

Actions Are Not Events:
An Ontological Objection to the Causal Theory of Action

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Jiajun Hu
August 2018

Copyright © 2018 by Jiajun Hu.

All rights reserved.

DEDICATION

For my parents
Derong Hu and Juxian Xia
my wife Shen Qian

ACKNOWLEDGEMENTS

Firstly, I would like to express my sincere gratitude to my advisor Prof. David Palmer for the continuous support of my Ph.D study, for his patience, motivation, and immense knowledge. His guidance helped me in all the time of research and writing of this dissertation.

Besides my advisor, I would like to thank the rest of my dissertation committee: Prof. Richard Aquila, Prof. E. J. Coffman, and Prof. Bruce MacLennan, for their insightful comments and encouragement.

ABSTRACT

One of the central puzzles in the philosophy of action is how to provide a coherent account of agency within a wholly natural worldview. More specifically, the challenge is to explain how a person's performing actions – the essential means by which she gets things done – could be a part of the natural order. In contemporary philosophy, a prominent and perhaps the most influential answer to this challenge is the so-called “causal theory of action” (henceforth, the CTA). Proponents of the CTA believe that what makes it the case that behavioral events are actions is that they are caused by the person's prior mental events or states (e.g., desire-belief pairs, or intentions), and, in addition, that the prior mental causes constitute her reasons for the action in question. The CTA is an *event*-causal theory of action that attempts to understand human agency in terms of event-causality and treat human actions as particular events in the network of event causation.

In this dissertation, I raised an ontological objection to the CTA that actions are not events. Since actions are not events, any attempt to account for the notion of action by appealing to how actions enter into the event-causal relations necessarily fails. The ontological objection consists in three separate arguments sharing the same spirit, each of which gives the same conclusion: actions are not events. Near the end of the dissertation, I anticipate and reject a Davidsonian response that appears to neutralize the force of my ontological objection. The response is that the difference between followers of the CTA and me is just another manifestation of the perennial dispute between nominalists and realists. I argue it is not. I close the dissertation by briefly explaining what a non-causal account of agency might look like.

TABLE OF CONTENTS

Chapter 1: Introduction	1
Some Objections to the CTA	2
The Ontological Objection to the CTA.....	4
Kim versus Davidson.....	9
Act-Universals	13
The Logical Form of Action Sentences	20
A Summary of Chapters	22
Chapter 2: The Intention-Based Argument.....	23
Two Immediate Challenges to SV	24
Bratman’s Videogame Case.....	28
Intentions and Events.....	39
An Objection to Consider	49
Chapter 3: The Ability-Based Argument.....	68
Ability to Act Intentionally.....	68
Minimal Ability	75
A Revised Account of Minimal Ability.....	79
A Timing Argument.....	85
Chapter 4: The Responsibility-Based Argument	95
Responsibility and Moral Quality.....	95
Moral Quality and Events	108
An Objection: Two Faces of Moral Responsibility	120
Chapter 5: A Davidsonian Challenge	126
Chapter 6: Conclusion.....	137
References	144
Vita.....	150

CHAPTER 1: INTRODUCTION

One of the central puzzles in the philosophy of action is how to provide a coherent account of agency within a wholly natural worldview. More specifically, the challenge is to explain how a person's performing actions – the essential means by which she gets things done – could be a part of the natural order. In contemporary philosophy, a prominent and perhaps the most influential answer to this puzzle of agency is the so-called “causal theory of action” (henceforth, the CTA). The CTA is first proposed and defended by Donald Davidson (1963). It is now the dominant view in philosophy of action, to such an extent that J. David Velleman (1992) calls it the “standard story” of action. While the CTA has various versions, philosophers who favor it all share a basic conviction about what makes something an action and what explains an action. That is, they believe that what makes it the case that behavioral events are actions is that they are caused by certain of the person's prior mental events or states (e.g., desire-belief pairs, or intentions), and, in addition, that the prior mental causes constitute her reasons for the action in question. The CTA is an *event*-causal theory of action that attempts to understand human agency in terms of event-causality and treat human actions as particular events in the network of event causation.

An important implication, and also a merit in my opinion, of the CTA is that it promises that human agency can fit nicely into the ontology of scientific naturalism. This is because events and event causality are usually taken to be respected “building blocks” of such a naturalistic ontology. If human agency consists in event causality and, let us assume with scientific naturalists, mental events are physical occurrences (which concerns a more fundamental question of philosophy of mind), then there is no obvious obstacle to locating and accommodating the common phenomenon of human actions in the natural causal network. The success of the CTA,

therefore, can be seen as a basis for reconciling the scientific perspective of the world with our idea of ourselves as agents who produce changes in the world. In fact, it might even be the first step towards answering the genuine philosophical question of “how morally responsible agents can be part of the natural order as our natural scientific worldview conceives of it” (Bishop 1989, pp. 15).

1. Some Objections to the CTA

The success of the CTA remains a controversial and debated topic. Since Davidson’s first presentation and defense of it, the CTA is not short of critics. A well-known objection raised against it is the so-called “problem of causal deviance.” Davidson anticipates such a problem in his famous case of the nervous climber:

A climber might want to rid himself of the weight and danger of holding another man on a rope, and he might know that by loosening his hold on the rope he could rid himself of the weight and danger. This belief and want might so unnerve him as to cause him to loosen his hold (Davidson 2001, pp. 79).

The case apparently undermines the CTA because it seems to illustrate a situation in which an event satisfies the causal conditions for an action according to the CTA but fails to be an action at all. Most philosophers in the debate agree that the problem of causal deviance poses a serious philosophical challenge to the CTA. To save the CTA from such a problem, an immediate response might be to qualify the CTA by requiring that the person’s prior mental events or states cause the behavioral events *in the right and non-deviant way*. However, the question is how to spell out what the right and non-deviant causation is, without appealing to concepts that cannot be reduced to event-causal terms (e.g., an unanalyzable notion of agent causation).

Various answers have been given to this particular question. First of all, Davidson (2001, essay 12) is pessimistic about the hope of discovering an event-causal account of non-deviant

causation. In the face of the difficulty, Davidson adjusts his goal and aims for a limited analysis of the notion of action or agency instead of a complete one. Such a limited causal analysis would leave the phrases such as “in the right way” or “in a non-deviant way” unspecified. As he writes:

Two ideas are built into the concept of acting on a reason ... the idea of cause, and the idea of rationality. A reason is a rational cause. One way rationality is built in is transparent: the cause must be a belief and a desire in the light of which the action is reasonable. But rationality also enters more subtly, since the way desire and belief work to cause the action must meet *further, and unspecified, conditions*. (Davidson 2001, pp. 233, my emphasis)

It is unclear if Davidson’s pessimism about a complete analysis based on event-causal terms is compatible with his commitment to the general spirit of the CTA. Some philosophers (e.g., Bishop 1983) would answer no. Alvin Goldman (1970), who shares Davidson’s pessimism to a certain degree, believes that giving an account of non-deviant causation is a matter of empirical discovery instead of a task that can be achieved by armchair philosophers.

However, more proponents of the CTA remain hopeful and have tried to propose various specifications of non-deviant causation (e.g., Peacocke 1979, Bishop 1989, Mele 2003b). Among those suggestions, an influential proposal is that behavioral events that are actions must be *responsive* to the contents of the relevant prior mental states or events (besides satisfying the condition that has already been explicated by the CTA). This is called the “sensitivity condition” which has been used to spell out what non-deviant causation consists in.

An alternative challenge to the CTA is sometimes called the “absent agent” or the “disappearing agent” problem. In the introductory chapter of *Causing Human Actions*, Jesús H. Aguilar and Andrei A. Buckareff summarize the problem in this way:

If, as the CTA proposes, a behavior is actional in virtue of how it is caused by some mental state or event, and no mental state or event seems to be identical with an agent, then it is hard to see how such behavior involves an agent who is in charge of producing and controlling it. In fact, it appears that the agent has been erased from the picture that the CTA offers about the production and control of an action, effectively making the agent an absentee in such process. This is the problem of the absent agent (2010, pp. 12).

Although suggestive, it is not immediately clear, without further clarification, what the problem of the absent agent is and how exactly it undermines the CTA. To some, the problem is not that the agent is literally missing in the causal picture of human actions provided by the CTA. Rather, the CTA should be accused of making the agent a mere passive participant in the production of her actions. On this interpretation, the agent might still exist, but she does not play any role in producing and explaining actions. In other words, the agent becomes a bystander or even a “victim” with respect to the events that are her actions or their production. Other philosophers understand the problem more radically. According to them, the problem of the CTA is not just rendering the agent as a mere passive bystander but making the agent *completely* disappear in the picture. As Jennifer Hornsby puts it, the CTA “in any of its versions ... is not a story of agency at all” (Hornsby 2004, pp. 2).

Various formulations of the problem of the absent agent have been raised in the literature (Melden 1961, Nagel 1986, Velleman 1992, Hornsby 2004, Steward 2012). In response, defenders of the CTA have tried to answer the problem of the absent agent by appeal to additional accounts of what constitutes an agent. For example, it may be argued that an agent is just someone who possesses those mental states or events (such as desires, beliefs, or intentions) that can cause and explain actions. Or, a more sophisticated view can be that a person is an agent of her action if and only if she reflectively endorses the mental causes of the action.

2. The Ontological Objection to the CTA

Although the problem of the absent agent has not convincingly shown that the CTA is false, it nevertheless raises this interesting question: What aspect or feature of the CTA allegedly makes the agent disappear (according to critics of the CTA)? In other words, if the problem of the absent agent exists, what is the source of this problem? Some philosophers believe that the

source of the problem is the largely unquestioned orthodoxy that actions are events. To them, any account of actions that treats actions as a special kind of event will necessarily make the agent disappear.

In this dissertation, I agree with those philosophers who object that actions are not events. My aim is to build an objection to the CTA by arguing that actions are not events. This particular objection to the CTA is based on a more general and fundamental point about the ontology of actions. That is, it involves a reflection on the question: What are actions? Or, what kind of things are actions? By denying that actions are events, we need to search for an alternative answer to this fundamental question, different from the one that is assumed or even taken for granted by the CTA. In some sense, the CTA can be understood as a causality-based answer to the following question: Since actions are events, how should we distinguish between actions and non-actional events (e.g., mere bodily movements)? Now, the objection I am going to raise in this dissertation challenges this assumption of the CTA. If, contrary to what proponents of the CTA assume, actions do not constitute a subclass of events, there is no hope for the CTA to be true since only events can enter into the event-causal relations as effects.

Due to the fact that my objection involves questioning the ontology of action, I shall call it the “ontological objection” to the CTA. If my objection turns out to be a solid one, it not only poses a challenge to any particular version of the CTA. In fact, it serves as an objection to the CTA *per se*. So, unlike the problem of causal deviance, if the ontological objection is sound, there is no remaining doubt about whether the CTA should be rejected or revised. The whole idea of understanding actions and human agency in terms of event-causality must be abandoned since actions cannot be *event*-caused.

The main argument starts with a basic idea that actions belong to the sort of things that we can intentionally do, that we can have abilities to do, and that we can be morally responsible for. This idea is about the ontological status of actions, namely, about what kind of things actions are. It contains three different and independent requirements:

- (1) We can intentionally perform actions.
- (2) We can have abilities to perform actions.¹
- (3) We can be morally responsible for performing actions.

The “can” in (1)-(3) refers to a very weak modal requirement. For example, as far as requirement (1) is concerned, it does not mean that it is metaphysically or physically possible for us to perform actions intentionally. It merely means that actions’ being the kind of things they are does not make it impossible for us to perform actions intentionally. So, the truth of (1) is compatible with the situation that we can never do actions intentionally due to some other reasons (e.g., that we and all living creatures are psychologically unable to intend to act and an intention as such is a necessary condition of intentional action). Therefore, (1) does not entail that there are at least some actions that are intentionally done in our real world although it is a highly plausible view. However, the negation of (1) would imply that no intentional action exists in our world (or any world), which I believe is a result unacceptable to almost everyone. Likewise, the negation of (2) or (3) also has a highly implausible implication: either that we have no ability to perform any action or that we are not morally responsible for performing any action.

¹ By “can have abilities to perform actions,” I naturally mean that it is possible for us to have abilities to do actions. Likewise, when I argue in this dissertation that events are not things that we can have abilities to bring about, I mean that it is impossible for us to possess or acquire abilities to bring about events given that they are events and we are human agents.

The plausibility of (1)-(3) can go beyond what the unacceptable implications of their negations show. A natural reading of (1)-(3) would take them to express some metaphysical facts. However, upon reflection, it seems that (1)-(3) express logical facts. If (1)-(3) are false, it raises an immediate worry what use the conception of action could have beyond purely metaphysical interests. This would lead to a pragmatist pressure to revise either the conception of action or the conception of intentionality/abilities/moral responsibilities so as to make sure that actions are not totally disconnected from intentionality/abilities/moral responsibilities.

With this being said, I am not going to provide further supports or proofs regarding the truth of (1)-(3). They are some basic assumptions that I take for granted for the sake of my argument. I do not deny that, *maybe*, the truth of (1)-(3) (together or separately) turns out to be more controversial than I expect. However, I do believe that most readers are willing to accept them as true or at least highly plausible. Now, given that (1)-(3) are true, I will argue for the following three theses in chapters 2-4:

- (4) Events are not things that we can intentionally bring about.
- (5) Events are not things that we can have abilities to bring about.
- (6) Events are not things that we can be morally responsible for.

These theses mean that, even if (some) actions are events and it makes sense to speak of particular actions, we cannot perform those actions intentionally. We cannot have abilities to perform those actions. And, we cannot be morally responsible for performing those actions.

It is surely more natural to speak of “to perform actions” and “to bring about events.” But it creates a difficulty for me that (1) and (4) (also, (2) and (5)) do not pair exactly with each other to give what I want to conclude. To avoid this problem, I propose that, in order for the thesis that actions are events to stand any chance, it seems necessary to allow a conditional

interchangeability of two kinds of expressions: if an action is an event, then for this action/event, “to perform the action” and “to bring about the event” have the same meaning. I will sometimes use “to perform *particular* actions” and “to bring about events” interchangeably in the rest of this dissertation based on the context. So, if the action of my punching a person’s face is the same as the event of my punching a person’s face, then replacing “perform the action of my punching a person’s face” with “bring about the event of my punching a person’s face” would not change the truth of any sentence that contains such expressions. If it is an accident for me to do the action of my punching a person’s face, it is an accident for me to bring about the event of my punching a person’s face. If it is wrong for me to perform the action of my punching a person’s face intentionally, it is wrong for me to bring about the event of my punching a person’s face intentionally. Therefore, for my opponents to reject the pairing of (1) and (4) (also, (2) and (5)), they need to deny that actions are events, which is fine with me.

If (1)-(6) are true, they constitute three independent arguments—the intention-based argument, the ability-based argument, and the responsibility-based argument—to reject the view that actions are events, since (1) and (4), (2) and (5), (3) and (6) all independently lead to the conclusion that actions are not events. It should be emphasized that each reason from its own conceptual perspective is an independent argument for the conclusion. The force of each reason does not rely on the force of the other two, not only because each pair of premises (e.g., (1) and (4)) can work separately, but also because my arguments for (4)-(6) are also independent of each other. For example, my reason why we cannot be morally responsible for events is not related in any way to my reason why we cannot have abilities to bring about events.

3. Kim versus Davidson

Events have some essential features, e.g., events can enter into event-causal relations as causes and effects. This must also be assumed in order for the CTA to be true. But my ontological objection will focus on the particularity of events. I believe, as most philosophers would agree, that events are particulars.² That is, events are things that occur at a specific place and time instead of things that can recur or be exemplified or instantiated at different places and times. Even though the opposite idea that events are universals has its defenders as well, most notably defended by Richard Montague (1969) and by Roderick Chisholm (1970), this thought has not found many supporters in recent work.

But there are different theories about events as particulars. The two most influential theories are Kim's fine-grained view and Davidson's coarse-grained view. According to Kim (1976), an event is a particular object's exemplifying a property at a time. It is something that has a structure, "built upon" more fundamental categories such as objects, properties, and times. Accordingly, two events are identical with each other just in case they have the same constitutive objects, properties, and times. This view is labelled "fine-grained" because events that are intuitively the same might be regarded by it as being distinct. For example, by its standard of individuation, Mary's walking to the campus around 2 pm yesterday is not the same event as Mary's moving to the campus (by walking) around 2 pm yesterday.

Davidson, inspired by Anscombe and in contrast to Kim, understands events in a rather coarse way. Instead of taking Mary's walking and her moving as two distinct events, he thinks "Mary's walking to the campus around 2 pm yesterday" and "Mary's moving to the campus around 2 pm yesterday" are merely two different descriptions of the same event. Instead of being

² Not all particulars are events, e.g., objects.

a structured thing based on more fundamental ontological categories, Davidson thinks that events, as concrete particulars, constitute a separate ontological category. An event can be described or characterized in many different ways without losing its identity, making it really coarse-grained.³

In this dissertation, instead of being neutral about events individuation, I adopt Davidson's coarse-grained view. Among other criticisms (e.g., Bennett 1988, 1991) against Kim's view, I think Davidson's criticism is convincing and points to a problem that is most relevant to my project. Davidson objects that Kim's so-called events are usually particulars but not necessarily so.

Some actions are difficult or unusual to perform more than once in a short or specified time, and this may provide a specious reason in some cases for holding that action sentences refer to unique actions. Thus with 'Jones got married last Saturday', 'Doris wrote a cheque at noon', 'Mary kissed an admirer at the stroke of midnight'. It is merely illegal to get married twice on the same day, merely unusual to write cheques simultaneously, and merely good fortune to get to kiss two admirers at once. Similarly, if I say, 'There is an elephant in the bathtub', you are no doubt justified in supposing that one elephant at most is in the bathtub, but you are confused if you think my sentence contains a singular term that refers to a particular elephant if any (Davidson 2001, pp. 167-8).

Take the event of Mary's walking to the campus around 2 pm yesterday as an example. The key question is what is "around 2 pm"? I think there are three alternatives. First, if it means a time period that is normally long enough for a *completed* walking to the campus (e.g., 1:50 pm to 2:15 pm), then Davidson's point makes perfect sense: it is merely physically *unusual* for a person to walk to the campus within 25 mins twice.

³ But what is Davidson's principle of individuation of events? At first, Davidson proposes that an event is marked out as distinct from other events by its unique causal history. That is, two events are identical to each other just in case they have the same causes and effects. But Davidson, persuaded by Quine's criticism (Quine 1985), abandons his earlier view and adopts Quine's suggestion that "events, like physical objects, are identical if they occupy the same places at the same time" (Davidson 1985).

If it means a time period that is not long enough for a completed walking to the campus in most cases (e.g., 1:59:50 pm to 2:00:10 pm), then this event of Mary's walking is probably just a segment of the whole walking. However, if it is a segment, there is no reason to think that there cannot be two segments of walking within 20 seconds. Maybe Mary changes her mind and starts to walk back only for a very short moment. Thus, there are two separate segments of walking to the campus within that 20 seconds, one before the short moment of changed mind and one after. So, Davidson's criticism still applies.

If around 2 pm is shortened to a time slice (e.g., at 2 pm exactly), then it is no longer clear that Mary's walking at that particular moment is an event. It seems more reasonable to take it to be a state of Mary at that particular moment, namely, that she *is* walking to the campus. Is there any event that can happen at a particular moment? I doubt it. For example, I kick a ball and cause an event: the ball's hitting my neighbor's wall. A closer look at the details of this event can tell us that it is a physical process—including the deformation of the ball (also the wall)—that takes time to complete. This will not change even if the hitting happens between two rather hard objects like billiards balls. In such a case, deformation of billiards balls would happen much faster but still takes time.

How about some *apparent* events that can only occur at a time-slice, such as a moving (or even static) object's occupying a particular spatial location? There are two difficulties with this kind of time-slice "events." The minor one is that it assumes that a physical object (e.g., a car) has a clear-cut spatial boundary, which is not true according to microphysics. The major problem is that time-slice "events" do not appear to be real events since they do not enter into event-causal relations. We can imagine a car that is making straight line motion at a constant speed. The car occupies a particular spatial location s_1 at a particular time t_1 . On the one hand,

this time-slice “event” does not last long enough to exert any physical influence onto the world. Think about pushing an object to speed up. The gained momentum of the object is the result of the force acting *for a certain amount of time*: $F \cdot t = m$. If t is zero, the object’s physical state will not be changed at all.

On the other hand, what can be the cause of this alleged event? If anything, it must be caused by the immediately prior “event” of the car’s occupying location s_0 at time t_0 . Even if someone wants to claim that it is actually caused by the event that initially causes the car to be in motion, it seems that the first mover can only cause the current “event” *through* those “events” between them—causation is transitive. So still, the current “event” must be caused by its immediate antecedent.

But how can we identify such an immediately prior location of space and time considering that space and time are usually taken to be continuous?⁴ Even setting this aside, this causal picture seems to contain too many causal relations. Even if space and time turn out to be discontinuous, there would be almost countless particular spatial and temporal locations covered by a very short period of moving. If each “event” of the car’s occupying a space at a time is caused by the prior “event” of the car’s occupying the immediately prior space at the immediately prior time, there would be almost countless causal relations accordingly. It is absurd to have so many causations for a situation in which the car is only moving at a constant speed solely due to its inertia.

Based on these difficulties of regarding the car’s occupying s_1 at t_1 as an event, I think it is rather a state of the car or a fact about the car. Therefore, no matter how “around 2 pm” is

⁴ Maybe quantum theory about time and space can resolve this difficulty.

interpreted, Kim's view is either susceptible to Davidson's objection that the alleged event is not necessarily a particular or vulnerable to my objection that the alleged event is not a real event.

The above discussion might not be convincing enough to change Kim's followers' mind about event individuation. It is at most a plausible argument against his fine-grained theory of events, especially regarding the particularity of events. So, my commitment to Davidson's coarse-grained view is still a basic assumption that I accept for this dissertation. I only try to show that it is a reasonable assumption (and a necessary one since my ontological objection needs to focus on the particularity of events). For those who do not find this assumption plausible or reasonable, the force of my ontological objection to the CTA would be confined within Davidson's ontology of events. Given the influence of Davidson's coarse-grained view and his contribution to the debate about the CTA, such a restricted goal would still have enough philosophical significance.

