige, Christoph (2001). "Instrumentalism" in Millgram 2001: 49-76.

Feinberg, Joel (1965). "Action and Responsibility." Reprinted in White 1968: 95-119.

Fodor, Jerry (1998). *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.

Fodor, Jerry (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

Gaut, Berys (2002). "Justifying Moral Pluralism" in Stratton-Lake 2002: 137-160.

Gibbard, Allan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.

Ginet, Carl (1991). *On Action*. Cambridge: Cambridge University Press.

Goldman, Alvin (1970). *A Theory of Human Action*. Englewood Cliffs, N.J.: Prentice Hall.

Goldman, Alvin (1971). "The Individuation of Action." *Journal of Philosophy* 68: 761-774.

Goodman, Nelson (1983). *Fact, Fiction, and Forecast.* Fourth edition. Cambridge, MA: Harvard University Press.

Greenough, Patrick and Michael P. Lynch, eds. (2006). *Truth and Realism: Current Debates*. Oxford: Oxford University Press.

Hanna, Robert (2000). "Conceptual Analysis" *Concise Routledge Encyclopedia of Philosophy*. London: Routledge: 106.

Hart, H.L.A. (1948/49). "The Ascription of Responsibility and Rights." *Proceedings of the Aristotelian Society* 49: 145-166.

Hart, H.L.A. (1973). *Punishment and Responsibility*. Oxford: Clarendon Press.

Herman, Barbara (1993). *The Practice of Moral Judgment*. Cambridge, MA: Harvard University Press.

Hornsby, Jennifer (1980). *Actions*. London: Routledge and Kegan Paul.

Jackson, Frank (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.

Kahneman, Daniel, Paul Slovic, and Amos Tversky eds. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

Kant, Immanuel (1965). *Grundlegung zur Metaphysik der Sitten* (*Groundwork for the Metaphysics of Morals*, 1785). Translated by H.J. Paton. New York: Harper & Row.

Kant, Immanuel (1991). *Metaphysik der Sitten* (*The Metaphysics of Morals*, 1797). Translated by Mary Gregor. Cambridge: Cambridge University Press.

Kauppinen, Antii (2007). "The Rise and Fall of Experimental Philosophy." *Philosophical Explorations* 10.2: 95-118.

Kavka, Gregory (1983). "The Toxin Puzzle." *Analysis* 43: 33-36.

Kim, Jaegwon. (1966). "On the psycho-physical identity theory." *American Philosophical Quarterly* 3: 231-232.

Kim, Jaegwon. (1976). "Events as property exemplifications," 159-177 *Action theory*. M. Brand and D. Walton (eds.). Dordrecht: D. Reidel Publishing Company.

Kokis, Judite, Robyn Macpherson, Maggie Toplak, Richard F. West, and Keith Stanovich (2002). "Heuristic and Analytic Processing: Age Trends and Associations with Cognitive Ability and Cognitive Styles." *Journal of Experimental Child Psychology* 83: 26-52.

Knobe, Joshua (2003a). "Intentional Action and Side Effects in Ordinary Language." *Analysis* 63: 190-193.

Knobe, Joshua (2003b). "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16: 309-324.

Knobe, Joshua (2004). "Intention, Intentional Action, and Moral Considerations." *Analysis* 64: 181-187.

Knobe, Joshua (ms). "Anscombe 50 years later."

Knobe, Joshua (ed.) (2008). *Experimental Philosophy*. New York: Oxford University Press.

Kuhn, Thomas (1962). *The Structure of Scientific Revolutions*. Chicago: The University of Chicago Press.

Lakatos, Imre (1970). "Falsification and the Methodology of Scientific Research Programmes" in Imre Lakatos and Alan Musgrave (eds.) *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press, 91-195.

Lepore, Ernest (1985). "XXX" in Ernest Lepore and Brian McLaughlin (eds.) *The Philosophy of Donald Davidson: Perspectives on Actions and Events*. London: Blackwell, XXX-XXX.

Lewis, David (1972). "Psychophysical and Theoretical Identifications." *Australasian Journal of Philosophy* 50: 249-258.

Libet, Benjamin (2004). *Mind Time: The Temporal Factor in Consciousness*. Cambridge: Harvard University Press.

Loux, M.X. and Zimmerman, D.W. eds. (2003). *The Oxford Handbook of Metaphysics*. Oxford: Oxford University Press.

Lynch, Michael (2004). *True to Life*. Cambridge, MA: MIT Press.

Machery, Edouard, Ron Mallon, Shaun Nichols, and Stephen Stich (2004). "Semantics, Cross-Cultural Style" *Cognition* 92: B1-B12.

Malle, Bertram (2006). "Intentionality, Morality, and Their Relationship in Human Judgment." *Journal of Cognition and Culture* 6: 61-86.

McCann, Hugh (1998). *The Works of Agency: On Human Action, Will, and Freedom*. Ithaca, N.Y.: Cornell University Press.

Mele, Alfred (1992). *Springs of Action: Understanding Intentional Behavior*. Oxford: Oxford University Press.

Mele, Alfred (1997). *The Philosophy of Action*. Oxford: Oxford University Press.

Mele, Alfred (2003). *Motivation and Agency*. Oxford: Oxford University Press.

Mele, Alfred (ms). "Practical Mistakes and Intentional Action." 1-26.

Millgram, Elijah ed. (2001). *Varieties of Practical Reasoning*. Cambridge, MA: MIT Press.

Millgram, Elijah (2005). *Ethics Done Right: Practical Reasoning as a Foundation for Moral Reasoning*. Cambridge: Cambridge University Press.

Nadelhoffer, Thomas (2004a). "On Praise, Side Effects, and Folk Ascriptions of Intentionality." *Journal of Theoretical and Philosophical Psychology* 24: 196-213.

Nadelhoffer, Thomas (2004b). "Blame, Badness, and Intentional Action: A Reply to Knobe and Mendlow." *Journal of Theoretical and Philosophical Psychology* 24: 259-269.

Nadelhoffer, Thomas (2006). "On Trying to Save the Simple View." *Mind and Language* 21.5: 565-586.

Nadelhoffer, Thomas and Eddy Nahmias (2007). "The Past and Future of Experimental Philosophy." *Philosophical Explorations* 10.2: 123-149.

Nell, Onora (1975). *Acting on Principle: An Essay on Kantian Ethics*. New York: Columbia University Press.

Newstead, Stephen, Simon Handley, Clare Harley, Helen Wright, and Daniel Farrelly (2004). "Individual Differences in Deductive Reasoning." *Quarterly Journal of Experimental Psychology* 57A: 33-60.

Nichols, Shaun and Joseph Ulatowski (2007). "Intuitions and Individual Differences: The Knobe Effect Revisited." *Mind and Language* 22.4: 346-365.

Penner, Terry (1973). "The Unity of Virtue." *Philosophical Review* 82.1: 35-68.

Pettit, Philip (1999). "A Theory of Normal and Ideal Conditions," *Philosophical Studies* 96: 21-44.

Pust, Joel (2000). *Intuitions as Evidence*. New York: Garland Press.

Quine, W.V.O. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Quinn, Warren (1989). "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing" *Philosophical Review* 98: 287-312.

Quinn, Warren (1993). "The Puzzle of the Self-Torturer" in *Morality and Action*. Cambridge, UK: Cambridge University Press.

Shapere, Dudley (1966). "Meaning and Scientific Change." In R.G. Colodny (ed.), *Mind and Cosmos: Essays in Contemporary Science and Philosophy*. Pittsburgh: University of Pittsburgh Series in the Philosophy of Science, vol. 3, 41-85.

Stratton-Lake, Philip (2002). *Ethical Intuitionism: Re-evaluations*. Oxford: Oxford University Press.

Rawls, John (1999). *A Theory of Justice*. Revised edition. Cambridge, MA: Harvard University Press.

Rawls, John (2000). *Lectures on the History of Moral Philosophy*. Edited by Barbara Herman. Cambridge, MA: Harvard University Press.

Rosch, Eleanor (1978). "Principles of Categorization" in *Cognition and Categorization*. Edited by Eleanor Rosch and B. Lloyd. New York: Lawrence Erlbaum Associates.

Shwayder, D.S. (1965). *The Stratification of Behavior*. London: Routledge and Kegan Paul.

Sinnott-Armstrong, Walter (ed.) (2008a). Moral Psychology: The Evolution of

Morality: Adaptations and Innateness. Cambridge, MA: MIT Press.

Sinnott-Armstrong, Walter (ed.) (2008b). Moral Psychology: The Cognitive Science

of Morality: Intuition and Diversity. Cambridge, MA: MIT Press.