4. Act-Universals

That events are particulars is the main reason why actions are not events. That is, actions are not events mainly because actions are not particulars.⁵ But what are they if not particulars? Even though I do not positively argue for this position, I take actions to be universals and, when I perform an action, an act-universal is exemplified partly by me in a particular situation.⁶ It does not make sense to ask this question: What is *this* action I just performed?⁷ Just like it does not make sense to ask me when I am happy: What is *this* happiness of mine? Strictly speaking, it is

⁵ But why do I not choose "Actions Are Not Particulars" as the title and conclusion of my dissertation? Although that actions are not events is entailed by the claim that actions are not particulars, there is no influential alternative theory of action that takes actions to be non-event particulars. Because of that, the event-based theory of action is interesting enough to be my main target of criticism.

⁶ If I drive to the campus, it seems absurd to think that I am the only thing that exemplifies the act-universal in question. At least, it involves a car among many others.

⁷ As Bach argues, "we need not count actions but only agents and events" (1980, pp. 119).

ambiguous to say that I perform an action. It can mean that there are many act-universals (different kinds of acting) and one of them is exemplified by me (and other things). Or it can mean, mistakenly, that there is an act-universal and I perform an instance of it. The very point I try to make is that an instance of an act-universal is not something that we can *perform*.

So, on my account, actions are things that can recur or be exemplified at different places and times. This may sound surprising, but it should not be. In our daily life, it is very common to speak of actions with the implication that they can recur. For example, after having a satisfactory meal, I may say to myself “I will do this again.” By saying so, I am expressing an intention to repeat what I just did, namely, eating a meal at this restaurant. This kind of expression carries the impression that actions can be done repeatedly. Philosophers who believe that actions can only be particulars might be able to find a way to explain away the impression. But it is equally, if not more, justified to leave open the possibility that actions are universals, and to adopt the terminology of “act-universals” to refer to those actions that allegedly can be done at different temporal-spatial locations.

Also, speaking of act-universals exclusively does not stop us from making various claims about actions. It does not create any problem for us to understand these sentences: I performed an action of hurting Jones; I have the ability to hurt Jones; and I am morally responsible for hurting Jones. If anything, it is alleged act-particulars that create problems for understanding those sentences, as I will show in this dissertation.

This brings up a question: Why do I not use the terminology of “act-types” instead? A short answer would be that act-types are usually paired with act-tokens but I do not want to implicate the existence of act-tokens. That is because, according to Alvin Goldman who first raises the distinction between act-types and act-tokens, act-tokens are just particular actions:

An act-type is simply an act-property, a property such as mowing one's lawn, running, writing a letter, or giving a lecture. When we ascribe an act to an agent, we say that the agent exemplified an act-property (at a certain time). ... A particular act, then, consists in the exemplifying an act-property by an agent at a particular time. I shall call such particular acts: "act-tokens" (Goldman 1976, pp. 10).

Since act-tokens are particular concrete instances of act-types, I choose not to use both concepts in order to be consistent with my view that actions are not particulars. If I use expressions such as particular actions or act-particulars in this dissertation, especially when I discuss some other philosophers' views, it is not meant to be a sign of acceptance of the existence of particular actions or act-tokens. I only mean an *alleged* particular action, i.e., an event that is mistakenly taken by many philosophers to be an action.

But someone may ask: Can I really avoid the implication of the existence of particular actions at all by speaking of act-universals? Don't act-universals have instances which are particulars as well? I answer yes to the second question but no to the first. I do not deny that act-universals can be exemplified (when actions are performed) and have corresponding instances. However, the particular instances of act-universals are not necessarily particular actions or act-particulars. In this respect, act-universals are similar to some other universals, e.g., properties. The color of white, as a universal and a property, can be instantiated and has its own instances. But the instances of whiteness are not particular whiteness.⁸ Rather, they are white objects such as white tables or cars. They merely exemplify the color of white without being particular whiteness in themselves. Likewise, the particular instances of act-universals are not necessarily particular actions or act-particulars.

In another respect, act-universals are not like properties (unlike Goldman, I do not speak of act-properties). It seems that properties are exemplified or possessed by objects and those

⁸ Most people will accept this analogous example unless they are trope theorists.

objects are taken to be the instances of the properties. So, properties are usually exemplified by their own instances. For example, the color of white is exemplified by a white object (e.g., a white car) which is also an instance of whiteness. But it is not easy to apply this relation of exemplification between properties and their instances to the idea of “act-properties.” On Goldman’s account, it is agents that are said to exemplify an act-property (at a certain time) when an act is ascribed to the agent. He contends that attributing an act-property to an agent is not problematic because “just as *owning* a Jaguar is a property that can be exemplified by John at time *t*, so *buying* a Jaguar is a property that can be exemplified by John at time *t*” (Goldman 1976, pp. 10). When it comes to the instances of act-properties or act-types (Goldman take them to be the same thing), however, act-tokens take the position. So, even though act-properties are exemplified by agents, their instances are not agents but particular actions or act-tokens. How can a property be exemplified by something other than its instances? This unusual feature of act-properties raises the worry whether act-universals are best understood as a sort of properties.

Recently, Jennifer Hornsby (2012) proposes and defends a view that actions should be understood by reference to a category of activity where activities are not particulars. According to Hornsby, to act is to engage in a certain type of activity: “one needs to think of a person’s raising her arm as a type of causal activity in which she engages” (2012, pp. 234). Hornsby then makes two points regarding this notion of activities. First, these activities are types (and universals) rather than particulars.

Now each of raising one’s arm, strolling, walking, reading .. is considered *an* activity by virtue of its being a *type* of activity. And when ‘raising an arm’, ‘strolling’, ‘walking’, ‘reading’ name types of activity, they are not count nouns (Hornsby 2012, pp. 237).

Second, to say that an activity of some type is going on or being engaged in by an agent is not to say that there is any particular event or process. This is so even if some event, as an instance of the activity being engaged in, must occur at the same time. Hornsby believes that

activities are, in this respect, analogous to *stuff* with non-count names (e.g., gold). She uses beer as an example:

Now if it is actually true that there's beer in the fridge, there will no doubt be there some thing or things that can be reckoned an instance or instances of beer ... Still to say that there is beer somewhere is not to say that there is any sort of instance of it there. Rather it is to say that there is there stuff of a certain type—sc. beer. Where 'beer' names a sort of stuff, and 'thing' applies to the particulars that are occupants of the spatial world, the word 'beer' in 'There is beer in the fridge' does not stand for any thing. ... likewise, 'Activity [of some type] was going on at a time' does not say that there was any particular in progress then. Even if it can only actually be true that someone was strolling at a time if there was at that time some event that can be reckoned as instance of strolling on the part of that person, still to say that the person was strolling is not to say that there was such an event (Hornsby 2012, pp. 237).

Hornsby's two points need to be clarified a little. Her first point is the metaphysical claim that activities are not particulars. Her second point is semantic. That is, to speak of an on-going activity does not necessarily mention or refer to any particular event even though there must be some particular event that exemplifies the activity (so every time chess is played, there must be a particular chess game understood as an event or process).

In the light of Hornsby's discussion of actions and activities, it seems plausible to understand act-universals as activities (but I am not officially endorsing and defending this theory). According to this understanding, to perform an act-universal is to engage in a particular type of activity. Act-universals are activities or one's engaging in the activities. Here, we do not need to distinguish between an activity and one's engaging in the activity. Think about speech as a type of activity and giving a speech. If giving a speech only means a kind of action without the reference to a particular speech, then there seems to be no difference between the meaning of "speech" and that of "giving a speech." This is different from table as a kind of object and moving a table as a type of action. They obviously mean different things. This contrast suggests that activities are essentially and inherently active in the sense that it cannot be an activity unless it involves some agent's engaging in it. When an activity is concerned, it necessarily involves the

idea that it is an activity of someone or someone is engaging in it. Without one's engaging in the activity, the activity is no longer active and stops being an activity.

Now, understanding act-universals as activities seems to support the earlier claim that, similar to properties, an instance of an act-universal *might* not be a particular action. Even though there must be some particular event (or process) that exemplifies the activity, the event (or process) itself does not need to be a particular activity. Consider an example of playing soccer. Every time soccer is played, there is a particular event—a particular soccer game—that exemplifies the activity of soccer. It is true that soccer players bring about the particular soccer game by engaging in the activity of soccer, but it is awkward to say that those players are also engaging in a particular activity. If so, the event of this particular soccer game is simply not a particular action.

It is surely possible to refer to the particular game as an activity if “activity” is used loosely to mean a mere collection of human behaviors and movements of the relevant physical objects (e.g., the soccer ball). But the particular game cannot be a *bona fide* activity because it is not something that people can engage in. Even on the level of particulars, the goal that guides the players' behaviors is still the goal of playing soccer as a *type* of activity. The same thing can be said for the rules, the procedures, etc. But if the players are not pursuing a different goal, or following a different set of rules and procedures, how can they be engaging in a different activity? This explains why it would not sound right if a soccer player says that yesterday I engaged in a particular soccer-playing activity and today I am engaged in another one.

I realize that the argument above is not adequate to turn the tide against the view that there are particular actions as instances of act-universals. However, that is not what I intend to achieve here. My intention is not to show that an instance of an activity is not a particular

activity, or that an instance of an act-universal is not a particular action, by endorsing and defending Hornsby's activity-based theory of action. Rather, by appeal to Hornsby's theory (as a seemingly plausible theory) and my further explication, I illustrate a real and serious possibility that instances of act-universals—the events that exemplify them—*might* not be actions.

By arguing that the events as instances of act-universals might not be particular actions, I avoid begging the question against my opponents who believe they are. What I try to do is to find a neutral ground so that I can present the argument for the view that events cannot be actions, including those events that instantiate act-universals. And philosophers who think otherwise also have the opportunity to dispute my argument or to provide their own argument for the view that events can be particular actions. Therefore, when I formulate principles and arguments in this dissertation, I speak of act-universals and events that exemplify act-universals.

Instead, notions such as act-tokens and act-types seem to presuppose an event-based understanding of actions. Therefore, it would beg the question against my view to ask me to adopt or accommodate those notions at the beginning. For the same reason, I choose not to use Davidson's expression "acting intentionally under a description" because this expression seems to presuppose the conjunction of the view that actions are events and a coarse-grained theory of events. After I complete the ontological objection to the view that actions are events, I will discuss why adopting Davidson's expression does not defend and save this view, which I take to be a legitimate question. But it would not be fair to ask me to accommodate this expression in formulating my own argument.

5. The Logical Form of Action Sentences

There is an obstacle I need to consider and respond to before I discuss my ontological objection to the view that actions are events. This obstacle is an influential argument proposed by Davidson in his paper “The Logical Form of Action Sentences” (2001, essay 6). There, Davidson makes a semantic claim that the logical form of adverbially modifiable action sentences involves existential quantification over actions that are understood as particular events. Let us consider the following action sentences:

(1) John sings a song.

(2) John sings a song in the library.

Davidson points out that if we analyze (1) and (2) in a traditional way based on subject-predicate semantics:

(1*) Sings (John, a song)

(2*) Sings-in (John, a song, the library)

then, it would be difficult to explain the plain fact that (2) entails (1). That is because, “sings” denotes a different predicate from the one denoted by “sings-in” and, consequently, no logical connection is revealed between (1) and (2) by rephrasing them into (1*) and (2*).

By contrast, Davidson makes the proposal that action sentences such as (1) and (2) should be analyzed as existential claims about actions as particular events.

(1**) There is an event e such that e is a singing of a song by John.

(2**) There is an event e such that e is a singing of a song by John and e is in the library.

If (1) and (2) are rephrased in this way, treating actions as events, it easily solves the question of why (2) entails (1).⁹ Then according to Quine's criterion of ontological commitment (as controversial as it is), we are ontologically committed to the existence of actions as events. As Davidson remarks, concerning the thesis that actions are events, "our common talk and reasoning about actions is most naturally analyzed by supposing that there are such entities" (2001, pp. 109).

Since the publication of his paper, Davidson's semantic consideration is widely thought to support the view that actions are events. However, this argument has its critics as well. In general, there are two kinds of objection to Davidson's semantic argument. First, by appealing to the ontology of events, Davidson's solution leads to other problems and difficulties (e.g., Aune 1977, Trenholme 1978). Second, there are alternative ways to make sense our ordinary language about actions that do not involve the existential claim about events (e.g., Aune 1977, Horgan 1978).

More recently, Maria Alvarez (1999) tries to undermine Davidson's semantic argument in a way slightly different from the second kind of objection. She does not deny that the semantics of action sentences requires assuming the existence of events. What she denies is that actions are those events that we are ontologically committed to. According to Alvarez, "the simplest way of formalizing actions sentences which captures all and only licit inferences requires quantification *only* over the results of actions" (Alvarez 1999, pp. 213).

Here, I do not intend to give any detailed discussion beyond this brief survey of critical thoughts on Davidson's semantic argument for seeing actions as events. I only try to show that

⁹ According to Davidson, this method of analyzing sentences is not limited to action statements with adverbial modifications, but also applies to some other statements, e.g., causal claims.

the success of the semantic argument is not without skeptics. Thus, it is not necessarily an insurmountable obstacle standing before my project, even though I need not be the person who surmounts it.

6. A Summary of Chapters

The main body of this dissertation has three chapters. In chapter 2, I provide an intention-based argument for the conclusion that events are not things that we can intentionally bring about. In chapter 3, I provide an ability-based argument for the conclusion that events are not things that we can have abilities to bring about. In chapter 4, I provide a responsibility-based argument for the conclusion that events are not things that we can be morally responsible for. Each chapter constitutes an independent argument against the view that actions are events. And if I am correct in believing that actions are not events, CTA is false because actions are not things that can enter into event-causal relations.

In chapter 5, I discuss a response that could be raised against my ontological objection to CTA, especially on behalf of Davidson. The response is based on Davidson's view that an action is done intentionally by its agent *only under a certain description*. I will show that this Davidsonian response does not neutralize the force my ontological argument.

CHAPTER 2: THE INTENTION-BASED ARGUMENT

As I state in the introduction, I will argue that events are not things that we can intentionally bring about in this chapter. My argument is based on an investigation of the relation between intention and intentional action. A commonsense view regarding this relation is the so-called the “Simple View” (or SV for short):

SV: A person performs an action *A* *intentionally* only if he *intends* to *A*.

So, a necessary condition of a person’s performing an action intentionally is her intending to perform that very action. This commonsense view, although intuitive, has been challenged by some influential counterexamples and, as a result, is widely taken to be false.

I argue, however, that SV is *not* undermined by these counterexamples. This is important because the truth of SV, together with another claim, would lead to the truth of my conclusion in this chapter. The other claim is that we cannot intend or form an intention towards events. In other words, if actions are events, then we cannot intend to perform them. The argument is formulated as follows:

- (1) A person performs an action *A* intentionally only if she intends to *A*.
- (2) Events are not things that we can intend to bring about.
- (3) Therefore, events are not things that we can intentionally bring about.

This chapter has four parts. In section 1, I defend premise (1) by defending SV against two immediate challenges. In section 2, I defend (1) by defending SV against its most influential counterexample: Bratman’s videogame case. In section 3, I defend premise (2) of the argument by explaining why we cannot intend or form an intention towards events. In section 4, I consider and reject an objection that is based on an alternative formulation of the Simple View: A person

performs a particular action *A* of type *B* intentionally only if he intends to perform an action of type *B*.

1. Two Immediate Challenges to SV

Premise (1) of my argument, i.e., SV, is the commonsense view of the relation between intention and intentional action, namely, that a person performs an action *A* intentionally only if she intends to *A*. As Frederick Adams remarks, the attractions of SV are many and we stand to lose heavily if we abandon it. For that reason, we should not give it up without compelling reasons (Adams 1986, pp. 284).¹ But the opponents of SV believe there *are* compelling reasons to abandon SV, namely, that it is subject to clear counterexamples. Two kinds of cases seem to pose an immediate challenge to SV: spontaneous and subsidiary actions and unintended side-effect actions.

The first kind of case concerns spontaneous and subsidiary actions. When I drive to the campus, there seem to be a series of subsidiary things that I do intentionally but without specifically intending to do them (e.g., checking my mirrors, shifting gears). Also, if another vehicle suddenly cuts into my lane, I may spontaneously hit my brakes. Doing so would appear to be an intentional action, but it happens so fast that I do not have enough time to *form* an intention to hit the brakes.

These cases, however, do not *clearly* illustrate a situation in which a person performs an action intentionally without intending to do that action. For one thing, it is not obvious that the person does *not* intend to perform the action in question. In driving to the campus, I may not

¹ The critics of SV are likely to agree. As Bratman remarks, “The Simple View has its virtues. It recognizes the distinctiveness of intentions and provides a straightforward and initially plausible account of the relationship between such intentions and intentional action. It is a view toward which common sense initially leans, as well as a view implicit in many discussions of intention in moral philosophy” (1999, pp. 112).

clearly and explicitly represent those subsidiary actions as intended entities in my mind, but this is not sufficient to prove that there are no such intentions. On Ginet's account (2016), for instance, the conscious awareness of the intention is not a necessary condition of the possession of the intention:

To have at the time of acting the intention to thereby do such-and-such one need not have a conscious thought that one so intends. We more or less continually have intentions that we do not consciously think about at the times we have them but would readily admit to having if asked (Ginet 2016, pp. 227).

Suppose the critic insists, however, that, in hitting the brakes spontaneously, I lack the time to form an intention to hit the brakes. This might well be true. But as Adams and Mele (1989) argue, a person may *acquire* an intention without *forming* one, and we need to be careful not to conflate them:

Notice that if forming an intention is itself an action, it does not follow from the fact that no prior intention is *formed* in a particular case that the agent *had* no prior intention in that case. For, it seems, some intentions are *passively* acquired – as is the case with many of our beliefs, desires, and thoughts (Adams & Mele 1989, pp. 521).

On this view, while actively forming an intention involves a deliberating or decision-making mental process prior to the forming of the intention which might take some time, a person can still *passively* acquire an intention without going through the mental process of deliberation. Here is an example of how a person might develop the relevant skill to acquire intentions without deliberation.

David is a recently licensed driver. Among other things, every time he meets a yellow light, he needs to explicitly and consciously calculate and deliberate in order to form the correct intention regarding whether to stop or continue to drive (depending on his current speed and how close he is from the intersection). But after a few months of repeated training and more driving experience, he realizes that he no longer needs to explicitly deliberate about whether to hit the brake when the traffic light turns yellow. Then, when he sees a yellow light, he can immediately

acquire the intention to hit the brakes (or not) without explicitly deliberating about it. In other words, he has developed a habitual instinct to acquire intentions regarding his reaction in this kind of situation. Therefore, in cases of spontaneous action, it is plausible that the agent *has* the relevant intention, even though the intention may not be *consciously* possessed or *actively* formed.

Alternatively, proponents of SV could question whether the agents really perform the action intentionally in those cases. It is not obvious that they do. My hesitation here is most strongly motivated by modified scenarios in which spontaneous and subsidiary actions turn out to be mistakes that the agent makes. The modification of the case is supposed to make our intuition clearer and more distinct without changing the essence of the case.

In the case of hitting the brakes, let us add that there is a bomb attached to the brake and I know it will explode if I hit the brakes. So, if it is up to me, I will never hit the brakes intentionally. But due to the suddenness of the situation, I spontaneously hit the brakes. It seems clear here that I do *not* hit the brakes intentionally. If I had hit the brakes intentionally, it would have been done for a reason. And given that I have an overwhelmingly strong reason not to do so, I would have been guilty of irrationality. But my hitting the brakes is merely a spontaneous reaction and I do not make any mistake with respect to reasons or reasoning. Therefore, I do not do it intentionally. Likewise, suppose that I intend to wait for a friend when I enter my car, but due to a habit I have developed over the years, I automatically shift to the reverse gear after I sit in my car. Do I intentionally do so? The answer is no.

In both modified cases, there is a fact, namely that hitting the brakes will cause the explosion or that I am waiting for a friend, and this fact makes it true that my spontaneous or subsidiary action is a mistake. In other words, these facts constitute reasons not to act as I do.

However, these facts do not play any role in the causal history of my action. They merely render my action a mistake without causing the action. My action is performed in exactly the same way as it is in the original case in which no such fact exists. Since an intentional action cannot become unintentional (or nonintentional) *simply* because there is a reason not to perform the action, if my spontaneous or subsidiary action is not intentional in the modified case, it is not intentional in the original case as well. Hence, my hitting the brakes or shifting the gear in the original case is not intentional.

The second kind of case at odds with SV are cases of unintended side effects. Consider the following example by Gilbert Harman:

In firing his gun, [a] sniper [who is trying to kill a soldier] knowingly alerts the enemy to his presence. He does this [i.e., alerts the enemy] intentionally, thinking that the gain is worth the possible cost. But he certainly does not intend to alert the enemy to his presence (Harman 1976, pp. 433).

In this case, alerting the enemy is merely a side effect of what the agent *really* intends to do, namely, kill the soldier. The occurrence of this side effect is fully expected and foreseen by the agent, but it is *not* what the agent intends to do. He is aware that he will alert the enemy, but he is not committed to bringing it about by aiming at it or trying to achieve it. This case is, therefore, a purported counterexample to SV because although the sniper apparently alerts the enemy *intentionally*, he does *not* intend to do so.

In reply, I argue that the sniper's alerting the enemy *might not* be an intentional action. Here, I agree with Mele and Moser when they write:

There is a middle ground between *A*-ing intentionally and *A*-ing unintentionally. We locate "side-effect actions" of the kind in question on that ground. Insofar as such actions are not done unknowingly, inadvertently, or accidentally, they are not unintentional. Insofar as the agent is not aiming at the performance of these actions, either as ends or as means to (or constituents of) ends, they are not intentional either. We shall say that they are nonintentional (Mele & Moser 1994, pp. 45).

According to Mele's distinction, because the sniper knowingly alerts the enemy, this action is not unintentional. However, since alerting the enemy is not what the sniper aims for, as an end in itself or as a means to an end, he does not do it intentionally either. Rather, his alerting the enemy should be categorized as a nonintentional action. If this is right, then Harman's case is not a counterexample to SV, since the action is not an intentional action in the first place.

Mele's view provides an explanation for the so-called "doctrine of double effect": sometimes it is permissible to cause a harm as a side effect (or "double effect") of bringing about a good result even though it would not be permissible to cause such a harm as a means to bringing about the same good end. Although the truth of this doctrine is controversial (see McIntyre 2014), it is an intuitively appealing and practically useful principle that helps elicit the moral significance of the difference between unintended side-effect actions and intended actions. Mele's view offers a metaphysical perspective—about the intentionality of agency—for why a normally impermissible behavior might be permissible when it is done as an unintended side-effect action: a foreseen but unintended side-effect action might not be an intentional action. Thus, unintended side-effect actions do not constitute a counterexample to SV.

2. Bratman's Videogame Case

Besides the two kinds of cases that I just dismissed (spontaneous or subsidiary actions and unintended side-effect actions), Michael Bratman raises perhaps the most influential counterexample to SV: the videogame case. The case can be stipulated as follows. Suppose I am playing a videogame in which I am to guide a "missile" into targets. There are two targets and if I hit one of them I will win. However, the game is known to me to be so designed that it is impossible to hit both targets. If both targets are about to be hit simultaneously the machine just shuts down. Both targets remain visible to me, so I can see which target I hit if I hit either target.

Since I know it is difficult to hit either target, I decide to shoot at both targets simultaneously. I see the risk of shutting down the machine as outweighed by the increase in my chances of hitting a target. I proceed to try to hit target 1 and also to try to hit target 2. Suppose I do hit target 1 in just the way I was trying to hit it, and in a way which depends heavily on my considerable skills at such games. It seems that I hit target 1 *intentionally*.

If SV is true, given the intentionality of my hitting target 1, I must intend to hit target 1. Then, due to the symmetrical structure of the case, I must intend to hit target 2 too. At this point, Bratman proposes two rational constraints on intentions: strong consistency and rational agglomerativity. They can be expressed as follows:

Strong Consistency: It is rational of a person to intend to *A* only if *A*-ing is consistent with her beliefs.

Agglomeration: It is rational of a person to intend to *A* and intend to *B* only if she intends to *A* and *B*.²

By “*A*-ing is consistent with her beliefs,” Bratman means that it is possible for her to *A* given that all her beliefs are true. Two clarifications are needed about this consistency condition. First, this condition does not require that she believes that it is possible to *A*. This is because this reading faces a sort of counterexample in which a person is agnostic regarding the possibility of doing something while doing it successfully is still consistent with her beliefs. For example, without believing if the mall is closed or not, it would be irrational of me to believe that buying a

² Both Strong Consistency and Agglomeration, if put more precisely, should include a qualification of “other things equal.” This qualification suggests that it might be, all things considered, rational of a person to violate these two rational constraints given a stronger reason to do so.

birthday gift for my wife is still possible.³ If it turns out that the mall is actually closed, then it would be impossible to buy the gift. But if I decide to drive to the mall when this information is unavailable to me, this seems to be a rational decision because buying the gift successfully is consistent with my beliefs.

Second, this consistency condition expressed by “it is possible for her to *A* given that all her beliefs are true” is still ambiguous. Under one natural interpretation, this condition requires that she does not believe that she will not *A*, and if so, her *A*-ing is not consistent with her *true* belief that she will not *A*. However, it can also be understood in a weaker sense, especially for those philosophers who are less stringent about the belief constraint on what one can rationally intend to do (e.g. Davidson, Adams, and McCann; Adams provides a survey of various views on this issue in his 2007 paper). This condition may require merely that the agent does not believe that it is impossible to *A*. According to the second understanding, her *A*-ing is consistent with her belief that she will not *A*. For instance, even if my belief that this lottery ticket will not win is true, it does not eliminate the metaphysical possibility that the lottery ticket might be a win (assuming drawing the winning number is a random process). Even if it turns out to be a loss, it could have been a win. This seems to be a reasonable position because my effort to win the lottery would not be rendered irrational by my belief that I will not win. Many people who buy lottery tickets believe so. Instead, what would make one’s intention (and effort) to *A* irrational is her belief that *A*-ing is simply impossible.

Due to the ambiguity, it is not clear what exactly Bratman means by the consistency condition. Fortunately, we need not solve the puzzle here. In the videogame case, I not only

³ I assume an ordinary notion of belief so that I either believe or not believe something (instead of construing beliefs as degrees of confidence).

believe that I will not hit both targets, I also believe that I cannot hit both targets. No matter how the consistency condition is interpreted, the combination of Strong Consistency and Agglomeration entails the following principle which applies to the videogame case (I call it the Consistency of Actions):

Consistency of Actions: It is rational of a person to intend to *A* and intend to *B* only if the conjunction of *A*-ing and *B*-ing is consistent with her beliefs.

Since I know that I cannot hit both targets, according to Consistency of Actions, it would be irrational of me to intend to hit target 1 and intend to hit target 2. But as Bratman points out correctly, there is no irrationality committed by me in the videogame case:

It seems clear that I need be guilty of no such form of irrationality: the strategy of giving each game a try seems perfectly reasonable. If I am guilty of no such irrationality, I do not have both of these intentions. Since my relevant intention in favor of hitting target 1 is the same as that in favor of hitting target 2, I have neither intention. So the Simple View is false (Bratman 1999, pp. 114– 15).

Therefore, it is false that I intend to hit target 1 and intend to hit target 2. Due to the symmetry of the case, I either have both intentions or have neither. Therefore, I have neither intention. The upshot seems to be that, in the videogame case, I intentionally hit target 1 without intending to do so. If so, SV is false.

I agree with Bratman that intentions are subject to various rational constraints due to the functions or roles they play in human actions. One of those constraints is a belief-related consistency as Strong Consistency represents. I think Strong Consistency is true (setting aside the complexity that its formulation is ambiguous).⁴ I also agree with Bratman that the planning and coordinating role of intentions requires, for the sake of rationality, that an intention should not only be internally consistent and consistent with one's beliefs but also be consistent with

⁴ For a detailed discussion of different versions of belief-related constraint on intentions, see Adams (2007).

other intentions.⁵ However, when Bratman says “Given the role of intentions in coordination, there is rational pressure for an agent to put his various intentions together into a larger intention” (Bratman 1999, pp. 134), I disagree. I do not think there is always rational pressure for an agent to combine her separate intentions into a larger compound intention.⁶

For Bratman, the explanation of why two intentions need to be consistent with each other in a broad planning and coordinating context is because that there is rational pressure to combine separate intentions into a larger compound intention and the resultant compound intention needs to be consistent with one’s beliefs. Therefore, if the objectives of the two intentions are not compatible with each other (and the agent believes so), she cannot rationally intend them even separately. So Bratman needs both principles – Strong Consistency and Agglomeration – to account for the consistency constraint among different intentions. I disagree with this explanatory scheme because Agglomeration is false. And also, I suggest that there is a more plausible account of the consistency constraint among intentions.

First, an unqualified or unrestricted Agglomeration is unlikely to be true. It would require us to agglomerate totally irrelevant intentions, such as my intention to have rice for dinner tonight and my intention to go shopping this coming weekend. As Zhu (2010) realizes, there seems to be no “rational pressure” on me to form an intention to both have rice for dinner tonight and go shopping this coming weekend. Also, as Steven Sverdlik points out:

Indeed if [Agglomeration] were true without qualification that rational agents agglomerate their intentions then each rational agent would always have one enormous

⁵ The question of internal consistency arises because an intention may represent a complex plan which involves multiple more basic actions that need to be consistent with each other.

⁶ McCann (1987) provides an objection to Bratman’s videogame case based on a similar rejection of Agglomeration. But McCann’s argument is different from the one I am going to give. His argument lies in a notion of trying to A that necessarily involves an intention to A. Even though I find such a notion highly plausible, it raises controversies that I want to avoid (admittedly, both notions are pretty complicated and may even be subject to different interpretations). So, in this chapter, I try to provide a direct counterexample to Agglomeration.

compound intention that was guiding her behavior. It seems likely that such an intention would interfere with – rather than facilitate – an agent’s satisfying her desires (Sverdlike 1996, pp. 517).

This is not meant to be a knockdown argument against Agglomeration. To save it, a distinction can be made between two actions being incompatible with each other and two actions being competitive with each other (Zhu 2010). To see the difference, I borrow Ginet’s case of double doors:

I confront a set of double doors through which I wish to go. I know that the doors open only one way, either toward me or away from me, but I don’t remember which. I am in a hurry, so I simultaneously push away from me on one of the doors and pull toward me on the other. In pushing on the one door I am trying to push it open and in pulling on the other door I am trying to pull it open. But it is not the case that I intend to push open the one door *and* intend to pull open the other. It would be irrational of me to intend to do *both* of these two things when I know that doing both is impossible. But it is not irrational of me to *try* to do both things while knowing that at most one attempt will succeed, which entails my intending to make it the case that either I push open the one or pull open the other. (Ginet 2014, pp. 22)

Ginet’s case is similar to the videogame case. But there is a crucial difference between them. For both cases, it is true that doing one thing is incompatible with doing the other (and the person knows it). But the reason for this incompatibility is different. And it matters. In Bratman’s case, the incompatibility of hitting target 1 and hitting target 2 is due to the fact that these two actions are competitive with each other. Doing one of them stops or prevents me from doing the other. In Ginet’s case, the physical design of the door does make it true that pulling one door open is incompatible with pushing the other open. But it is so not because doing one thing would prevent the other. Rather, it is because the physical design makes one of the two actions impossible. Which one is made impossible is unknown to the agent. But as long as one of the two actions has already been made impossible, of course they cannot both happen. In such a case, even if I did not perform the other action, I would not be able to do the impossible action anyway. They are not competitive with each other simply because there is no *causal* relevance

between them. As a matter of fact, these two courses of action are causally independent of each other.⁷

The rationale for Agglomeration seems to make better sense when the two intended actions are not only incompatible with each other but also competitive with each other. Only in such a competitive context would the rational pressure for coordinating separate intentions start to operate and have an effect on the agent. Otherwise, it would cause huge or even insuperable difficulties for the agent to intend and perform any action since she would be required to check the consistency of this intention with all her other intentions – no matter whether they are relevant or not – in order to form so many additional compound intentions. To avoid this problem, Agglomeration needs to be revised accordingly:

Agglomeration*: If a person believes that *A*-ing is competitive with *B*-ing, it is rational of her to intend to *A* and intend to *B* only if she intends to *A* and *B*.

Agglomeration* is weaker than Agglomeration in the sense that the rational pressure of agglomeration is now conditioned on the two actions being competitive with each other. This qualification avoids the original worry raised by Sverdlik and Zhu that Agglomeration would require the agent to combine totally irrelevant intentions. So, in Ginet's case, the agent can rationally intend to push one door open and pull the other open at the same time without combining them into a compound intention (Ginet's case does not undermine the Simple View). Due to the lack of competitive relevance, the agent makes no rational mistake not to form the additional intention.

⁷ I come up with this line of thought from reading Zhu (2010). I also borrow the term “competitive” from Zhu. I am not going to elaborate and defend a clear definition of the competitive relation between two courses of action in this paper. For more details, please read Zhu (2010).

The second problem with Agglomeration* and Bratman's explanatory scheme is that it is unclear why there is such rational pressure on the agent to agglomerate her separate intentions into a compound one (when they are competitive intentions). What is the source of the rational requirement of having a compound intention when a person possesses separate intentions? Unfortunately, Bratman does not give any explicit answer to this question.⁸

Maybe the truth of Agglomeration* lies in a principle like this: A person should intend logical consequences of what she has already intended. But this principle would put too much burden on human psychology since there are a large number of goals whose occurrences or existences are entailed by what we have already intended. Also, given that an intention contains a psychological or practical commitment to the intended objective (or its occurrence), it is simply not clear that, if a person is committed to two goals, she is therefore subject to rational pressure to give a certain commitment to the compound goal. Practical commitments are different from epistemic commitments in a crucial way, that is, epistemic commitments do not have any cost other than the psychological burden mentioned above. It makes better sense to think that I am under rational pressure to either believe the logical consequences of what I have already believed or give up one of my current relevant beliefs *when questions about the logically entailed beliefs are raised*. But practical commitments come with a cost when the intended objectives are competitive with each other. And I see no reason to rule out a priori the possibility that one's

⁸ Zhu (2010) offers an explanation for why an agent is under "rational pressure" to agglomerate all her so-called executively competitive intentions into a larger intention. However, although Zhu successfully shows that agglomerating intentions in Bratman's way could produce some benefits (e.g. a unified background framework for practical reasoning, or an integrated intention that plays some guiding, monitoring, and coordinating functions), this does not entail that the agent *must*, for the sake of rationality, agglomerate those intentions into a larger or compound intention. For example, a less unified deliberative background framework may still do the job even with less efficiency. Also, it is not clear that only an integrated intention can guide, monitor, and coordinate one's action in the context of competitive intentions.

commitment or effort to bring about the conjunction of *A-ing* and *B-ing* turns out to compete or even conflict with her commitments or efforts to bring about *A-ing* and *B-ing* taken separately.

Here is a case that illustrates how the conjunction of *A-ing* and *B-ing* becomes a competitive goal to *A-ing* and *B-ing* taken separately. Suppose Ben is a criminal defense lawyer, and he has two clients – Peter and Mary – who are indicted for the same crime (let us stipulate that the court allows Ben to represent both defendants). Ben realizes that the best defense strategy for each of them is to shift suspicion to the other. It gives each of them a 50% chance to be acquitted (I shall call it “strategy of implication”). However, if they do so, the jury would naturally suspect that at least one of them committed the crime and feel reluctant to find both of them not guilty. In that case, there is only a 5% chance for both of them to be acquitted. In other words, their respective acquittals are not independent of each other (if they are independent of each other, it should be 25% chance for both of them to be acquitted). Now, Ben also knows that the best strategy to get both of them acquitted is different. If both of them do not accuse each other but keep silent, the jury has a 30% chance to find both of them not guilty (I shall call it “strategy of silence”). But the downside of using the strategy of silence is that if the jury decides not to acquit both Peter and Mary the jury will find both of them guilty (70% chance). So the second strategy, even though giving a much better chance for both of them to be acquitted, lowers Peter’s or Mary’s individual chance of acquittal to 30% from 50% of the first strategy. All Ben cares about is each client’s chance of acquittal and he has no particular concern about how likely both of them will be released, therefore, Ben chooses to use the strategy of implication to defend Peter and Mary respectively.

Ben's two relevant actions are to defend Peter successfully and to defend Mary successfully.⁹ Let us use "*P*" and "*M*" to denote these two actions respectively. My case is different from the videogame case in two ways. First, it is *not* impossible for Ben to *P* and *M*. Either defense strategy has a non-zero chance of getting both defendants set free. But as far as the strategy of implication is concerned, even though it does not make *P*-ing and *M*-ing incompatible with each other, these two actions are competitive with each other in the sense that one occurrence would make it more difficult and less likely for the other to occur. So, this case depicts a situation within the scope of Agglomeration*. Second, an intention to *P* and *M* seems to be a totally heterogenous intention which is not rationally required even if Ben intends to *P* and intends to *M*. In order for Ben to intend to *P* and *M*, Ben would need to have a very different motivational set. For instance, if Peter and Mary are a couple and their best mutual interest is to be released together, Ben as their lawyer may be motivated to form a separate intention to successfully defend both of them. Also, in order for Ben to intend to get both Peter and Mary set free, he should be ready to adopt the strategy of silence which is incompatible with the strategy of implication.

In the lawyer case, there seems to be no rational pressure on Ben to form an intention to *P* and *M* if he intends to *P* and intends to *M*. In fact, there seems to be rational pressure not to have these intentions at the same time. That is because the conjunction of *P* and *M* is in a competition with *P* and *M* taken separately. If Ben is committed to realizing *P* and *M* taken separately, then he cannot commit himself at the same time to the compound goal of *P* and *M*. Although the realization of the compound goal is not incompatible with the realization of two

⁹ By 'to defend Peter successfully' and 'to defend Mary successfully,' the action involved is not the collective action of acquitting Peter (or Mary) by the jury. Instead, I mean Ben's complex action of defending Peter (or Mary) in court.

separate goals, the defense strategies that are required by commitments to these goals are in conflict. The compound goal requires the strategy of silence while the separate goals require the strategy of implication. Ben simply cannot adopt both strategies.

Therefore, we have a case in which the conjunction of *A-ing* and *B-ing* is in competition with *A-ing* and *B-ing* taken separately even though the former is logically entailed by the latter put together. To intend to *A* and *B* is partly to be practically committed to this compound goal, which might be incompatible with the commitments given to *A-ing* and *B-ing* taken separately. This shows something I suggested earlier, namely, that logical entailment or logical conjunction does not necessarily produce the kind of rational pressure suggested by Agglomeration*. Between an intention to *A* and an intention to *B* on the one side and an intention to *A* and *B* on the other, there is no “cohesion” that rationally pulls them together. As the lawyer case shows, there is a strong intuition that no such rational pressure exists. In fact, there is a strong intuition that there is rational pressure not to intend to *P* and *M* when Ben intends to *P* and intends to *M*. So, Agglomeration* is false.

Bratman may emphasize that Agglomeration* is not meant to be a categorical principle. Instead, it is qualified by a *ceteris paribus* clause. That means, the rational requirement proposed in Agglomeration* is defeasible: “there may be special circumstances in which it is rational of an agent to violate them” (Bratman 1999, pp. 32). This might well be what happens in the lawyer case. Given the fact that intending to *P* and *M* together and intending to *P* and *M* separately require different and incompatible defense strategies, it seems that Ben has a stronger reason not to intend to *P* and *M*, and this reason overrides the rational pressure to have such a compound intention. If so, it does not falsify Agglomeration* even if it is overall rational of Ben to intend to *P* and intend to *M* without intending to *P* and *M*.

But to clarify, my argument against Agglomeration* is not based on the intuition that, as a matter of fact, Ben rationally intends to *P* and intends to *M* without intending to *P* and *M*. The lawyer case may impress us in such a way and I am inclined to regard this impression as accurate. But what grounds my rejection of Agglomeration* is the intuition that there is simply no rational pressure to have the additional compound intention even if Ben has two separate intentions. These intentions represent different practical commitments to different goals. They can naturally come apart without failing to meet any rational demand. Even if the defense strategy required by intending to *P* and *M* is not the strategy of silence but something else which is not incompatible with the strategy of implication (e.g. Ben can raise the chance of a double acquittal without hurting either client's chance by simply reminding the jury of their legal obligation to consider each client's defense separately), it seems rationally permissible for Ben to omit to adopt the additional strategy. Ben simply does not care about whether Peter and Mary can be set free together and has made no commitment to this further goal. Therefore, the moral of the lawyer case is not that, with all things considered, Ben's failure to intend to *P* and *M* is rationally okay. Rather, there is no rational pressure on him to intend to *P* and *M* whatsoever.

So, I conclude that Bratman's videogame case fails to undermine SV. Until a convincing objection to SV is found, SV remains plausible and attractive. Therefore, as Adam remarks, we should treat SV as true.

3. Intentions and Events

Recall my main argument in this chapter:

- (1) A person performs an action *A* intentionally only if she intends to *A*.
- (2) Events are not things that we can intend to bring about.
- (3) Therefore, events are not things that we can intentionally bring about.

In sections 1-2, I defend the truth of (1), i.e., SV. In this section, I argue for the truth of (2). To do so, I argue that intentions cannot be directed at events. To illustrate, if we suppose a person, at some time during the morning, formed an intention to drive to the campus this afternoon, the object of the intention does not seem to be the particular event (or action) of her driving to the campus. It is difficult to single out a particular event when no event has been brought about. Let us stipulate that she does drive to the campus in the afternoon, and correspondingly, an event of her driving to the campus then occurs. Even if the event is taken to be a particular action of her driving to the campus, such a particular action, retrospectively, is not what was intended to be performed when she possessed the intention this morning. It seems true, instead, when a person forms an intention to perform an action, what she intends is always an action of a certain type specified by a description of the type. It does not matter if the specification is in details or in more general terms. The difficulty of intending to perform a particular action cannot be solved by further specifying the type of action that is intended. Even if the content of the intention is specified as “drive safely to the eastern entrance of the campus around 2 pm,” the intention is still directed at no particular action. Specification is not particularity.

The example above seems to support premise (2), especially for those cases in which the intended action is not being performed at the moment of intending (even though it will be performed). At this point, I want to introduce a distinction between two kinds of intentions: pure intentions and intentions in action. An intention in action is an intention that is being executed or acted upon. A pure intention, by contrast, is an intention that is not being executed or acted upon. Sometimes, pure intentions are also named “prior intentions” or “intentions for the future.” I avoid using these time-based names for pure intentions because it seems possible to have an

intention *for the present* without the intended action being in progress (e.g., I may intend to get out of bed right now even though I am lying still in bed).

When pure intentions are concerned, the truth of premise (2) seems to be very plausible and thus widely accepted among philosophers.¹⁰ For example, when Bratman formulates SV, he writes:

Suppose I intentionally start my car. On the Simple View it follows that

(1) I intend to start my car.

The point I want to note here is that I can have the intention reported in (1) whether or not I actually do start my car. The form of (1) is not

(2) aRb

where b is replaced by a singular term denoting an actual, particular action of my starting my car (Bratman 1999, pp. 112).

Unfortunately, Bratman does not explicate the reason why he thinks that the object of the intention is not an actual particular action. So, let us turn to Adams and Mele's (1989) discussion of the issue, which would shed some light on why pure intentions are not directed at actions as events:

S's intention to do *A* (unless it is a *de re* intention) does not represent the act token that *S* performs. It is true that *S* intends to perform *an* act token (dive #5, say) of type *A* [the type of dive called the 'Louganis']. But it is not true that *S*'s intention carries specific information about that dive – for prior to making it (after dive #4, say), there is no dive of type *A* (by *S*) to have information about. Thus, the intention represents an act type (*A*-type). The only time information about *S*'s performance of a token of act type *A* is received is *in perception* of *S*'s actual movements – of *S*'s *A*-ing. But that is not what was represented in *S*'s original intention. It could not be, for *S* had not *A*-ed (Adams & Mele 1989, pp. 516).