Sinnott-Armstrong, Walter (ed.) (2008c). *Moral Psychology: The Neuroscience of Morality: Emotion, Brain Disorders, and Development*. Cambridge, MA: MIT Press.

Slote, Michael (1989). "Moderation and Satisficing" reprinted in Millgram 2001: 221-235.

Sosa, Ernest (2006). "Intuitions and Truth" in Greenough and Lynch 2006: xx-xxx.

Sosa, Ernest (2007). "Experimental Philosophy and Philosophical Intuition" *Philosophical Studies* 132: 99-107.

Sosa, Ernest (forthcoming). "A Defense of the Use of Intuitions in Philosophy" in Bishop and Murphy forthcoming.

Stich, Stephen (1990). *Fragmentation of Reason*. Cambridge, MA: MIT Press.

Stich, Stephen (1996). *Deconstructing the Mind*. New York: Oxford University Press.

Stich, Stephen and Jonathan Weinberg (2001). "Jackson's Empirical Assumptions" *Philosophy and Phenomenological Research* 62: 637-643.

Stanovich, Keith and Richard F. West (1998). "Individual Differences in Rational Thought." *Journal of Experimental Psychology: General* 127: 161-188.

Stanovich, Keith and Richard F. West (2000). "Individual Differences in Reasoning: Implications for the Rationality Debate." *Behavioral and Brain Sciences* 23: 645-726.

Steinbock, Bonnie and Alastair Norcross eds. (1994). *Killing and Letting Die*. New York: Fordham University Press.

Strawson, Peter F. (1997). *Entity and Identity: and Other Essays*. Oxford: Clarendon Press.

Thalberg, Irving (1971). "Singling Out Actions, Their Properties and Components," *Journal of Philosophy* 68: 781-787.

Thalberg, Irving (1977). *Perception, Emotion, and Action.* London: Basil Blackwell.

Timmons, Mark (1997). "Decision Procedures, Moral Criteria, and the Problem of Relevant Descriptions in Kant's Ethics." *Jarbuch für Recht und Ethik* (5): 389-417.

Timmons, Mark ed. (2002a). *Kant's Metaphysics of Morals: Interpretive Essays*. Oxford: Oxford University Press.

Timmons, Mark (2002b). "Motive and Rightness in Kant's Ethical System" in Timmons 2002a: 255-288.

Thomson, Judith Jarvis (1971a). "The Time of a Killing" *Journal of Philosophy* 68(5): 115-132.

Thomson, Judith Jarvis (1971b). "Individuating Actions." *Journal of Philosophy* 68: 774-781.

Thomson, Judith Jarvis (1977). *Acts and Other Events*. Ithaca, N.Y.: Cornell University Press.

Torrens, Donna, Valerie Thompson, and Kenneth Cramer (1999). "Individual Differences and the Belief Bias Effect: Mental Models, Logical Necessity, and Abstract Reasoning." *Thinking and Reasoning* 5: 1-28.

Vargas, Manuel (2006). "Philosophy and the Folk: On Some Implications of Experimental Work for Philosophical Debates on Free Will" *Journal of Cognition and Culture* 6(1-2): 239-254.

Weil, Vivian M. and Irving Thalberg (1974). "The Elements of Basic Action." *Philosophia* (4): 111-138.

Weinberg, Jonathan (2007). "How to Challenge Intuitions Empirically without Risking Skepticism." *Midwest Studies in Philosophy* 31: 318-343.

Weinberg, Jonathan, Shaun Nichols, and Stephen Stich (2001). "Normativity and Epistemic Intuitions," *Philosophical Topics* 29: 429-460.

White, Alan ed. (1968). *The Philosophy of Action*. Oxford: Oxford University Press.

Wilkes, Kathleen (1988). *Real People*. New York: Oxford University Press.

Wittgenstein, Ludwig (1953). *Philosophical Investigations.* 3rd edition, translated by G.E.M. Anscombe. Upper Saddle River, N.J.: Prentice-Hall.

Woodward, P.A. ed. (2001). *The Doctrine of Double Effect: Philosophers Debate a Controversial Moral Principle*. South Bend, IN: University of Notre Dame Press.

Von Wright, Georg Henrik (1963). *Norm and Action*. London: Routledge and Kegan Paul.

Von Wright, Georg Henrik (1971). *Explanation and Understanding*. London: Routledge and Kegan Paul.