Their argument seems to be based on a cognitive difficulty. When an action has not yet occurred, the agent cannot have certain *specific* information about the particular action, i.e., dive #5. Of course, this does not mean that the agent has no information about dive #5 at all. At least,

¹⁰ As far as I know, Davidson is the only one who explicitly advocates that pure intentions can be directed at particular actions.

the agent knows that the dive #5 is going to be a dive of a certain type, namely, a “Louganis” type of dive. But a person can perform a dive of Louganis type in many different ways, where each of them would constitute a particular *instance* of Louganis type of dive that can be described as such. So the problem is *not* that the agent does not have any information at all about the particular dive she is going to perform. Rather, the problem is that the information available to her before the dive is done is not *specific enough* to single it out among all potential instances of the type of dive the agent will perform in the immediate future.

But if a person does not have enough specific information about a particular action, then she cannot have an intention to perform the action. The reason lies in the functional role of intentions. Adams and Mele put it this way:

Intentions have two main cognitive functions: (i) they fix the type of action that is intended, and (ii) they figure in the guidance and correction of the bodily movements that bring off an action of that type (Adams & Mele 1989, pp. 515).

We cannot use requirement (i) to rule out the possibility of intending a particular action without begging the question. But it seems highly plausible to require that an intention carry adequate and specific enough information about the intended action so that it can figure in the guidance and correction of the bodily movements that bring off the action. If a person’s mental state fails to contain such information, it necessarily fails to serve the essential function of guiding the agent to perform the action, and thus fails to be a genuine intention at all.¹¹ Therefore, since not enough specific information about the particular action is available to an agent before the action is performed, it is impossible for the agent to have a pure intention to perform a particular action.

¹¹ It is not required that an intention must guide the agent to do the intended action. But an intention must be able to do so if the agent has the will and the opportunity to execute the intention.

To illustrate why epistemic indeterminacy makes intending impossible, let us suppose that I intend to marry my future wife (suppose I am not married now and “my future wife” refers to *the* particular woman I am going to marry, even though I have no idea who that particular woman is, and therefore in some sense *I* do not refer to that particular woman). Is this a genuine intention or not? The answer may seem to be no because it seems that the identity of my future wife is not determined until I marry her. If I have already married this woman, then there is no place for such an intention. If not, the alleged intention does not carry any specific information that is required to guide me to perform the action of marrying such a woman.

However, this answer is too quick. Even if the identity of my future wife is not ontologically and semantically determined before the ceremony is completed (it is possible to have a certain accident that stops it from happening), when the ceremony is very close, I may have enough evidence to justify my belief that a particular woman (say, my fiancée) is my future wife. In this sense, the identity of my future wife is not epistemically undetermined for me. Even if I am not 100-percent certain that my fiancée is my future wife, I am certain enough to act on this information and execute my intention to marry her. This makes the supposed intention to marry my future wife more similar to another intention of mine, an intention to give an heirloom (e.g., a ring) to my future wife. When the intention was formed, the content of the intention was not epistemically determined. But since there is a way to determine, at least epistemically, the content of the intention before the intended action is performed, it can be a genuine intention after all. To put it another way, the intention comes with a plan to guide me to perform the action. The first step of the plan is to epistemically determine the content of the intention.

However, we cannot say the same thing about an alleged intention to perform a particular action in future. Suppose I allegedly intend to have my next coffee-drinking this morning. It may

well be the case that “my next coffee-drinking” refers to a particular action. But no matter how close I am to performing such an action, I cannot epistemically pick out the particular action and determine what is intended exactly. The difficulty seems to lie in a difference between the indeterminacy of my next coffee-drinking as an event and that of my future wife as a person (or an entity in general). No matter how close I am to bringing about the event of my next coffee-drinking, there are still too many or even countless different ways to do it. Even if I have entered a particular Starbucks and decided that I am going to order a venti Latte and finish it at the corner table, how I order my Latte and how I drink it still vary. I may narrow down the scope of potential instances of my next coffee-drinking, but I cannot narrow it down to just a few. Therefore, unlike my future wife as a person, I cannot even epistemically determine the identity of my next coffee-drinking until I ontologically bring it about in the world. In this sense, any alleged intention to perform a particular action that has not existed would be *inherently* undetermined. Thus, it cannot be a genuine intention at all.

Most philosophers are convinced that pure intentions cannot be directed at particular actions. The only exception, as far as I know, is Donald Davidson. Davidson has a peculiar way of understanding the concept of intentions. For him, to intend to perform an action is to judge that the action as a whole is desirable (Davidson calls this kind of judgment “all-out judgements” or “unconditional judgements”). Since he also believes that “It is only when I come to an actual action [i.e., a particular action] that it makes sense to judge it as a whole as desirable or not” (Davidson 2001, pp. 97), intentions must be about particular actions. We do not need to go into the details of Davidson’s account of intentions. For our purpose, we only need to realize that Davidson acknowledges the difficulty of intending to perform a particular action when the action has not existed. The problem, according to Davidson, is that when the intended action does not

yet exist, it seems impossible to have a demonstrative reference to the action. Thus, it seems no intention can be directed at a particular action.

However, in a later essay, Davidson proposes a solution to this difficulty:

For some reason I thought that it made sense to make an all-out judgement about a past or present action (even though the judgement might turn out to be wrong) because we could pick out a particular past or present actions, whereas there are not yet any future actions to be picked out. This now seems to me simply wrong. Future actions, like any other entities, have unique descriptions, and so can be 'picked out' like other things. It is not even the case, as I apparently thought, that future actions can't be identified through the use of indexical devices. Of course they can: 'the next time I drink a Pernod in Paris' picks out, I hope, a unique future action (Davidson 1985: pp. 197-8).

Here, Davidson tries to correct a mistake in his earlier thought that we cannot pick out a particular action in future due to lack of acquaintance and impossibility of demonstrative reference. I agree with Davidson that this might be a mistake. To see this, consider John Searle's first axiom of reference—whatever is referred to must exist—it does not say that whatever is referred to must exist *now*. In a footnote, Searle clarifies that “exist” is to be understood tenselessly and “one can refer to what has existed or what will exist as well as to what now exists” (Searle 1969, pp. 77). Davidson seems to mistakenly think in “Intending” that demonstrative reference is the only kind of reference. But here, Davidson corrects this mistake by suggesting that even if it is impossible to have a demonstrative reference to an object that does not exist in the present, if the object is going to exist in future, it might still be referred or picked out by definite descriptions such as “the next President of United States” or “my next Pernod drinking in Paris.”

Someone may argue against Davidson that the referents of these definite descriptions are indefinite or undetermined. The meanings of “the next President of United States” and “my next Pernod drinking in Paris” do not change if they are followed by “whoever he or she is” and “whenever it is.” When a person uses these descriptions, she does not have any particular person

or particular action in her mind. Therefore, it can be argued that these examples of what Donnellan (1968) calls “indefinite definite descriptions” do not refer at all until the identity of the referent becomes determined. However, it seems possible to reply that they do refer to a particular person or action. After all, only one object can turn out to satisfy the description in future. The problem is not that they do not refer to a unique referent, it is that the identity of the referent is at first unbeknownst to the person when she uses these descriptions. In other words, the referent is epistemically undetermined rather than semantically undetermined.

I do not hold a particular view on the issue whether indefinite definite descriptions can refer to a unique object or not. But for the sake of argument, let me assume they do and Davidson is right to reverse his earlier view and contend that a future action yet to occur can be picked out by the agent.¹² My objection to Davidson is rather that it remains impossible for an agent to have an intention towards a future particular action. This is because the agent does not have the sufficiently specific information about the action that is required to serve the essential cognitive and guiding functions of intentions. Even if the referent—the particular action that is going to be performed by the agent—is merely epistemically undetermined rather than semantically undetermined, the agent still lacks the needed specific information about the action. Therefore, even if the agent has a mental state representing a particular action, the mental state would necessarily fail to figure in the guidance and correction of the bodily movements that bring about the action. It cannot be an intention.

¹² One problem with this view is that maybe there will not be any particular action that fits the description. For example, maybe Davidson never get the opportunity to have his next Pernod drinking in Paris. Then, the description, if it refers, refers to something that does not exist at all. This seems to cause a difficulty for Davidson since he is still committed to Searle’s first axiom of reference. It is not just that the referent of “my next Pernod drinking in Paris” is undetermined. Rather, whether it refers at all is also undetermined.

How about intentions in actions? As Davidson remarks, the demonstrative reference to an on-going action is possible, therefore, it might be expected that intentions in action do not face the difficulty that intending particular actions in the future faces. Some other philosophers (e.g., Brandom 1994, Wilson 1989) also share this view that intentions in action are special in that they can be directed at particular actions.

In response, McDowell gives the following example and argues against the idea that intentions in action can be directed at particular actions:

[I]t may be true of one that one is crossing a street though one is not going to cross the street; perhaps halfway over one is going to be knocked down by a bus. One's intention in action may be correctly specifiable by saying "I am crossing the street," even if one is not going to make it to the other side. So there may never be a relevant particular action of the sort that figures in the specification of one's intention in action – for instance a street-crossing. There may be no action for it to be directed *de re* at.

Reality comes to contain, as it were, a particular action of the type describable as "my crossing the street" only when it gets to be true to say "I have crossed the street." But by then the time for the relevant intention of crossing the street to be an intention in action is over (McDowell 2011, pp. 7).

According to this passage, it is impossible to have an intention in action to perform a particular action for the very same reason why a pure intention to perform a particular action is impossible. Even if the action is in progress, no such action has come into existence before the action is completed. So, strictly speaking, the action does not exist yet. Therefore, the argument raised in my discussion of pure intentions starts to kick in.

However, someone may object that McDowell's argument is based on a false dilemma: It is false that either an action has come into existence or not. A third possibility should be considered, namely, that an action is *coming into existence* when it is in the process of being performed. Since an intention in action is defined as an intention that is being executed or acted upon, this third possibility seems to fit the definition better than the two options in the dilemma. To say that an action is coming into existence is partly to say that only a part or stretch of it has

come into existence. Opponents of McDowell can argue, however, that pointing to the earlier stretch of the action seems to be enough to pick out this particular action. No matter how the action develops later, it is the one that is referred by the agent by pointing to that stretch of the action. Even if the person eventually fails to cross the street and therefore it is incorrect to say that this person crossed the street, it does not entail that there is no relevant particular action at all. After all, the person does something that could plausibly be described as “street-crossing.” Admittedly, it turned out to be a failed street-crossing, but people can argue that a failed street-crossing is still a street-crossing.

However, when intentions in action are concerned, the person who allegedly intends to perform a particular action is in the middle of performing that action. At such moments, part of the resultant action has come into existence whereas part of it has not. Our discussion needs to accommodate this complexity by distinguishing between those two parts of the action. As McDowell says:

When an intention of crossing the street is an intention in action, it has behind it, at any point, a stretch of action with the particularity that comes with completion, signalled by the truth of a claim in the perfect tense: one has moved to here (to put it in a way that requires imagining oneself engaged in the action). But an intention that is still in action also looks forward to a stretch of action that may never emerge into reality as a particular achievement: moving from here to the other side (McDowell 2011, pp. 7).

Let us call these two stretches of the action “prior stretch” and “future stretch” respectively. We realize that to intend to perform an on-going action does not contain or entail an intention to bring about the prior stretch of the action that has already come into existence. It is pointless and does not make sense to intend to bring something into existence if that thing has already existed. Rather, to intend to perform an on-going action must be to have an intention about the future stretch of the action. It is either to intend the future stretch of the action to have certain general (desirable) feature so that the resultant action as a whole has certain (desirable)

feature, or to intend the future stretch to unfold in certain particular way so that the resultant action as a whole turns out to be a certain particular action.

So, in McDowell's example, either the person intends the future stretch of action to unfold successfully so that his street-crossing as a whole is successful, or he intends the future stretch of action to unfold in a particular way so that it results in the occurrence of a particular action of street-crossing in the end. Obviously, it is the latter intention that counts as an intention to perform a particular action. Now, it should become clear that the problem raised in my discussion of pure intentions does apply to intentions in action as well. The problem is not that the agent cannot refer to or mention the future stretch of action in question. The real problem is rather that the agent does not have enough specific information about the future stretch until it has already been realized. If so, the alleged intention would necessarily fail to figure in the guidance and correction of the bodily movements, thus cannot play a role in bringing about the action. Even if a particular action, which fits the description, will come into existence later, it is explanatorily irrelevant to the alleged intention. To think otherwise would lead to the absurd idea that my alleged intention to do the next thing I am going to do (let us assume "the next thing I am going to do" refers to a particular action, i.e., the next action I am going to perform) can guide and explain the next thing I end up doing.

So, in this section, I have defended the view that an intention, no matter whether it is a pure intention or an intention in action, cannot be directed at an action if this action is taken to be an event (thus, a particular action).

4. An Objection to Consider

Before moving on, let me summarize the discussion above. In section 1 and 2, I defend SV: A person performs an action A intentionally only if he intends to A. In section 3, I argue that

an intention cannot be directed at an event, and therefore, even if there are particular actions, their performance cannot be intended by us.

In this section, I anticipate and give a response to a potential objection to my argument. In my view, most philosophers would agree that a person cannot intend to bring about an event or intend to perform a particular action, even though they may have slightly different reasons to back up this conclusion. In contrast, SV may still appear controversial to some philosophers. The question is not whether SV is false. Rather, it can be challenged that SV, as I formulate it, does not accommodate the possibility that an action—as an event—can be performed intentionally. It may seem to be the case especially considering that it is impossible to intend to perform a particular action. In other words, given most people's belief that intentions are always directed at actions of a certain type or act-universals as I tend to call them, SV needs to be formulated accordingly. To reflect on this line of thought, let us consider the following principle as a revised formulation of the Simple View:

SV*: A person performs a particular action *A* of type *B* intentionally only if he intends to perform an action of type *B*.

By saying that a person intends to perform an action of type *B*, it does not mean that the intention is directed at any particular action (let alone the particular action *A*). It only means that the person intends to do something, *whatever it is*, to bring about the exemplification of *B* as an act-universal. In short, it is an intention directed at an act-universal rather than a particular action. According to SV*, what is required for a particular action to be an intentional action is not an intention to perform that *very* action. Instead, what is required is an intention to perform an action of a certain type (without specifying which particular action) when the particular action in question is one of such type.

Obviously, SV* is different from SV. SV requires of the agent something she cannot do—intending to perform a particular action—in order to perform a particular action intentionally. SV* does not. My opponents may argue that, although the Simple View seems to be true, SV is not. SV arguably proposes a too demanding necessary condition for performing a particular action intentionally. If so, the Simple View might be best captured by SV* rather than SV. That is, SV* might seem to appropriately state exactly what is required of a person to perform a particular action of a certain type intentionally. If SV* as a revised and better formulation of the Simple View, it would be possible for an action to be performed intentionally even if it is an event. And my whole argument in this chapter seems to fail.

How do I respond to this objection that SV* best captures the Simple View? I will argue that SV best captures the relevant necessary condition for intentional actions in the Simple View. So, there is no good reason to revise SV into SV*. To see SV's plausibility, I need to bring the concept of trying into the picture. Here, I defend a principle about the relation between trying and intending: trying to *A* entails intending to *A*. Let us follow Mele in calling it the "Strong Intention Thesis" or SIT for short.

(SIT) A person tries to *A* only if she intends to *A*.

Hugh McCann (1987), proposes an argument for SIT. I think his argument is basically sound and convincing. But I shall discuss and defend McCann's argument in the light of two objections. One is raised by Mele (1989) based on the distinction between intentions and intentions to try. The other is based on Bratman's distinction between a strong sense of acting with an intention and a weak sense of it. I begin with Mele's summary of McCann's argument followed by Mele's objection:

(1) An agent tries to *A* only if she does something *B*, that she "conceives of as a means" to *A*-ing, and "does it because [she] so conceives it." (2) To *B* because one conceives of

B-ing as a means to *A*-ing is to *B* for the purpose of *A*-ing. (3) To *B* for the purpose of *A*-ing is to *B* with the intention of *A*-ing. Therefore, (4) an agent tries to *A* only if she does something *B*, with the intention of *A*-ing and, hence, only if she intends to *A*.

McCann anticipates the objection that someone who *B*-s for the purpose of *A*-ing may hope – but does not intend – to *A*, if she does not believe that she (probably) will succeed in *A*-ing. He offers two replies (p. 204). First, the objection has the untoward result that there is no intention with which the agent *B*-ed. Second, “when we seek to fulfill our hopes, we do so precisely by acting with the intention of fulfilling them – with the intention of doing what we hope to do.

One might think that both replies can be countered with the same contention, namely, that the pertinent intention in both cases is an intention to *try* to *A*” (Mele 1989, pp. 101).

McCann’s argument is meant to show that trying to *A* entails or simply is acting with a certain intention, namely, the intention of *A*-ing. Mele has no objection about the first part.

However, Mele points out that the intention with which a person acts when she tries to *A* is not necessarily the intention of *A*-ing. It might well be the intention to try to *A*, according to Mele.

To be precise, Mele does not positively argue for this view: trying to *A* entails intending to try to *A*. I think Mele’s positive view is best captured by this: trying to *A* entails either intending to *A* or intending to try to *A*. For Mele, trying to *A* entails a certain intention, and in those cases in which the intention in question is not an intention to *A*, it is an intention to try to *A*. It is not clear that Mele wants to rule out the possibility that there are cases in which a person tries to *A* with the intention of *A*-ing associated with no other intentions (including an intention to try to *A*).

McCann’s reply to Mele’s challenge is that mere intending to try to *A* is “vacuous” unless it is associated with a more substantial intention, namely, the intention to *A*. McCann gives his reasons as follows:

The reason for this is the fact mentioned earlier, that trying is not a species of action. Unlike the case with acts like moving a finger or pulling a trigger, which consist in bringing about a specific sort of change in the world, there is no one sort of change that counts as a “try” (McCann 1987, pp. 203).

I find McCann's reason unconvincing because he assumes an understanding of "change" that seems too narrow. Granted, there is a sense of change that makes it sound right for someone to say "Nothing has been changed by you" or "You failed to make any change" to another person who tried to do something (e.g., to save a patient) and failed. But saying this suggests that change must be or contain something desirable or valuable (or, the opposite). So, "You failed to make any change" actually means you failed to make any *real* change, namely, any change with valuable or desirable attributes. But strictly speaking, even a change without any value is still a change. So, if someone tried to do something but failed, she at least tried and the occurrence of her trying was the change she brought about in the world. Therefore, I see no reason to deny that a person can merely intend to try. Also, it seems possible to value the effort of trying *per se* so that it makes better sense to intend to try.

Mele rejects McCann's response by appeal to the following example between two college basketball players standing at mid-court:

Alice: "Do you intend to make a basket from *here*?" Betty: "No. But I do intend to try."
Betty's response strikes me as quite unproblematic. But, for McCann, Betty's avowed intention is not an intention at all unless it is "supplemented by a more substantive intention" – specifically, an intention to sink the shot (Mele 1989, pp. 102).

I agree with Mele that Betty can rationally intend to try while not intending to sink the shot. And I reject McCann's claim that an intention to try to *A* without an intention to *A* is vacuous. However, I still believe that McCann's argument is basically to the point and SIT is true accordingly. Therefore, I need to respond to Mele's suggestion that trying to *A* may well be only acting with the intention to try to *A* (at least in some cases). To do so, I shall correct a commonly accepted but mistaken view about trying. Many philosophers believe that it is possible for a person to rationally try to do something even if she believes doing so is impossible. Consider a case in which Josh *apparently* tries to move the heavy log by pushing the log as hard

as he can even though he believes he cannot move the log since it is too heavy. Josh's motivation may be that he wants to show his friend that he cannot move the log.

Many philosophers (Mele is one of them) would say that Josh is making genuine effort to move the log so that he is trying to move it. I argue that if Josh is not being irrational (Josh is not neglecting or ignoring the relevant belief, and Josh has good reason to do what he does in the case), Josh cannot try to move the log when he believes the log is too heavy for him to move. Josh is surely doing something, namely, pushing his hands on the side of the log and exerting a force towards it as best as he can. But this does not necessarily mean that he is genuinely trying to move the log. Instead, Josh is merely doing something that is usually taken to be a genuine endeavor to move the log in normal circumstances (e.g., when he does not believe the log is too heavy for him to move). If Josh were really trying to move the log, given his belief that it is impossible to move it merely by pushing it, he would have used a crane or some other instrument instead of using a method he knows is not going to work.

To motivate my analysis of the log example, let us imagine that rather than pushing the log, Josh pushes the Eiffel Tower.¹³ This time, Josh wants to demonstrate that the Eiffel Tower is too heavy for him to move by doing the same thing as he does in the log case. Now, is Josh genuinely trying to move the Eiffel Tower by pushing on it with all of his weight? It is absurd to answer yes. This method of moving things (i.e., pushing the thing), no matter how hard Josh tries it, is clearly not going to work in this case. So, when Josh chooses this method, he is probably trying to do something else. For example, it could be that he is just trying to show that he cannot move the Eiffel Tower (by merely pushing it).

¹³ In Adams (2007), a similar case is used to show that one cannot rationally try to do what one believes to be impossible.

If someone is still not convinced by the example above, let us imagine a different case. Suppose you are a lab researcher and to open the lab door you need a password. One day, you want to go to the lab as usual and you enter the password but it does not work. You keep entering the password for a few times but the door remains locked. Then, you are told by one of your colleagues that the password has been changed by someone else and no one here has been given the new password. Now, with this new information, you naturally form the belief that it is impossible for you to open the door by entering the old password. However, if you keep entering the old password, is this effort a genuine (rational) endeavor to get the door opened? Of course not!

What is the difference between the log case and the password case that makes the former *seem* to be a genuine case of knowingly and rationally trying something that is impossible? The only difference seems to be that, in the password case, it is highly *unlikely* to mistake the act of entering a useless password for a genuine endeavor to open the door while it is much easier to make such a confusion in the case of pushing a heavy object, especially when the heavy object is not beyond a person's physical capability by too much. In other words, pushing things has been taken to be the typical and appropriate method of moving things by most people, which may lead people to always conceive it as the right way to move things even if it is not sometimes.

There are two lessons that we can learn here. First, doing something that is usually taken to be a genuine endeavor to A is not always a genuine endeavor to A. This observation corroborates my objection to the videogame case. No doubt, I am shooting at target 1 and my shooting is guided by my perception of target 1. This is usually taken to be a genuine endeavor to hit target 1. But it is not necessarily true that I am really trying to hit target 1. It might be the case that I am only trying to hit either target 1 or target 2, and this trying includes my aiming and

shooting at target 1 (and target 2). Suppose a businessman donates a large amount of money to the Clinton campaign. This is usually taken to be a genuine endeavor to get Clinton elected. However, if he is also donating an equal amount of money to the Trump campaign at the same time, can he still honestly claim that he is trying to get Clinton elected? Or is what he is trying just to make sure one of them gets elected?

Second, and more importantly, it strongly suggests that my trying to *A* is subject to the rational constraint of strong consistency as well.¹⁴ Let us call it “Strong Consistency of Trying” (or SCT for short):

(SCT) It is rational of a person to try to *A* only if *A*-ing is consistent with her beliefs.

In the password case, as soon as you realize that the old password is expired, even if nothing else is changed in the case, even if you are doing exactly the same thing as you did a second earlier, you cannot rationally and honestly claim that you are still trying to open the door by typing in the old password.

How to explain the truth of SCT that trying to *A* is also subject to the strong consistency constraint as intending to *A* does? In specific, what can explain why trying to *A* (for the sake of rationality) entails *A*-ing is consistent with one’s beliefs. Since trying to *A* is acting with a certain intention (both Mele and McCann agree with this), it seems most natural to conclude that trying to *A* must contain or involve the particular intention that also (for the sake of rationality) entails the consistency between *A*-ing and one’s beliefs. McCann has a straightforward answer here. Since trying to *A* is acting with the intention to *A*, and since the intention to *A* (for the sake of rationality) entails *A*-ing is consistent with one’s beliefs, trying to *A* must be subject to this consistency constraint. This explains why SCT is true.

¹⁴ Strictly speaking, it is a counterpart of the strong consistency of intentions and beliefs.

The difficulty for Mele is that he has no good explanation for the truth of SCT. If trying to *A* can be acting only with the intention to try to *A*, then it only entails that it is rational of a person to try to *A* only if trying to *A* is consistent with her beliefs. This does not further entail SCT because it is possible for trying to *A* to be consistent with one's belief when *A*-ing is not. For example, such a person may mistakenly believe that trying to *A* does not rationally require that *A*-ing must be consistent with her beliefs. If so, she may find trying to *A* to be possible given her beliefs without finding *A*-ing to be possible at all. Therefore, if trying to *A* is acting with the intention to try to *A* as Mele maintains, there is no explanation for why trying to *A* is subject to the consistency constraint that SCT manifests.

Now, I consider another potential objection to McCann's argument. Bratman also rejects SIT. This is motivated by his belief that trying to *A* is not subject to the same rational constraints as intentions are. Otherwise, Bratman cannot maintain that I try to hit target 1 but do not intend to hit target 1 based on rational constraints such as Strong Consistency (of intentions) and Agglomeration (or Agglomeration*). So, Bratman makes it clear that, unlike intention, there is no such rational constraint on trying or endeavoring. To defend this claim, Bratman proposes two senses of "acting with an intention." Using the example of opening the curtain with the intention of getting more light, he writes:

I think we need to distinguish two readings of (2) [namely, I open the curtains with the intention of getting more light]. On the *strong* reading, (2) does indeed entail that I intend to get more light, perhaps as part of a larger plan of action. On the *weak* reading, (2) only entails that I act *in order to* get more light; it does not entail (though it does not preclude) that, strictly speaking, I *intend* to get more light" (Bratman 1999, pp. 129).

On Bratman's account, the weak sense of acting with an intention of *A*-ing is *equivalent* to trying to *A* or endeavoring to *A*. And he uses his videogame case to defend the distinction between these two readings: "one lesson of the video-games example is that, in this endeavoring sense, I may act with the intention of *A*-ing and yet not intend to *A*" (Bratman 1999, pp. 129).

The implication of the distinction is that, although the strong sense of acting with an intention of *A*-ing is subject to the rational constraints on intention to *A* (since it entails the corresponding intention), it is *possible* that the weak sense of acting with an intention of *A*-ing (i.e., trying to *A* or endeavoring to *A*) is not. In fact, Bratman claims that trying to *A* is subject to neither Strong Consistency nor Agglomerativity (or Agglomerativity*).

If Bratman is correct, then even if trying to *A* is identical to acting with the intention to *A*, it does not follow that trying to *A* contains or entails the intention to *A*. I have two responses to Bratman. First, Bratman's distinction between two senses of acting with an intention is solely justified by his analysis of the videogame case. Since I have already undermined his analysis of the videogame case, he cannot use the videogame case to support the distinction anymore.

My second objection is based on my disagreement with Bratman concerning how to best defend the strong consistency of intentions and beliefs in face of apparent counterexamples such as that I can rationally try to move the heavy log even though I believe I cannot. Bratman uses this case to corroborate his distinction between two senses of acting with an intention. According to Bratman, this case is an *apparent* counterexample to the strong consistency because the following inference may be made:

- (1) I try to move the log.
- (2) Therefore, I act with the intention of moving the log.
- (3) Therefore, I intend to move the log.

So, the case seems to show that I can rationally intend to move the log even though I believe that I cannot. If so, it falsifies the strong consistency constraint on intention. Bratman's reply to this alleged counterexample is that the inference "involves an equivocation on 'with an intention'" (Bratman 1999, pp. 133). A weak sense of acting with an intention is required for the

inference from (1) to (2) to be valid, while a strong sense is required for the inference from (2) to (3) to be valid. So, this inference commits the fallacy of equivocation, thus is invalid.

But this is not a plausible way to explain away the apparent difficulty. Bratman's analysis assumes something I have shown to be false. That is, Bratman shares Mele's belief that it is possible for me to rationally try to move the heavy log even though I believe it is too heavy for me to move. I have shown that trying, like intentions, is subject to strong consistency constraint as well. Therefore, premise (1) is false. There is no apparent counterexample to the strong consistency to start. So, if Bratman wants to use the explanatory power of his distinction between two senses of acting with an intention as a reason to support his distinction, he cannot.

Give that SIT is defended, what is the connection between SIT and SV? It is a principle that is sometimes called "Ubiquity":

Ubiquity: A person performs an action *A* intentionally only if she tries to *A*.

Ubiquity has been defended by some philosophers (e.g. Hornsby 1980, 2010; Grice 1989). The reason why they try to defend it is not that Ubiquity is intuitively implausible, but because it faces the famous challenge raised by Wittgenstein: "When I raise my arm I do not usually *try* to raise it" (*Philosophical Investigations*, §622). Wittgenstein seems to contend that a person does not try to do something if doing it does not involve some sort of difficulty or any possibility of failure. This challenge is based on the observation that we would not *say* that a person tried to raise his arm if his ability to do so was not being questioned.

Paul Grice responds to Wittgenstein's challenge by first accepting that we do not speak of trying if the agent's ability or success is not in question. However, that is not because it would be false to say so, but because saying so would be misleading – though (trivially) true (Grice 1989, chapters 1-3). In other words, Grice believes that to *use* the word "try" is appropriate only when

the possibility of failure is real or is taken to be real. But it does not follow that not all intentional actions require trying. As Hornsby puts it:

[W]hen an agent is intentionally Φ -ing, doubts or difficulties about her succeeding do not bear on the question of whether she is trying to Φ ; they bear rather upon whether there could be any point in thinking of her as trying to Φ , and upon whether it would be appropriate to say that she is trying to Φ (Hornsby 2010, pp. 19).

To see this, consider an analogy given by Schroeder (2001). An entry in the log “Today the captain was sober” insinuates, but does not imply, that normally the captain is drunk. If the captain is not a drunkard, the entry is extremely misleading but nonetheless true.

If Grice is correct in denying any semantic significance to Wittgenstein’s observation that we do not think of or speak of trying in those cases that do not involve doubts or difficulties, there is no real challenge to Ubiquity. In response to Grice, Schroeder (2001) tries to explain why it is not merely a matter of pragmatic insinuation that we do not think of or speak of trying in cases without doubts or difficulties. “Given that p, it is found inappropriate to say that q,” Schroeder asks, “How are we to settle whether this is a matter of (A) *semantic presupposition* ... or whether it’s merely a matter of (B) *pragmatic insinuation*?” (emphasis in the original)

Schroeder gives two criteria for determining whether it is a (B)-type case:

For one thing, in (B)-type cases the utterance in question would not be inappropriate under all circumstances. E.g. if the question is ‘who of the crew is sober at the moment?’, then saying ‘The captain is sober &c.’ would carry no suggestion that this was an exceptional occurrence. Again, if the question is ‘That book costs five pounds. Have you got that amount on you?’, then one can answer ‘I have five pounds on me’, without insinuation that that is all one has. ... could there be a situation in which this utterance would not be taken to indicate the possibility (or acknowledge the actuality) of failure? I think not. ... For another thing, if Grice’s construal, and the ubiquity thesis, were correct, the answer to the question whether Sam’s emptying his glass involved trying, should be a confident ‘Of course!’, backed up, if necessary, by a trivially analytic statement: ‘If you’ve done something, then you’ve tried to do it.’ But that’s exactly where the analogy fails: this latter statement isn’t analytic, let alone trivially analytic. It’s a controversial philosophical thesis (Schroeder 2001, pp. 271-8).

Schroeder's second reason lacks force. Even if something is an analytic statement, it does not entail that its truth cannot be controversial at the same time. For instance, some philosophers argue that it is conceptually (therefore analytically) true that moral responsibility requires causal indeterminism. It would be absurd to object to those philosophers by arguing that, if it were an analytic truth it would not even be debated for so long.

Schroeder's first reason makes better sense but still fails. That is because there are cases in which it would be appropriate to speak of trying without insinuating doubts or difficulties. As Hornsby says:

In order to see that 'try to' applies in many cases where we wouldn't think to apply it, it helps to see that learning what a person is trying to do is one way to come to know something about the person's reasons for acting. Often enough, someone who is asked 'Why are you Φ -ing?', if she is Φ -ing as a means of Ψ -ing may say any of 'I want to Ψ ' or 'I intend to Ψ ' or 'I am trying to Ψ ' or 'I am Ψ -ing' (Thompson 2008, 97–100). So, when one is told that some agent is trying to Ψ , one may be alerted, not to something that requires special effort or in which the agent is unlikely to succeed, but to something which the agent has a reason to do, yet is not visibly doing now (Hornsby 2010, pp. 19).

In the light of Hornsby's remark that trying may be applied without bearing the pragmatic insinuation of doubts or difficulties, Schroeder's first reason to undermine Grice's response fails. Therefore, Grice's response based on the distinction between semantic presupposition and pragmatic insinuation stands as a good way to explain away Wittgenstein's observation that we do not think of or speak of trying when no doubt or difficulty or possibility of failure is either real or taken to be real. The discussion above, although short of a conclusive defense of Ubiquity, shows that Ubiquity remains a plausible principle that can be used to bridge the gap between SIT and SV (SIT and Ubiquity entail SV).

To object to the truth of SV (and make SV* look like a good alternative), opponents of SV have two options. Either they argue that Ubiquity is not true or at least that Ubiquity should be revised to the following thesis:

Ubiquity*: A person performs a particular action *A* of type *B* intentionally only if she tries to perform an action of type *B*.

Or, they argue that the entailment between trying and intending is not best expressed by SIT but rather by a different principle:

SIT*: A person tries to perform a particular action *A* of type *B* only if she intends to perform an action of type *B*.

I doubt the second option will be chosen. The reason is that it seems impossible that a person can ever try to perform a particular action. In section 3, when I explained why it is impossible for a person to intend to perform a particular action, the difficulty lies in the lack of certain specific information about the particular action. Without certain specific information, the agent's performance of the action is not guided by the content of the intention. If she performs it, it is not due to the guidance function of the intention. From another perspective, the agent does not know how to bring about the existence of a particular action. For example, if a person intends to drive to the campus around 2 pm, she knows how to do it (if she has a car, knows how to drive, and knows the route). She has a plan or procedure, more or less elaborated, to implement in order to drive to the campus around a specific time. The plan will reliably guide her to perform an action of driving to the campus around 2 pm. By following the plan, she does what she can to drive to the campus around 2 pm. In other words, she tries to drive to the campus around 2 pm. In contrast, due to the lack of specific information, there is no plan or procedure which can reliably bring about an event of driving to the campus around 2 pm, even though there is something she can do to bring about a particular action of driving to the campus around 2 pm. Therefore, she cannot try to perform that particular action even if she eventually brings it about. So, I set aside the option that the critic will try to revise SIT into SIT*.

I now focus on the critics' attempt to revise Ubiquity into Ubiquity*. By doing so, they can maintain that SV* is true and SV is not. I argue, however, that there is no good reason to revise Ubiquity into Ubiquity*. On the face of it, Ubiquity is plausible in itself and seems to best capture the point of the ubiquity of trying in intentional actions. If a person wants to perform a particular action intentionally, it seems reasonable to demand that she tries to perform that particular action rather than merely trying to perform an action of the type of which the particular action is one of its potential instances. This suggests that, without any good reason, the revision of Ubiquity into Ubiquity* is motivated solely by the attempt to justify the revision of SV into SV*.

In response to my claim that there is no independent motivation for the revision of Ubiquity into Ubiquity*, let us consider some potential reasons that might be raised to support the revision. Suppose a person tries to perform an action of a certain type, say, driving to the campus. As a result, a particular action of driving to the campus is performed. In order to determine whether the particular action is intentional or not, here are some seemingly relevant features:

- (1) The person's trying to perform an action of driving to the campus causes her to perform the particular action of driving to the campus.
- (2) The person's trying to perform an action of driving to the campus substantially contributes to the occurrence of the particular action of driving to the campus.
- (3) The person's trying to perform an action of driving to the campus is the best she can do to contribute to the occurrence of the particular action of driving to the campus.
- (4) The person is most likely not surprised by her performing the particular action of driving to the campus since she tries to perform an action of driving to the campus.

(5) The person's performance of the particular action of driving to the campus is explained, at least partly, by her trying to drive to the campus.

(1)-(5) seems to constitute a fair and plausible assessment of the case. Let us take them to be true, for the sake of argument. Based on this, some may argue that Ubiquity* is better than Ubiquity in capturing the relevant necessary condition for intentional actions because, first, (1)-(5) seem to be adequate for intentionally performing the particular action of driving to the campus; and, second, (1)-(5) can be satisfied without the agent having tried to perform that very particular action. Therefore, Ubiquity seems to propose a too strong condition for intentional action regarding trying.

In reply, I dispute the first point. That is, I argue, even if (1)-(5) are true, the person could still fail to satisfy the necessary condition for intentional action regarding trying. Consider the following example. Suppose there is a black box which has an indefinite number of balls in it (like a lottery box). A person intends to reach her hand into the black box and draw one ball out of the box. She does what she intends, and therefore she tries to do it. She does not try to draw out any particular ball in the black box. Even if she wants to do it, she does not have the information about the particular ball (regarding its shape or its location in the box) for her to try to do it. My question is: If she ends up drawing a particular ball, ball #37, can we sensibly say that she draws out ball #37 intentionally? Apparently, the answer is no. She does not intentionally draw out ball #37. What she does intentionally is to draw out some ball or other, since, other things being equal, drawing out one ball out of the box, though not any particular ball, is what she tries to do.¹⁵

¹⁵ Someone may object that this is not a good analogous case. It would be false to argue that drawing out ball #37 is to drawing a ball out of the black box as performing the particular action

The case of ball-drawing shares those features about the case of driving to the campus.

They are:

(1') The person's trying to draw out a ball causes her to draw out ball #37.

(2') The person's trying to draw out a ball substantially contributes to the occurrence of her drawing out ball #37.

(3') The person's trying to draw out a ball is the best she can do to contribute to the occurrence of her drawing out ball #37.

(4') The person is most likely not surprised by her drawing out ball #37 since she tries to draw out a ball.

(5') The person's drawing out ball #37 is explained, at least partly, by her trying to draw out a ball.

This shows that even if a person satisfies the features listed above, it may well be the case that she still fails to satisfy the necessary condition of intentional action regarding trying.

Therefore, (1)-(5) is not adequate for intentionally performing the particular action regarding trying. As long as the person's trying is not directed at the action *A*, even if she tries to perform an action of a certain type of which *A* is an instance (and as a result, features such as (1)-(5) are instantiated), she still does not perform *A* intentionally.

of driving to the campus is to performing an action of driving to the campus. Unlike the particular action of driving to the campus, drawing out ball #37 is actually a type of action, even though it is more specified in contrast to drawing out a ball out of the black box. There could be different particular instances of drawing out ball #37. In this sense, the example is not a perfect analogy. Hence, it cannot be used to dispute the revision of Ubiquity into Ubiquity* directly. I agree that the analogy is not perfect. However, as my argument unfolds, the reader should see that, for my argument to work, it does not require a perfect analogy.

Having argued that there is no good reason to prefer Ubiquity* over Ubiquity, I finish by giving a reason why Ubiquity* should *not* be preferred over Ubiquity. If what is required, regarding trying, for performing a particular action *A* of type *B* intentionally is merely trying to perform an action of type *B*, what if we conceive an agent who can try to perform particular actions. Maybe God, due to the foresight and all the powers he possesses, can intend to and try to perform a particular action *A* he is about to perform.¹⁶ Now, compare God's performance of a particular action to an ordinary person's performance of a particular action. The contribution made by God to the occurrence of his action seems more substantial than that of the ordinary person. Since God can try to perform the particular action in question, God has more control over its occurrence. Also, God's trying is much more reliable in bringing about *that* particular action. In contrast, even if an ordinary person's trying causes the occurrence of the particular action she does, the occurrence is mostly due to luck and factors beyond the person's control. The person does not have a reliable plan and method to bring about the particular action in question, which is exactly why she cannot try to perform it.

With all these points in mind, how can we have the conceptual space to account for the intentionality of these two different particular actions? Since intentionality is a matter of assessing an agent's contribution to her action, we should expect different assessments of God's action and of the ordinary person's action in terms of intentionality. But the best we can say about God's performance seems to be: God performs the particular action intentionally. In other words, God is linked to his action directly through the relation of intentional acting. If so, it

¹⁶ So, in this chapter, I do not reject the conceptual or even the metaphysical possibility that an agent can intend to perform particular actions. The difficulty I raised in section 3 for intending to perform particular actions applies only to finite agents such as human beings who have no specific enough information about particular actions to guide their actions. So, this is physical and psychological incapability that is caused by the lack of specific information.

would be unacceptable to say the same thing about the ordinary person's action. It would fail to represent the huge difference between God's action and an ordinary person's action if we report that the person performs the particular action intentionally as well. To solve the difficulty, we should simply adopt Ubiquity in requiring of the agent a trying to perform the very action that is allegedly done by her intentionally. If an agent fails to satisfy this condition, she does not perform the action intentionally.

In conclusion, I have made two points about the revision of Ubiquity into Ubiquity*. First, there is no good reason to prefer Ubiquity* over Ubiquity. Second, Ubiquity is better than Ubiquity* in accommodating the conceptual or metaphysical possibility that some infinite agent (e.g., God) can intend to perform particular actions. Therefore, we should reject the revision of Ubiquity into Ubiquity*. Since such a revision is the only plausible way to promote the revision of SV into SV*, the latter revision is not justified as well. SV stands as the most plausible expression of the Simple View. If so, the objection I consider in this section fails to undermine my main argument in this chapter.

In this chapter, based on the intention-based argument, I have thus shown that events are not things that we can intentionally bring about. This conclusion, paired with the premise I take for granted in the introduction of this dissertation—we can perform actions intentionally—would constitute an argument against the view that actions are events.

CHAPTER 3: THE ABILITY-BASED ARGUMENT

In this chapter, I will argue that events are not things that we can have abilities to bring about. Therefore, even if (some) actions are events and it makes sense to speak of particular actions, we could not have abilities to perform those actions. But since I take for granted that we can have abilities to perform actions (see Introduction of this dissertation), actions are not events.

This chapter has four parts. In section 1, I will introduce the preliminary distinction between being *simply* able to A and being able to A *intentionally*. I argue that being able to A intentionally is the appropriate kind of abilities for my investigation. However, it is disputed what the required success rate of A-ing is for a person to have the ability to A intentionally. In sections 2-3, I discuss and revise the notion of minimal ability to elaborate a weaker form of I-ability in terms of the required success rate. Minimal ability would help us identify a necessary condition of the weaker form of being able to A intentionally. In section 4, based on the necessary condition elicited in sections 2-3, I provide a timing argument for the conclusion that we cannot have even minimal abilities—i.e., a weaker form of I-abilities—to bring about events. So, we cannot have I-abilities to bring about events.

1. Ability to Act Intentionally

As stated in the Introduction, I take for granted that we can have abilities to perform actions. However, it is not clear what kind of abilities is most appropriate when this statement is discussed in the investigation about whether events are things that we can have abilities to bring about or not. A starting point is a distinction raised by Mele (2003) between being simply able to A and being able to A intentionally.

There is an ordinary sense of ‘able’ according to which agents are able to do whatever they do. In this sense of ‘able’, an agent’s having A-ed at a time is conceptually sufficient for her having been able to A then. If Ann backed her car into mine, she was able to do

that, in this sense. That is so whether she intentionally or accidentally backed her car into mine. ... It can also be said that there is a kind of ability about which claims such as these are true. I call it *simple ability*. I have not claimed that simple ability to *A* is found *only* in cases in which agents *A*. Rather, my claim is that an agent's *A*-ing at a time is sufficient for his having the simple ability to *A* at that time. ... Being simply able to *A* is distinguishable from being able to *A intentionally*. It is controversial how much control an agent who *A*-s must have over her *A*-ing in order to *A intentionally*. Even so, there are clear illustrations of a difference between control that is appropriate for intentional action and control that falls short. Ann has enough control over her body and dice to roll a die intentionally, but, like any normal human being, she lacks control over dice needed for rolling a six intentionally with a single toss of a fair die. Therefore, although she is able to roll a six with a single toss of a fair die, she is not able to do that intentionally" (Mele 2003a, pp. 447-8).

Let me adopt Mele's abbreviation which abbreviates "simple ability to *A*" as "S-ability" and "ability to *A intentionally*" as "*I*-ability." Also, corresponding "able" expressions are "S-able" and "*I*-able." Now, a few observations. Mele's illustration of the distinction between "S-ability" and "*I*-ability" implies that *I*-ability is stronger than S-ability: "being able to *A intentionally* entails having a simple ability to *A* and the converse is false" (Mele 2003a, pp. 448). So, it is perhaps true that human beings can have *S*-abilities to bring about events so that, if actions are events, we might be *S*-able to perform them. But that does not entail that we have *I*-abilities to bring about events—a kind of abilities stronger than *S*-abilities.

For my investigation, *S*-ability is too weak. Although it is easy to support my assumption that we can have *S*-abilities to perform actions, it is also true that events *are* things that we can have *S*-abilities to bring about. It is an uncontroversial fact that we brought about events when acting, leaving it open whether those events are actions or whether those events are brought about intentionally. If so, according to Mele, we have *S*-abilities to bring about events. This weak sense of abilities would render my crucial premise false.

It is more promising to argue that events are *not* things that we can have *I*-abilities to bring about. But, is it still true that we can have *I*-abilities to perform actions? In other words, can we have abilities to do actions intentionally? I think yes. The basic idea here is that the

phenomenon of acting intentionally is the most essential feature of human agency (e.g., Anscombe 2000, Davidson 2001, essay 3). If it were the case that we cannot have abilities to perform actions intentionally, it would raise the question of whether we can have abilities to perform actions at all. It does not seem a sensible stance that human beings are only able to perform actions unintentionally or non-intentionally. So, I will assume that *I*-abilities is the appropriate notion of abilities for the purpose of my argument.

But this raises a worry that the force of the ability-based argument might not be independent of the intention-based argument proposed in chapter 2. As a reminder, this is the outline of the intention-based argument:

- (1) A person performs an action *A* intentionally only if she intends to *A*.
- (2) Events are not things that we can intend to bring about.
- (3) Therefore, events are not things that we can intentionally bring about.

The conclusion of the ability-based argument—events are not things that we can have abilities to bring about intentionally—would entail the conclusion of the intention-based argument. Moreover, it seems that premise (2) of the intention-based argument could be formulated more strongly as follows: events are not things that we can have abilities to intend to bring about. If so, the intention-based argument could lead to the conclusion of the ability-based argument.

However, these two arguments still present two substantially different obstacles for human agents to bring about events intentionally. The intention-based argument is about our inability to have the required intention (even if this inability is also impossible to overcome). The obstacle illustrated by the ability-based argument is a different one. As said in the Introduction, both arguments are supposed to “independently lead to the conclusion that actions are not

events.” So, in this chapter, I leave it open whether it is possible for us to intend to bring about events. By discussing the ability-based argument in such a dialectic context, I try to show not only that the ability-based argument is sound but also that it can work even if the intention-based argument turns out to be unsuccessful.

Now, back to *I*-ability. As Mele realizes, by illustrating the distinction between these two kinds of abilities, he does not provide an analysis of being able to *A* intentionally nor of the control intentional action requires. Therefore, even if the reader is convinced that there exist two kinds of abilities—one being stronger than the other—at this point, it is not yet clear whether human beings can have *I*-abilities. The only thing that is clear at this point is that it is impossible for human beings to have *I*-abilities while not having the corresponding *S*-abilities. For our purpose, we need an account of the notion of *I*-ability that goes beyond the intuitive illustration of its difference from *S*-ability. Before doing that, some preliminary clarifications are needed.

First, for Mele, both *S*-ability and *I*-ability are what are sometimes called “specific abilities.” The distinction between general abilities and specific abilities is widely adopted among philosophers.¹ As Mele remarks,

Although I have not golfed for years, I am able to golf. I am not able to golf just now, however. I am in my office now, and it is too small to house a golf course. The ability to golf that I claimed I have may be termed a *general* practical ability. It is the kind of ability to *A* that we attribute to agents even though we know they have no opportunity to *A* at the time of attribution and we have no specific occasion for their *A*-ing in mind. The ability to golf that I denied I have is a *specific* practical ability, an ability an agent has at a time to *A* then or to *A* on some specific later occasion” (Mele 2003a, pp. 447).

¹ Philosophers use different terms to refer to this distinction. Honoré (1964) speaks of general vs. particular abilities. Vihvelin (2013) speaks of narrow vs. wide abilities. Maier (2013) speaks of general abilities vs. options.

To be fair to Mele, he does not claim that the distinction between *S*-ability and *I*-ability is exclusive to specific abilities and does not apply to general abilities. I am inclined to think that, if the distinction between *S*-ability and *I*-ability is a sensible one, it should exist for general abilities as well. But for this investigation, this complication can be bypassed. There exist certain ability-related necessary conditions for the performance of an action that seem to be satisfied only by one's specific ability rather than general ability. For example, although being able to play golf in the general sense, Mele does not have the ability required for playing golf in his office. The needed ability here is apparently a specific one. Therefore, even if there is a kind of general *I*-abilities and it is possible for human agents to acquire this kind of abilities with respect to particular actions or events, it does not alleviate the worry that it might be impossible for human agents to acquire abilities to perform particular actions intentionally after all. Thus, I am going to narrow my discussion to the realm of specific abilities.

Second, it is awkward to claim, as Mele does, that the notion of *S*-ability expresses an *ordinary* sense of "able" or "ability." As an approximation, being *S*-able to *A* simply means one's chance of *A*-ing is higher than zero. In other words, it is possible that the agent *A*-s. That is why one performance would entail the possession of an ability of this kind. So, a normal person who has no special control over dice is *S*-able to throw a six with a single toss of a fair die. Likewise, a person is *S*-able to win a lottery (by buying a lottery ticket) even if the chance is, say, one out of a million. But this is not the kind of abilities when people use "able" or "ability" in an ordinary sense. We do not say that a person has an ability to win a lottery unless some method of cheating is available to her. Instead, we say "it is possible for her to win the lottery."

Third, it is not clear whether the notion of *I*-ability expresses an ordinary sense of "able" or "ability" either. Mele's example for *I*-ability is one's ability to roll a die intentionally. Since

this seems to be a typical case of ability in an ordinary sense, it might create the impression that *I*-ability is meant to be an ordinary ability.² In fact, it seems to suggest that *I*-ability is the kind of ordinary abilities that we talk about most in real life when it is said that I am able to drive to the campus within the next 20 minutes or that I have an ability to cook a meal for myself. To have the ordinary ability in this sense seems to require that the agent can perform the action reliably. If the chances of my driving to the campus (successfully) within the next 20 minutes are merely 50%, then I cannot be said to have an ability to do so in the most ordinary sense.³ However, is this ordinary ability identical to *I*-ability? Not necessarily.

In terms of reliability (i.e., the success rate or chances of *A*-ing), the situation for *S*-ability to *A* is simple and straightforward. The required success rate for *S*-ability is simply higher than zero. As long as the agent has nonzero chances of *A*-ing, she is *S*-able to *A*. This explains why one performance of *A*-ing by a person is sufficient (although not necessary) for that she has an *S*-ability to *A*.

The problem is that it is not immediately clear what the required success rate is for *I*-ability, which is the key to the ability-based argument. As Mele points out, “A proper account of being able to *A* intentionally hinges on a proper account of *A*-ing intentionally and the control that involves” (Mele 2003a, pp. 448). However, philosophers have conflicting intuitions about how much control or what success rate is required for acting intentionally. According to

² My best guess is that Mele does take *I*-ability to be the kind of ordinary ability that we talk about most in real life. One clue is his choice of example. Another clue is that he discusses three kinds of ability: *S*-ability, *I*-ability, and ensurance-level ability that is more reliable than *I*-ability (it puts the agent in a position to make a sincere promise about the action). It seems clear that ensurance-level ability is stronger than, thus different from, the ordinary ability that we talk about most (although ensurance-level ability is arguably a kind of ordinary ability).

³ I do not suggest that there is a fixed success rate (e.g., 90%) for the ordinary ability in question here. A more plausible view seems to be that the required success rate is sensitive to the context, especially relative to the difficulty of the action. To illustrate, 90% might be required for the ordinary ability to drive to campus, but a much lower success rate (say, 40%) could suffice for the ordinary ability to score 3-point in basketball.

Christopher Peacocke, for instance, for a person who attempts to hit a croquet ball through a distant hoop while believing his chances of success are tiny, it is undisputed that “if the agent does succeed in getting the ball through the hoop, then he does so intentionally” (Peacocke 1985, pp. 69).

Gilbert Harman takes it to be “quite controversial” whether a person who tries to sink a long and difficult putt “can be properly said to have intentionally sunk the putt” if she sinks the putt to everyone’s surprise, including her own (Harman 1986, pp. 91-2). Likewise, Brian O’Shaughnessy compares a dart champion and a novice. Both of them intend to hit the bullseye on a dart board. He claims that “while they both can succeed in doing what they tried and intended doing, only the champion’s act counts as an intentional hitting of the target” (O’Shaughnessy 1980, pp. 568).

The issue here is whether intentional actions require a relatively high success rate or not. Harman and O’Shaughnessy tend to answer yes while Peacocke disagrees. The problem for Mele is that he confesses he knows of no way to settle the controversy between Peacocke and Harman (or O’Shaughnessy). In other words, Mele cannot rule out the possibility that a person may be able to act intentionally even if her chances of so acting are not high at all. If this possibility is real (as Peacocke argues), then the ability required for intentional actions is not identical to ability in the most ordinary sense. This means that, even if human beings cannot have a relatively high success rate with respect to performing particular actions, unless I find a way to settle the debate between Peacocke and Harman (or O’Shaughnessy) in favor of the latter, I cannot conclude that human beings cannot have the *I*-ability to perform particular actions. It is possible that an *I*-ability to *A* might well allow a comparatively low success rate of *A*-ing, thus being a notion of ability weaker than abilities in the most ordinary sense.

For the rest of this chapter, I proceed with the following assumption. It seems obvious that human beings cannot have a relatively high success rate with respect to performing particular actions insofar as the latter are viewed as particular events. Given the nature of particular actions—i.e., as events—any slight change in a particular action would make it become a different one. Thus it is simply impossible for any human agent to bring about the performance of any particular action reliably. This is most obvious for more complicated actions such as driving to the campus. But even for a simple particular action or event such as raising one's hand in a particular manner at a particular time, no one can have a relatively high success rate in bringing about this event that is actually brought about. We do not and cannot have the extremely accurate control over our bodies that is required for reliably bringing about the event or particular action in question. Anyone who doubts this should try to repeat a simple bodily movement of raising her hand in its finest details (setting aside the complication of different timing).

Second, I assume that Peacocke is correct in thinking that intentional action does not require ability in the most ordinary sense but only a comparatively weaker sense of ability. So, the questions left are to find an account of this weaker sense of ability, and then, to determine whether we can have abilities in that sense—weaker than the ordinary sense of ability—to bring about events intentionally.

2. Minimal Ability

Michael Costa (1986) has provided a *prima facie* plausible account of what he calls “minimal ability,” which is a necessary condition of acting intentionally with respect to an agent's control over the action in question. Since Costa also believes that an agent can act intentionally even if her chances of success are very low (say, one out of a hundred), his notion

of minimal ability seems to be the notion of I-ability in its weaker form that I seek to identify (henceforth, I shall adopt Costa's term "minimal ability"). The account is given in response to the contrast between the following two cases.

In case I, Smith is offered ten dollars if he can shoot a bottle placed on a distant wall. Smith is not an expert marksman, and he knows it. The chances that he will hit the bottle at this distance are remote (say, one in a hundred). Yet Smith has nothing to lose. He aims as carefully as he can and fires, hoping to hit the bottle. The shot hits the bottle, and Smith gets the ten dollars. In case II, Smith is playing a board game. He needs to roll a six with the die to win the game, thereby winning ten dollars. He throws the die, hoping to throw a six. The die is a fair die and Smith has no special manipulative powers with dice, yet he does throw a six and wins the ten dollars.

Comparing these two cases raises a puzzling question. As Costa writes:

I think intuitions are fairly clear that Smith shot the bottle intentionally in case I, but that Smith did not throw a six intentionally in case II. Case I shows that one can do something intentionally without having the ability (in any ordinary sense) to do it. ... But then why isn't Smith's action intentional in case II? Smith's actual chances of succeeding were much better in case II (one out of six) than they were in case I (one out of one hundred) (Costa 1986, pp. 144).

This puzzle is especially pertinent to the attempt to identify minimal ability. It may not be surprising that Smith can act intentionally even if he does not have an ordinary ability to do so when the ordinary ability requires a relatively high success rate.⁴ But how can it be that a lower success rate meets the criterion for minimal ability while a higher success rate does not? If our intuitions about case I and case II are not mistaken, the criterion for minimal ability cannot be a

⁴ This notion of "ordinary sense of ability" is obviously different from S-abilities, which according to Mele is also an ordinary sense of "abilities." As I mentioned above, Mele's notion of S-abilities is not "ordinary" as how most people would understand the word.

fixed success rate between zero and one (e.g., 5%). But how do we determine the required success rate for minimal ability? Costa proposes an account of minimal ability based on the rough idea that “one has the minimal ability to do an act if the *chances* of one’s doing the act are *improved* by one’s trying to do it” (Costa 1986, pp. 145). In case II, Smith’s chances of throwing a six were not improved by his trying to throw a six, but Smith’s chances of hitting the bottle in case II were improved by his trying to hit the bottle with the shot.

However, Costa thinks that trying is not the right concept due to two problems. First, a nervous person may have lower chances of succeeding when she tries to perform the action. Second, in case II, Smith’s trying to throw a six does improve the chances of throwing a six if it is compared with Smith’s trying nothing (the chances being zero). To avoid these objections, Costa’s final formulation of the notion of minimal ability is paraphrased thus:

S has the minimal ability to *A* if and only if:

(1) *S* has the ability to *A*, or

(2) *S* has the ability to *B* ($\neq A$) such that the probability that *S A*-s given that *S B*-s in order to *A* thereby, is greater than the probability that *S A*-s given that *S B*-s but not in order to *A* thereby.

There remains a question as to how to understand the ordinary notion of ability in Costa’s account. Costa realizes this problem. He tries to deal with it by assuming that we can have an analysis of ability independent of the analysis of minimal ability. So, there would be no explanatory circularity (Costa 1986, pp. 146). Since Costa does not engage in the task of analyzing ability, he only attempts to provide “the basis for understanding the notion of minimal ability in terms of [ability]” (Costa 1986, pp. 146). Like Costa, I need not engage with the

analysis of ordinary ability since, even without such an analysis, I have already assumed that we *cannot* have abilities in the ordinary sense to bring about events.

Costa's analysis of minimal ability seems to save the ability requirement for intentional actions by accommodating our intuitions about cases I and II. In case I, there exists an action that Smith is able to perform, namely, shoot a bullet towards the bottle, which is such that his shooting a bullet towards the bottle in order to hit the bottle improves the chances of hitting the bottle in comparison with a case where he shoots a bullet towards the bottle but not in order to hit the bottle. In case II, however, Smith is able to throw the die and his throwing the die in order to throw a six does not improve his chances of throwing a six in comparison with a case where he throws the die but not in order to throw a six.

Although Costa's account is not yet a complete analysis of minimal ability, it bears an important implication for my investigation. Given the assumption that human beings cannot have the ability (in its most ordinary sense) to *A* when *A*-ing is a particular action, the only way to satisfy the necessary condition of having the minimal ability to *A* is through condition (2). As a reminder, in order to have the ability-based argument that is independent of the intention-based argument, I assume in this chapter that human agents can intend to perform particular actions. If so, the problem for me is that it seems possible for human agents to satisfy condition (2) of the minimal ability to *A* (even when *A*-ing is a particular action).

It remains true that human agents cannot do anything to raise the probability or the success rate of *A*-ing—a particular action—to a relatively high level. But if, as I assume, human agents can intend to perform *A*, it seems reasonable to expect that one's effects to execute such an intention would at least minimally increase the probability of *A*-ing. If I (magically) foresee my performing a future particular action, say a particular instance of my raising my right hand,

and I form an intention to do exactly that, I can try my best to mimic the particular action appearing in my foresight and thus raise the probability of doing so. For condition (2) to be met, the increase of the probability need not be a lot. A slightly higher probability or success rate of *A*-ing would be sufficient.

In response to this problem, I argue in the next section that Costa's account of minimal ability is not correct. His account is on the right track, but its current formulation faces serious problems. Therefore, it needs a revision.

3. A Revised Account of Minimal Ability

Consider the following case raised by Charlotte Katzoff: Brown puts one live cartridge into a six-chambered revolver, and spins the chamber as he aims it at White at point-blank range. He pulls the trigger hoping to hit White and he does (Katzoff 1989, pp. 236). In this case, it seems pretty clear that Brown kills White intentionally. However, this intuition does not fit with Costa's account of minimal ability. According to Costa, Brown does not have the minimal ability to kill White because he lacks the ordinary ability to kill White since there is only one live cartridge in the revolver. Also, given that he is shooting at point-blank range, the probability of his killing White is not increased by his shooting in order to kill White (Katzoff 1989, pp. 236).

The problem is that Costa, in face of the potential objections to his rough formulation, mistakenly replaces the notion of trying with the notion of *B*-ing in order to *A*. To be fair, to try to *A* might just be to do what one can to *A* (Hornsby 2010). This account has two elements: one is about doing something else (i.e., *B*-ing) that one is able to do, the other is about one's intention to *A* (i.e., in order to *A*). So, for example, if Smith tries to shoot a bottle as in case I, he needs to do the "means" act (to borrow Costa's term)—say, adjust breath, aim carefully, and fire the gun—so that his chances of hitting the bottle will be improved. Also, he needs to do so with the

intention of hitting the bottle. In other words, trying to *A* entails intending to *A* (I have already defended this view, SIT, in chapter 2). If Smith lacks this necessary motivational mental state, even if he is doing something that can increase his chances of hitting the bottle, he is not trying to do so.

The intention and the “means” act work separately. Having the relevant intention explains why Smith does what he does, namely, adjust breath, aim carefully and fire the gun. But the intention itself is not related to the odds-improving effect of doing the “means” act. This is the lesson we learn from Katzoff’s counterexample. The “means” act that Brown does is a complicated one, constituted by loading the revolver, spinning the chamber, and firing the gun at point-blank range. With or without the intention of killing White, doing such a “means” act would result in the same chances of killing White. So, the chances are not improved by doing this act *in order to kill White*.

To put it another way, to do the “means” act is to exercise the minimal ability that is allegedly possessed by the agent. For a person to have a minimal ability is to have a corresponding “means” act in ready, and by performing the “means” act, the chances of doing what is intended get improved in comparison to the chances when the “means” act is not performed. Applying this thought to Smith’s throwing the die. If the relevant “means” act is simply the act of throwing the die in the normal way, then there would be no difference in terms of chances of throwing a six between doing so in order to throw a six and doing so not in order to throw a six. Throwing the die in the normal way is the typical “means” act of exercising the ability or minimal ability to throw the die rather than throw a six. For the chances of getting a six to be improved beyond one out of six, the “means” act must be something different from throwing the die in the normal way. A rational person who mistakenly thinks that she has the

minimal ability to throw a six must be ready to do some “means” act (e.g., hold the die in a particular manner, throw it at a particular angle with a particular force, or perhaps mutter a particular spell before throwing the die) which is different from grabbing the die and throwing it out at will (which is a typical “means” act for throwing the die in the normal way). So, the reason why Smith does not have the minimal ability to throw a six in case II is that, among all actions that Smith is able to perform, nothing’s performance could improve the chances of getting a six beyond one out of six.

Accordingly, Costa’s account of minimal ability needs to be revised as follows:

S has the minimal ability to *A* if and only if:

(1) *S* has the ability to *A*, or

(2) *S* has the ability to *B* ($\neq A$) such that the probability that *S A*-s given that *S B*-s is greater than the highest probability that *S A*-s given that *S* does not *B*.

The reason why I add “highest” to qualify the probability that *S A*-s given that *S* does not *B* is that not *B*-ing (i.e., not exercising the alleged minimal ability) does not result in a fixed probability of *A*-ing. Therefore, for a person to have the minimal ability to *A*, she should be able to improve the chances of doing an *A* upon the highest probability of that when such a minimal ability is not exercised. When Smith does not exercise his minimal ability to shoot the bottle, his chances of hitting the bottle might be zero if he does not shoot at all, or one out of two hundred if he shoots without aiming carefully. The point is: Smith’s having the minimal ability means he can do something to improve the chances of hitting the bottle beyond the best chances (e.g., one out of two hundred) when he does not exercise the minimal ability.

How about the original objection Costa raises for himself? Smith does not have the minimal ability to throw a six. However, there does exist a “means” act—throwing the die in the

normal way—the performance of which improves the chances of getting a six in comparison to the chances when this “means” act is not performed. I agree with Costa that this is a real problem for the rough formulation of the notion of minimal ability. However, I think Costa’s reaction to this problem is incorrect. The reason why Smith still has no minimal ability to throw a six is that, by throwing the die in the normal way, the chances of getting a six are improved only with equal improvements of the chances of getting other results. Not only the chances of getting a six are raised from zero to one out of six, but the same improvement happens to the chances of getting a one (or two, three, four, five). A minimal ability admittedly exists in such a case. However, it is not the minimal ability to throw a six, but rather the minimal ability to throw either a one, or a two, or a three, or a four, or a five, or a six. It is the minimal ability to throw the die.⁵

This judgment is backed up by the intuition that a person cannot get a head in one coin flip intentionally either. She is not minimally able to get a head with a single coin flip. The explanation must be that, by flipping the coin, the person not only increases the probability of getting a head from zero to 50% but also equally increases the probability of getting a tail. The lesson from this example seems to be that, if a person has the ability to *B* and *B*-ing would increase the chances of *A*-ing and the chances of *C*-ing *to the same degree*, and the relation between *A*-ing and *C*-ing is analogous to the relationship between getting a head and getting a tail, this person does not have the minimal ability to *A* (or *C*). What she has, based on his ability to *B*, is the minimal ability to either *A* or *C*. Since the relationship between getting a head and getting a tail is mutually exclusive, i.e., incompatible with one another, it suggests that a minimal ability to *A* requires that the increase of *A*-ing’s probability must be greater than the increase of the probability of anything that is incompatible with *A*-ing. The “anything” here should be

⁵ The person also has an ordinary ability to throw the die.

understood broadly. That means that, in the case of throwing the die, not only is getting a five incompatible with getting a six, getting either a five or a four (or getting either a five, or a four, or a three, or a two, or a one) is also incompatible with getting a six. I see no reason to include one but not the other.

However, since Smith's getting ten dollars is not incompatible with getting a six, whether the chances of the former are also improved or improved by how much is irrelevant to the question of whether Smith has the minimal ability to throw a six or not. Now, I propose that we need to add a further condition to complete condition (2) in the earlier account of minimal ability, namely, that by performing the relevant "means" act, the increase of the probability of *A*-ing is greater than the increase of the probability of doing anything which is incompatible with *A*-ing. A difficulty remains. The added condition does not specify how greater the increase of the probability of *A*-ing needs to be so that the agent can be said to have the minimal ability to *A*. Without specifying this, a person may have the minimal ability even if the increase of the probability of his *A*-ing is merely minimally greater than the increase of something incompatible. Consider someone who can flip the coin in a particular way so that the probability of getting a head and of getting a tail are increased from zero to 51% and 49% respectively. If a 50-50 split does not amount to the minimal ability to *A*, people should feel reluctant to ascribe the minimal ability to *A* based on a 51-49 split. A more plausible position seems to be that the increase of the probability of *A*-ing needs to be significantly greater than that of the probability of anything incompatible with *A*-ing. But it is not easy to specify how significant it needs to be.

Therefore, if the added condition is formulated as it was above, it can only serve as a part of a necessary condition for minimal ability instead of a necessary and sufficient condition. That is,

S has the minimal ability to *A* only if:

(1) *S* has the ability to *A*, or

(2) *S* has the ability to *B* ($\neq A$) such that the probability that *S A*-s given that *S B*-s is greater than the highest probability that *S A*-s given that *S* does not *B*; and that the increase of the probability that *S A*-s is greater than the increase of the probability that *S* does anything which is incompatible with *A*-ing.

This may not sound satisfactory if what I look for is a complete analysis of minimal ability. Fortunately, since I only need a necessary condition of minimal ability to rule out the possibility that human beings have such abilities with respect to particular actions, this account of minimal ability's necessary condition will be good enough.

Now, I consider another relevant example discussed by Mele and Sverdlik (1996) which corroborates my account of minimal ability: Brown wants to kill Smith now, and he believes that the only way he can do so is by throwing a six now. Smith is in another building. There is a bomb in that building and Brown can detonate it only by throwing a six with a single toss of a fair die. He tosses the die onto the table, hoping that it will land six-up. Unluckily for Smith, the die lands six-up. By throwing a six, Brown detonates the bomb, thereby killing Smith.⁶

Many people are inclined to believe that Brown intentionally kills Smith. How can this intuition be reconciled with the intuition that Smith does not intentionally throw a six in case II? What accounts for the difference between these two cases? Specifically, since *A*-ing intentionally entails one's having the minimal ability to *A*, how could Brown (just like Smith, without any special power over dice) have no minimal ability to throw a six but have the

⁶ This is a shortened version of Mele and Sverdlik's original formulation of the case (1996, pp. 279). I leave out some unnecessary details.

minimal ability to kill Smith? The revised account of minimal ability could easily answer these questions. Let us take the act of throwing the die to be the “means” act of Brown’s alleged minimal ability to kill Smith. Given that Brown does not throw the die, the chances of his killing Smith are zero. The chances of his not killing Smith are one hundred percent. Given that Brown throws the die, the chances of his killing Smith are increased from zero to one out of six. By throwing the die, Brown contributes to a significant and the sole possible increase of the chances of his killing Smith. In the meantime, the chances of his not killing Smith are decreased from one hundred percent to five out of six. Since not killing Smith is the only relevant event or state that is incompatible with Brown’s killing Smith, condition (2) is satisfied by Mele and Sverdlik’s example. In contrast, as we already know, condition (2) is not satisfied by case II.

In conclusion, since I assume that condition (1) cannot be satisfied by human agents with respect to particular actions, condition (2) in my account of minimal ability is a necessary condition for human agents to possess minimal abilities to perform particular actions. In the next section, I will argue that human agents cannot have minimal abilities to perform particular actions.

4. A Timing Argument

In this section, I propose a timing argument for the claim that human agents do not have minimal abilities to perform particular actions:

- (1) We cannot have the minimal ability to do anything at a particular time.⁷
- (2) The exact timing of an event is essential to the identity of the event.

⁷ By a particular time, I leave it open whether it is an interval of time that the event occurs (e.g., 11:37:00 am-15:42:00 pm on May 15, 2017) or a particular time-slice moment (e.g., 15:42:00 pm on May 15, 2017).

(3) Therefore, we cannot have minimal abilities to perform particular actions or bring about events.

By proposing this timing argument, I do not suggest that human agents' lack of minimal ability to control the timing of their actions is the only reason why they cannot have minimal abilities to perform particular actions. If particular actions are taken to be bodily movements, I am inclined to believe that human agents do not have minimal abilities to control their precise bodily movements, and also that the precise bodily movement is also essential to the identity of a particular action. However, the existence of other potential objections does not undermine or alleviate the force of the timing objection. In addition, it is not controversial that timing is involved in the ontology of events (whether it is essential to events' identity is another issue I need to address later). In contrast, it is debated whether a "particular action" necessarily involves a certain bodily movement.⁸ So, for the sake of brevity, the timing objection seems to serve the purpose of my investigation best.

I offer the following argument to defend premise (1):

- (i) If a person has the minimal ability to *A* at a particular time *t*, she needs to be able to *A* at *t* or be able to *B* so that, by *B*-ing (an act-universal), the probability of *A*-ing at *t* is increased more than anything that is incompatible with *A*-ing at *t* (e.g., *A*-ing at *t*').
- (ii) A human agent is not able to do anything *at a particular time*.
- (iii) A human agent is not able to do anything that increases the probability of the exact timing of its occurrence more than any other timing.

⁸ See Moore (2010: 31-5) for a discussion of the objection towards the view that bodily movements are essential to "particular actions."

(iv) Therefore, human agents do not have the minimal ability to do anything at a particular time.

In this argument, premise (i) is what is stated by the revised account of minimal abilities. As for premise (ii), its truth is obvious since it is clear that human agents cannot reliably—no matter how loosely we understand the term “reliably”—control the exact timing of any action. That is, if a person *A*-s at *t*, the exact timing of *A*-ing at *t* is largely a matter of luck in contrast to other possible timing of *A*-ing, e.g., a millisecond after *t*. Anyone who has doubts about this claim must not understand “exact timing” in the most strict and literal sense. Just imagine perhaps the most accurate control of timing of anything that human beings do: launching rockets. Even for a rocket’s launching, it is not reliably done at an exact particular time that is intended. When a rocket is designed to be launched at a particular time, the actual launch time never happens at that time exactly (at least a few milliseconds before or after). It may *not* be a matter of luck if the timing of one’s action is understood not as a particular time (or time interval) but a specified range of time (or time interval). For example, I may be able to reliably control the start of my next raising a hand to a second, say, between 14:50:00 pm and 14:50:01 pm today, which is a range of time. But specificity is different from particularity. Time is too finely grained for human agents to have any substantial control of their behaviors’ exact particular timing.

Since premise (ii) is correct, the only possibility of having the minimal ability to *A* at *t* lies in the possible falsity of premise (iii). That is, there must be some action *B* that a person is able to do, and by *B*-ing, the probability of *A*-ing at *t* is increased more than anything incompatible with *A*-ing at *t* (e.g., *A*-ing at *t*’). The problem is that to increase something’s probability more than anything incompatible requires control that singles out, most accurately, the targeted thing, i.e., a particular timing of the action. But how could human agents have the

kind of control that is required for this unique increase of the probability of a particular timing of the action? For example, it is 5:30 pm right now. If I try to raise my hand in the near future, without making any special efforts to do it at a particular time, the probability of my raising my hand at 5:31:37 pm is approximately as likely as my doing so in close temporal location (e.g., 5:31:36 pm). So, to increase the probability of raising my hand at exactly 5:31:37 pm, perhaps, I need to find a very accurate clock and start to raise my hand when I see the clock turns 5:31:37 pm. But by doing so, the timing the probability of which is increased more than anything incompatible is likely after 5:31:37 pm (e.g., 1.42 second later, that is, 5:31:38.42 pm). Starting to raise my hand 1.42 second earlier—5:31:35.58 pm—does not solve the problem considering the accuracy that is needed for accurately doing so.

It should be noticed that the probability that is involved here is not subjective probability but objective probability. To see the difference, consider the odds of a soccer match given by a gambling agency based on sophisticated research. The odds are likely to be more or less different from the real odds of this match's result. The former is merely a subjective probability while the latter is objective. When I make great effort to start raising my hand at a particular time, it may be true that, from my epistemic standpoint, the probability of the targeted particular timing is most likely increased more than anything incompatible with it. However, it does not follow that objectively speaking, the probability of this particular timing is increased more than anything incompatible with it.

Also, it is true that there must be a *B*-ing (a particular bodily movement) such that, if a person does it, the probability of a particular timing of *A*-ing is increased uniquely in the sense that it is increased more than anything incompatible with this particular timing. However, this does not entail that a person has the minimal ability to *A* at this particular time. The reason is that

the required bodily movement (i.e., *B*-ing) would be a particular one so that it is beyond one's ordinary ability to do so reliably. In other words, even if there is something a person can do to uniquely increase the probability of a particular timing, this thing can only be performed based on extremely good luck rather than abilities. Therefore, premise (iii) is true. So, the upshot is that human agents do not have the minimal ability to do anything at a particular time.

Premise (2) of the timing argument says that the exact timing of an event is essential to the identity of the event. We can call it the "essentiality thesis." According to this thesis, events are not modally fragile in terms of their timing.⁹ We typically think that events could have occurred at slightly different temporal locations without changing into numerically different events. For instance, had WWI began on July 29th in 1914 instead of July 28th in the same year, it would have remained the same war. If I had raised my hand a second later, it would not have been a different raising of the hand.

Two different arguments have been raised to support the essentiality thesis. One is proposed by Lawrence Lombard (1982). Lombard first asks the reader to think of a possible world in which two extremely similar events occur at different times:

Let us suppose that there is a possible world, w_1 , in which a certain object, x , changes *twice* in a certain way at distinct times, t_1 and t_2 . That is, there occur in w_1 two distinct events, e_1 and e_2 , which are as alike as possible given the fact that e_1 occurs at t_1 and e_2 occurs at t_2 ... With respect to their essential properties (apart from the possibility that their times of occurrence are essential) and their non-relational accidental properties, however, e_1 and e_2 are exactly alike (Lombard 1982, pp. 10).

Lombard argues, correctly, that both the essentiality theorist and the inessentiality theorist should have no trouble allowing the stipulation of w_1 . Next, he argues:

⁹ Whether events can survive other kinds of small changes (e.g., changes of spatial location) is a more general question that I need not discuss here.

If the inessentiality thesis is correct, however, it need not be the case that e_1 occurs at t_1 and e_2 occurs at t_2 in every possible world in which they both occur. If so, then there is a possible world, w_2 , distinct from w_1 , which is to the extent possible exactly like w_1 except for the fact that in w_2 e_1 occurs at t_2 and e_2 occurs at t_1 The fundamental difference between w_1 and w_2 is that e_1 and e_2 switch temporal places (Lombard 1982, pp. 11).

With this setup, Lombard argues that the inessentiality thesis will be true only if, on the assumption that w_1 is a possible world, w_2 is a possible world distinct from w_1 . Then he argues that, although it seems clear that w_1 is a possible world and w_2 is distinct from w_1 , w_2 is not a possible world. His argument depends on his endorsement of the following principle: Possible worlds cannot be alike with respect to the truth and falsity of propositions concerning objects, the properties which those objects have, and the times at which those objects have other properties, and yet be unlike with respect to the truth and falsity of propositions concerning events (Lombard 1982, 13).

Based on this principle, given the fact that w_1 and w_2 are alike with respect to the truth and falsity of propositions concerning objects, the properties they have and then lack, and the times at which they have and then lack those properties, Lombard argues:

[E]ach proposition concerning events is true in w_1 if and only if true in w_2 as well. Thus, since it is true in w_1 that e_1 occurs at t_1 and true in w_1 that e_2 occurs at t_2 , it is also the case ... that it is true in w_2 that e_1 occurs at t_1 and true in w_2 that e_2 occurs at t_2 . But, it is true also in w_2 , by hypothesis, that e_1 does *not* occur at t_1 and true in w_2 that e_2 does *not* occur at t_2 . Thus, w_2 is a world in which is true that e_1 does and does not occur at t_1 and true that e_2 does and does not occur at t_2 . That is, w_2 is not a possible world (Lombard 1982, pp. 13).

I am not convinced by his argument although it is *for* the essentiality thesis. The problem is his claim that the inessentiality thesis will be true only if, on the assumption that w_1 is a possible world, w_2 is a possible world distinct from w_1 . Lombard is correct in thinking that if the inessentiality thesis is true, there should be a possible world in which e_1 occurs at t_2 and e_2 occurs at t_1 . However, it does not follow that this change of temporal locations of e_1 and e_2 can be done without any other fundamental changes between w_1 and w_2 . As Lombard stipulates, w_2 is “to the

extent possible exactly like w_1 ” except for the change of temporal locations of e_1 and e_2 . This would require that an event (e.g., e_1) may change its temporal location “magically” without requiring or entailing any other change. It is highly controversial if this kind of magical change of an event’s timing is metaphysically possible. Perhaps, some other changes that accompany this change of timing are negligible. However, it cannot be the case that there exists no such accompanied change at all. At least, some traces of small change must exist in the event’s causal history and causal effects. For example, even if inessentialists are correct in that WWI could have broken out a day later, it seems unlikely that they must accept the possibility that such a delay is done “magically” without any other changes in either the causal history of WWI or among its causal effects. Therefore, Lombard’s argument depends on an implausible premise that the possibility of w_2 is a necessary condition of the inessentiality thesis. So, his argument fails.

Now, let us turn attention to the second argument for the essentiality thesis. It is an argument recently proposed by Chad Vance (2016). Vance argues that any account of “non-fragile” event individuation is subject to a problem—the recycling problem—initially raised against the view that individuals are individuated by their origins. Let us consider two possible worlds that are devised by Vance as a preliminary challenge to the view that the exact timing is not essential to the identity of an event:

w^* In the actual world, Sebastian strolls leisurely around the block through the streets of Bologna at midnight. Call this event S^* .

w_2 In world w_2 , Sebastian leaves his home a bit earlier than he does in w^* , and takes a leisurely stroll (S_1) around the block through the streets of Bologna at 11:55 pm. Arriving back at his doorstep five minutes after midnight, and having enjoyed himself so much. Sebastian pauses for a moment before deciding to do it once more. So, at 12:05,

Sebastian takes a second leisurely stroll (S_2) around the block through the streets of Bologna.

Vance's argument seems to be this. If the exact timing is not essential to the identity of an event, at least one of S_1 and S_2 should be identical to S^* . But it is difficult to determine which one since both S_1 and S_2 are comprised of the same object (Sebastian) instantiating the same property (a leisurely stroll), at approximately the same space-time location (around the block in Bologna at approximately midnight). Obviously, it cannot be the case that S_1 and S_2 are both identical to S^* . This raises a difficulty for the inessentiality thesis.

Then, Vance considers whether it helps to solve the difficulty by adding the following condition as a necessary condition for event identification: both events have roughly the same causal history. The reason why the *exact* same causal history is not required is that any change of timing would result in a different causal history.¹⁰ Does this further condition enable us to determine which of S_1 and S_2 is identical to S^* ? Vance answers no because:

[W]ith both strolls, the complete causal history of the universe prior to each stroll is nearly identical, and furthermore, both strolls have the same immediate cause (i.e., something like Sebastian's being in a certain location and making a decision to stroll around the block) (Vance 2016, pp. 9).

Vance anticipates that someone may object that S_1 and S_2 do *not* have even roughly the same causal histories. This alone does not suffice for a defense of the inessentiality thesis.

Presumably, it needs also be the case that S^* and S_1 do have roughly the same causal histories.

So, the objection needs to defend "causal history" in a way that yields the conclusion that $S^*=S_1$, $S^*\neq S_2$, and $S_1\neq S_2$. But it may seem arbitrary to require such a conclusion:

¹⁰ This is so unless the complete causal history of the whole universe is shifted in time. But this picture entails an absolute existence of time unrelated to changes in the universe, which is no longer a plausible concept of time.

The feature that makes it such that S_2 has a different causal history than S^* is the presence of the small perturbation of a prior 11:55 stroll in S_2 's causal history (which S^* lacks). But one wonders, then, how the small perturbation of the 5 min of events from 11:55 to midnight that precede S^* (which S_1 lacks) does not also prevent S^* from being the same event as S_1 (Vance 2016, pp. 9).

So, without an independent reason to treat the 11:55 stroll and the 5 min of events from 11:55 to midnight differently, as Vance argues, “the delivery of different verdicts in the two cases would be quite odd, and seemingly ad hoc (Vance 2016, pp. 9).

It seems to me that Vance's response is not convincing considering that there seems to be a clear difference between the 11:55 stroll and the 5 min of events from 11:55 to midnight. It seems plausible that, without the 11:55 stroll, S_2 would not occur since S_2 is caused by Sebastian's enjoying his 11:55 stroll and his deciding to do that *again*. But without the 5 min of events from 11:55 to midnight, S^* probably would still happen. There is no reason to think that the occurrence of S^* causally depends on what happened between 11:55 and midnight. So, it is not ad hoc to argue that S^* and S_1 have roughly the same causal histories while thinking that S_1 and S_2 do not.

This is not a fatal problem for Vance's argument. It can be fixed by slightly revising his second possible world w_2 . Instead of stipulating that Sebastian enjoyed himself and decided to do another stroll, it could be that when Sebastian arrived back at his doorstep five minutes after midnight, he totally forgot that he had already done a stroll and simply decided to do a leisure stroll—the second stroll—at 12:05 am. In this way, we make clear that Sebastian's second stroll does not depend on his first stroll that night. Now, Vance would be correct in maintaining that it is arbitrary or even ad hoc to reach a conclusion that $S^*=S_1$, $S^*\neq S_2$, and $S_1\neq S_2$, as opposed to the conclusion that $S^*=S_2$, $S^*\neq S_1$, and $S_1\neq S_2$.

Given Vance's case of recycled strolls, the inessentiality thesis is refuted. The exact timing is essential to the identity of events. If so, premise (2) of the timing argument is true. Since both premises of the timing argument are true, its conclusion—if actions are events, human agents cannot have minimal abilities to perform actions—is defended. Because minimal ability is merely a weaker form of *I*-ability in terms of agents' control and the required success rate of the action in question, I conclude that events are not things that we can have *I*-abilities to bring about.

In this chapter, based on the ability-based argument, I have therefore shown that events are not things that we can have abilities to bring about. This conclusion, paired with the premise I take for granted in the Introduction of this dissertation—we can have abilities (i.e., *I*-abilities) to perform actions—would constitute an argument against the view that actions are events.

CHAPTER 4: THE RESPONSIBILITY-BASED ARGUMENT

In the past two chapters, I have offered two independent reasons—the intention-based argument and the ability-based argument—that support the view that actions are not events. In this chapter, I am going to propose a (moral) responsibility-based argument for the view that actions are not events.¹ The outline of the argument is as follows:

- (1) In order for an agent to be responsible for X , X must have a particular moral quality.
- (2) Events are not things that can have a particular moral quality.
- (3) Therefore, events are not things that we can be (morally) responsible for.

It should be noticed that this responsibility-based reason is also intended to be independent of either the intention-based argument or the ability-based argument. It does not depend on any apparent connection between a person's responsibility and one's intentional action or one's ability but merely relies on a reflection on the nature of moral responsibility itself.²

1. Responsibility and Moral Quality

It is typically taken for granted that we can be responsible for performing actions and consequences of actions. For example, if a person intentionally attacks an innocent person and the person is injured as a consequence, then it is almost uncontroversial that she is responsible for her action of attack and for the injury that she causes.³ It is more controversial whether there are responsibilities for other things in this case. Is a bystander responsible for her omission to

¹ In the rest of the chapter, what I mean by “responsibility” (or “responsible”) is moral responsibility unless specified otherwise.

² It seems plausible that intentional action or ability to act is a necessary condition for moral responsibility for the relevant action. But my responsibility-based reason is clearly not about this point.

³ Some philosophers who are skeptical about responsibility may disagree, e.g., Pereboom (2001).

help the victim? Some philosophers may have doubts about responsibility for omissions because it is not clear if omissions can be treated equally as actions. It is also possible to question whether the bystander is responsible for the injury given the fact that the injury is more of a direct consequence of the action of attack than of a consequence of the omission to help. On the other hand, it remains an interesting question how to make sense and defend the intuition that the connection between the omission and the injury is less direct.

These skeptical philosophers may think that the bystander is only responsible for her decision not to help the victim. But it seems questionable as well whether that is true. It is not clear to me that we can be responsible for our decisions or intentions when they are merely part of our inner mental life. Imagine that this person decides or forms an intention to attack someone else (for no good reason), but before she does anything, she managed to restrain herself and decide not to execute the earlier decision or intention. Is she responsible for her previous vicious decision or intention in this case? A quick reflection may lead to a negative answer. We all form morally questionable intentions from time to time. What matters seems to be how we handle ourselves after such an intention is formed. Assuming that the person realizes that her intended action is morally wrong, will she make adequate efforts to refrain from executing the intention, critically reflect upon the decision and eventually make a new decision not to attack the other person? Also, an affirmative answer would lead to an absurdity if a person struggles and changes her mind many times before actually doing anything. She would be responsible for too many things even when she does nothing that affects the outside world and people around her. This

raises the question whether decisions or intentions are the kind of things that we can be responsible for.⁴

However, what kind of things we can be responsible for can be discussed on a more fundamental and general level, for instance, whether we can be responsible for events, states of affairs, or even perhaps objects. In this chapter, I am going to challenge the view that we can be responsible for actions as events. The basic idea is simple: events, like objects, are not the kind of things that we can be responsible for. Let us think about why responsibility for objects seems a dubious idea. Does it make sense to claim that a person is responsible for an object such as a basketball? It is puzzling how to make sense of this alleged responsibility for a basketball. This claim is surely not about a causal responsibility for the basketball, or more precisely, for the event of this basketball's coming into existence. The only plausible reading of this might be that this person is responsible for the current existence of the basketball. However, the object basketball is not the same thing as the state of affair that this basketball exists now. "being responsible for the basketball's existence" is not an interpretation but a modification of "being responsible for the basketball," therefore cannot justify the latter as a sensible statement. The lesson seems to be that it is a category mistake to speak of responsibility for objects.

I argue we can say the same thing about responsibility for events. For such an argument to work, I need to explicate what kind of things we can be responsible for, although a complete and thorough answer to this question is unnecessary for my argument. I only need to specify one necessary condition for things for which we can be responsible in order to rule out events. This

⁴ An alternative explanation of my intuition that a person is not responsible for her restrained intention or decision is that the person is responsible. But if she revokes it later, blaming her would be inappropriate. In other words, my example illustrates an excuse for responsibility (Wallace 1994, chapter 5).

condition is explained by premise (1): things for which we can be responsible must have a particular moral quality.

Premise (1) is best supported by the Strawsonian approach to the notion of responsibility: to be responsible for *X* (*X* can be an action, a decision, an omission, a consequence of one's conduct, or even one's attitude towards someone else) is to be an apt or appropriate target of certain reactive attitudes with respect to the case at hand, which includes "gratitude, resentment, forgiveness, love and hurt feelings" (Strawson 2003).⁵ The Strawsonian approach best illustrates what I mean by "moral quality" and why it is essential to the concept of responsibility.

To clarify, this Strawsonian view should be distinguished from Strawson's view about responsibility proposed in his influential paper "Freedom and Resentment" (1962). According to Strawson himself, to be responsible is *to be* subjected to those reactive attitudes. Strawson denies the possibility that there can be any external justification of the appropriateness of such practices of praising and blaming. However, most philosophers (e.g., Fischer and Ravizza 1998) would think that there is a distinction between being responsible and being held responsible. Thus, it raises the question of whether holding someone responsible by actually subjecting her to certain reactive attitudes is appropriate given some independent metaphysical facts about the agent. A well-accepted thought is that, only when a certain set of conditions are met regarding the agent's control over and knowledge of her action (or other things that she can be responsible for), is she responsible for her action (or other things that she can be responsible for) in the sense that it

⁵ Although I adopt an inclusive view that we can be responsible for performing actions, omissions, consequences, and even attitudes, I will focus on responsibility for actions in the rest of the chapter for the sake of brevity, and also because the goal of my investigation is to figure out whether we can be responsible for actions as events.

would be appropriate to subject her to certain reactive attitudes.⁶ This view is called “Strawsonianism.”

According to this view, that only given that a person is responsible for her action are certain reactive attitudes appropriate, how is this appropriateness defined and determined? As most philosophers recognize (including Strawson himself), various reactive attitudes and feelings—e.g., positive and negative appraisals, attitudes of admiration and resentment, explicit verbal and behavioral expressions of gratitude and reproach—constitute a whole continuum of inter-personal moral reactions that may be made appropriate in response to someone’s responsibility. It is hard to imagine that the appropriateness is not defined, at least partly, in terms of how much a particular reactive attitude along the continuum fits and is in proportion to the moral nature of what the agent is responsible for. In principle, a less significant moral transgression tends to call for less severe blame, and a more significant moral transgression tends to call for more severe blame. The same pattern would apply to morally praiseworthy cases.

I use “moral quality” to refer to this moral nature of the action that is used to define and measure the appropriateness of reactive attitudes. Basically, it is a measure of how good/bad or right/wrong something is. When “something” is an action, it is natural to describe it as a right action or a wrong one. When it is a consequence of action, it is natural to describe it as a good consequence or a bad one.⁷ First of all, however, both good/bad and right/wrong can sometimes be used in non-moral senses. For example, that a horse is good, that studying Chinese is a bad idea, that studying philosophy is the right choice, or that $1+1=3$ is the wrong answer. Moral

⁶ I adopt an inclusive attitude towards what can we be responsible for. For that, I include decisions, actions, omissions, consequences of actions, and attitudes. But for the sake of brevity, I am going to focus on responsibility for actions since the central question that concerns us is whether we can be responsible for actions as events.

⁷ It quickly becomes more complicated when it comes to omissions, decisions, attitudes. Some of them are more similar to actions, while the others are like consequences.

qualities are morally good/bad or morally right/wrong. They are closely connected with what we take to be moral values or moral principles.⁸ That something is morally good means it is valuable according to a moral value that is worth achieving (e.g., happiness or life). That something is morally right means that it is recommendable and praiseworthy according to a moral principle about what we should do (e.g., be honest or be charitable).⁹

We can in any case follow the commonly adopted terminology by calling good/bad “axiological” and right/wrong “deontic” categories.¹⁰ Philosophers may find interesting what the relationship between axiological concepts and deontic concepts is. Is one of them logically or conceptually more fundamental than the other? Some people (e.g., Aquinas) take moral goodness to be the more fundamental—conceptually and/or explanatorily—concept. Some other people (e.g., Kant) will take moral rightness to be more fundamental. Some (e.g., Ross) propose a hybrid view with multiple moral values and principles that cannot be reduced to each other. But I do not find it necessary to commit myself to either camp since all that matters to my argument is that we can sensibly use axiological concepts of good/bad and deontic concepts of right/wrong in a moral sense and apply them to measure the moral nature and significance of actions (and other things).

A few more clarifications. First, I follow a broad understanding of responsibility which allows for being responsible for things that have positive, negative, or even neutral moral

⁸ But what makes a value or a principle a moral one? Is it merely a contingent fact that some values and principles are grouped together and labeled as moral? I don’t think so. But it is difficult to give a satisfactory answer. In my opinion, it is due to the fact that certain values and principles play a unique role in organizing our lives, especially the part that involves inter-personal relationships. And this may not be a purely a priori question, which means that its answer could go beyond conceptual research. With that said, a complete answer to this question is not essential to my investigation. We can simply follow how moral values and principles are traditionally categorized.

⁹ “Right” can be used in the minimum sense of permissibility. However, to build the parallel to moral goodness, moral rightness refers to positive moral qualities that are illustrated and recommended by some relevant moral principles.

¹⁰ Some philosophers would use terms such as “evaluative” and “normative” to describe those two moral concepts.

qualities. Accordingly, the agent would be an apt or appropriate target of agreeable and pleasant reactive attitudes (e.g., gratitude), disagreeable and unpleasant reactive attitudes (e.g., resentment), or no reactive attitudes at all. For the first two situations, the person is either praiseworthy or blameworthy for what she is responsible for. On the last occasion, the person is merely responsible for something without being praiseworthy or blameworthy.¹¹

Second, I do not take moral quality to be a strictly quantifiable property of those things that we can be responsible for as are happiness or pleasure to some utilitarians. It seems unlikely that everything can be put onto a unified continuum of moral importance or significance so that values can be used to represent moral import as its common measure. This is similar to what Ruth Chang (1997) calls “incommensurability” of value (e.g., liberty and equality).¹² I do not deny that killing two innocent people is a more significant moral transgression than killing one innocent person, and *killing* an innocent person is worse than merely *hurting* an innocent person. It seems awkward, however, to morally compare, for instance, hurting an innocent person and cheating on one’s wife, or lying to a friend and vandalizing some public property. They seem to

¹¹ Someone may raise a concern about this broad scope of things for which we can be responsible. It seems to make us responsible for too many morally trivial things, such as raising my left hand or preparing a meal for myself. It becomes even empty to make a claim about one’s responsibility and appropriate reactive attitudes when the appropriate reactive attitudes are no reactive attitudes at all. Also, it is not clear if the metaphysical condition of responsibility—whatever it is—can actually be satisfied when the action in question is morally neutral or perceived by the agent to be morally neutral. For example, during the process of my forming an intention whether to raise my left hand or right hand to signal my desire to ask a question, is what I do more similar to a mental “coin-flipping” or to a robust moral deliberation with the exercise of free will or whatever moral capacity that qualifies me as an agent who is responsible for my decision and action? I think these are all potentially good reasons to raise the concern about the broad scope of things for which we can be responsible. However, it seems to me that whether we narrow the scope to morally non-neutral things or not does not affect the force of my argument. I choose a broader scope because it leads to the most straightforward formulation of the connection between responsibility and blameworthiness/praiseworthiness.

¹² Given that some values are incommensurable, it is natural to think that moral qualities of things for which we can be responsible could be incommensurable as well. How bad or good a conduct is might well just be how much certain relevant value is harmed or violated (or protected or respected) in this conduct. Since different values might be involved with different conducts, the incommensurability of values entails the incommensurability of moral qualities.

belong to different categories of moral transgression which makes it difficult or even impossible to find a standard to measure their moral significance. This point makes better sense when we realize that the reactive attitudes differ in kinds and not just in degrees. For example, hatred may sound worse and more severe than resentment, but they may well be different kinds of reactive attitudes calling for different kinds of moral transgression. Increasing the degree of resentment does not eventually make it closer to and eventually become hatred. That I hate someone seems to implicate a hope on my part for something terrible to happen to this person, while resentment, even if very strong, does not involve such an evil will. Given this difference, hatred is more appropriately held against someone who also has an evil will towards me while resentment is more appropriate for someone who is responsible for hurting my interest but without an apparent evil will.

Third, moral qualities are supposed to be objective or socially entrenched features instead of subjective ones. That means, even if the victim happens to take a particular moral transgression very seriously as a result of, say, some subjective personal emotions, it does not make it appropriate for the victim to hold stronger and more severe reactive attitudes towards the wrongdoer. So, even if I adore small cute animals such as squirrels, I am not justified in holding correspondingly stronger resentment towards careless drivers who run over squirrels from time to time as if they were running over human beings. Of course, people should be allowed to react more strongly if they have objective personal reasons to take the moral transgression more seriously. For instance, it is appropriate to feel more resentment towards a squirrel-killer if the squirrel is my pet. Also, an even stronger reaction would be justified if the squirrel-killer kills my pet squirrel intentionally in order to hurt my feeling. If so, I may not only resent him for killing my pet, it even seems justified for me to hate him for holding such an evil attitude towards me.

But these escalated reactive attitudes are justified by some personal and objective reasons rather than mere subjective emotions or preferences.

Now, let us return to premise (1): in order for an agent to be responsible for X , X must have a particular moral quality. The logic behind it can be explicated in this way: responsibility is about the appropriateness of certain reactive attitudes, and it seems safe to assume that reactive attitudes are appropriate only if they are in proportion to the moral quality of X . As Justin Oakley puts, “To be responsible for something is to be open to creditworthiness or blameworthiness for it, but whether one is actually creditworthy or blameworthy for it depends also on its goodness or badness (or rightness or wrongness)” (Oakley 1992, pp. 124). But if so, it is hard to imagine that the goodness or badness (or rightness or wrongness) does not come with degrees.

A five thousand dollar punishment would simply be disproportional and thus inappropriate for a fifty buck crime. Of course, I am not suggesting that the moral quality of a thing for which a person is responsible is the sole factor in determining what reactive attitudes are appropriate. Some other factors may play a role as well. For example, the way in which the wrongdoer responds to her transgression also influences our judgment of what kind and degree of reactive attitudes are appropriate (Smith 2007, pp. 482). But, at least, the moral quality of X is one of the factors that are essential to the question of the appropriateness of corresponding reactive attitudes. If a thing has no moral quality, then it is conceptually impossible to adopt that someone is responsible for it since it does not make sense to say that any particular reactive attitudes are appropriate in that case.¹³ That explains why, as stated by premise (1), a necessary

¹³ It should be noticed that having no moral quality is different from having a neutral moral quality. An event does not have a neutral moral quality but has no moral quality at all. For something to have a neutral moral quality is for it to be morally indifferent. But an event of murdering might also be an event of keeping a promise. It seems false to say that the event is morally indifferent so that it does not have a neutral moral quality. It is just that the event has

condition for responsibility for X is that X has a particular moral quality that can be used to measure the appropriateness of moral blame or praise.

The “can” in the formulation of premise (1) may however give the wrong impression that the possibility of determining or assessing X ’s moral quality constitutes a strong requirement for responsibility. Under this impression, someone may think, mistakenly, that inability to assess the moral quality of X due to lack of information would render one’s responsibility for X impossible. But I believe that a person can be responsible for something even if this thing’s moral quality is merely unknown to us (all people who are related to and involved in this moral practice of responsibility) due to lack of relevant information. For example, a doctor makes a life and death decision about one of her patients. When some relevant conditions are satisfied, she is responsible for her decision and the immediate consequence of it (e.g., her patient’s survival or death). This is so even if the moral quality of her decision is inaccessible to anyone involved in this situation due to some epistemic obstacles (e.g., the patient survives, but all people involved including the doctor, her patient, and her patient’s family lose their memories due to an accident).¹⁴

Likewise, it seems possible that something’s moral quality can be determined in the future but cannot be determined at the moment of action or right after that. Maybe the doctor who makes the decision lives in an indeterministic universe and whether her decision turns out to have a positive or negative moral quality is yet to be determined. If so, it is still possible for the

different moral qualities under different descriptions and, without specifying which description, the event itself has no moral quality. To illustrate, think about a three-dimensional object, e.g., Earth. Does it have length? No. It would be weird to say that Earth is longer than Mars. It is not because Earth has a length of zero. It is just that Earth can be measured by length in too many ways (Earth’s longest perimeter or diameter is surely longer than Mars’s).

¹⁴ Here, I do not want to rule out the possibility that the decision has a particular moral quality as an event. But speaking of the moral quality of a decision does not imply that we are speaking of the moral quality of an event since a decision can be understood as something other than a particular event (just like actions, see the Introduction).

moral quality to be determined. There is at least no conceptual obstacle for a particular moral quality to be involved in that case. In such cases, instead of saying that there is no appropriate reactive attitude, the appropriate response is that we suspend holding any reactive attitude until *X*'s moral quality is determined.

Someone may ask: what if the whole universe ceases to exist before the consequence of her decision unfolds. Is this a case of “can be determined” or “cannot be determined”? I am inclined to regard it as the former because, in such an extreme case, the decision has a neutral moral quality—as far as its consequence is concerned—since no good or bad consequence is or will be caused by it.

With these things being said, the “can” I use in the formulation of premise (1) only refers to a weak modal requirement: it is not a category mistake to take the object to have such a moral quality. Premise (1), understood as such, aligns pretty well with my premise (2) that it is a category mistake to attribute a particular moral quality to an event. Roughly speaking, particular events are “multifacted” in the sense that they can be described in different ways. Therefore, multiple moral qualities can be attributed to a singular event under different descriptions.¹⁵ So, it does not make sense to ask what reactive attitudes are appropriate when it comes to responsibility for events. But a defense of premise (2) must wait until the next section.

The conclusion of this line of thought can be generalized beyond Strawsonianism. Even if being responsible is about some moral responses or reactions that are broader or different than Strawson's emotion-based reactive attitudes so as to include sanctions, awards, and punishments,

¹⁵ As Helen Steward (1997, chapter 1) argues, the “structureless” entities—particular events understood by Davidson—seem to best represent the feature of multifactedness and possessing a “secret life.”

as long as the moral reactions still need to be appropriate and in proportion to the moral quality of what the agent is responsible for, the force of the argument above remains.

An immediate worry is that Strawsonianism, even if broadly understood, is not the only theory in the market of theories of responsibility. An alternative view about the nature of responsibility holds that to be responsible is something like having a moral credit or debit assigned to one in a sort of a metaphorical ledger.

Praising someone may be said to constitute judging that there is a 'credit' in his 'ledger of life,' or a 'positive mark' in his 'report-card of life,' or a 'luster' on his 'record as a person'; that his 'record' has been 'burnished'; that his 'moral standing' has been 'enhanced.' Blaming someone may be said to constitute judging that there is a 'discredit' or 'debit' in his 'ledger,' or a 'negative mark' in his 'report card,' or a 'blemish' or 'stain' on his 'record'; that his 'record' has been 'tarnished'; that his 'moral standing' has been 'diminished.' Someone is praiseworthy if he is deserving of such praise; that is, if it is correct, or true to the facts, to judge that there is a 'credit' in his 'ledger' (etc.). Someone is blameworthy if he is deserving of such blame; that is, if it is correct, or true to the facts, to judge that there is a 'debit' in his 'ledger' (etc.) (Zimmerman 1988, pp. 38).

The ledger view involves no reference to the practices of Strawson's reactive attitudes (or other forms of reactive moral responses), or at least, as Fischer and Ravizza remark:

[T]he reactive attitudes and the associated practices of praising and blaming take on a secondary role, following from these primary assessments of moral worth like practical consequences. For the Strawsonian, this view of responsibility as involving primarily a theoretical judgment inappropriately diminishes the role the reactive attitudes play in the definition of moral responsibility (Fischer & Ravizza 1998, pp. 9-10, fn 12).

To see how the ledger view diminishes the role of reactive attitudes, we can imagine that each person's moral "ledger" is kept by God in order to make a final moral judgment upon us after we die and, based on which, appropriate praise or blame and reward or punishment would be given. Such a Christian version of our afterlife could be false. Still, it helps us understand the notion of the moral ledger. Even if the afterlife judgment and reward or punishment are removed from the picture, keeping a ledger which records a person's moral credit and debit is still a valid and sensible idea and practice. This is comparable to having a transcript that records each

student's academic performance even if the transcript is not used for the ground of granting a scholarship or offering a job. We might only have a purely theoretical interest in knowing how well a student does academically or how well a person does morally. It seems that the fundamental difference between the ledger view and Strawsonianism is that the latter emphasizes the practical aspect of the moral practice of holding responsible when human nature, human emotions, and interpersonal interactions play a primary role, while the former tends to abstract out all such day-to-day complexity and delicacy by intellectualizing (or maybe over-intellectualizing) our practice of moral responsibility and the practical aspect of holding responsible becomes secondary or even contingent to the theoretical judgment.

However, it seems that the ledger view still necessarily involves the question of whether the moral credit or debit put on the agent's ledger is appropriate and in proportion to the moral quality of what the agent is responsible for. A too "big" credit would be inappropriate if the agent is merely responsible for a "small goodness." And it is surely absurd to have a ledger that does not allow for credits of different "sizes" and only the number of credits or debits matters. Therefore, even the ledger view of responsibility necessarily appeals to the moral quality of what the agent is responsible for. To use the transcript analogy again, it is similar to the necessary fact that a transcript cannot be used to record a student's non-academic performances—i.e., those works that do not have an "academic quality"—since appropriate grades (like a moral credit or debit) must be given and recorded on the transcript in accordance to the academic quality of the student's works.

In conclusion, both the Strawsonian view and the ledger view support the necessary condition of moral responsibility illustrated in premise (1). The general lesson is that, as long as responsibility necessarily involves something—reactive attitudes, other forms of moral reactions

or responses, or even abstract moral credits or debits, etc.—that needs be appropriate according to the moral quality of the thing for which an agent is responsible, premise (1) is true.

2. Moral Quality and Events

Premise (2) of my main argument says that events are not things that can have a particular moral quality. To clarify, the reason why it is impossible for an event to have a particular moral quality is that, I shall argue, it is a category mistake to ascribe a particular moral quality to an event. It does not make sense to say that an event is a morally good or bad event or a morally right or wrong event. This view is similar to a view that is skeptical about good simpliciter. Good simpliciter skeptics (e.g., Geach 1956, Foot 1985, Thomson 1997) believe that there is no such thing as being good *simpliciter* or being of final value. Ascriptions of good are either ascriptions of good in relation to other things (e.g., good for someone, good as things of a particular kind, or good to read or eat) or simply meaningless and nonsensible.

The discussion of this view starts with Geach's famous claim that "good" (and "bad") is always an attributive, not a predicative, adjective. According to Geach, an adjective *A* is predicative if "*x* is an *A y*" splits up into "*x* is an *y*" and "*x* is *A*." For example, "red" is predicative in this sense since "is a red book" logically splits up into "is a book" and "is red." In contrast, attributive adjectives fail to follow this logical rule. "Big" and "small" are attributive: "*x* is a big flea" does not split up into "*x* is a flea" and "*x* is big," since *x* is big only relative to fleas. Geach argues, "good" shares this feature of attributive adjectives. For instance, "*x* is a good assassin" surely does not entail "*x* is an assassin" and "*x* is good." It is not even clear what the latter expression means. This leads to Geach's conclusion that "there is no such thing as being just good or bad, there is only being a good or bad so-and-so" (Geach 1956, pp. 34).

As Thomson suggests, the point is not just that “some substantive has to be understood” for ascriptions of good to be meaningful and sensible. Thomson asks us to reflect on odd expressions such as “*x* is a good corpse”: “What on earth does the speaker mean? I have no idea, because I have no idea what he means to be saying the corpse is good for use in, or good to, or good for, and so on” (Thomson 1997, pp. 278-9). So, Thomson would probably feel more comfortable with the expression “*x* (that corpse) is good for teaching human anatomy” or “*x* is good to scare away kids.” I agree with Thomson that Geach’s point should be understood in a broader way. It is not just that meaningful ascriptions of good must accompany a substantive (a kind of substance, e.g., assassins). Rather, “*x* is good” must accompany appropriate standards of evaluation, which can be provided or implied by mentioning *x* as a certain kind of thing (e.g., cars or assassins), or by putting *x* into a good-to or good-for relation with a certain activity (e.g., reading or eating) or a certain goal (e.g., health or knowledge), or simply by the context of the ascription of good.¹⁶

An implication of what we learn from the discussion above is that, sometimes, even “*x* is a good *K*” does not amount to a meaningful ascription of good. This happens when *K* is a too empty kind that does not imply any standard of evaluation. As Geach realizes, “though we can sensibly speak of a good or bad human act, we cannot sensibly speak of a good or bad event, a good or bad thing to happen. ‘Event,’ like ‘thing,’ is too empty a word to convey either a criterion of identity or a standard of goodness” (Geach 1956, pp. 40-1). This conclusion or its spirit is shared by both Foot (1985) and Thomson (1997):

¹⁶ Just like “flea” is not merely meant to mention a substantive but rather to provide or imply a particular standard of size. But it seems more difficult to find a mismatching example for “big” as like “a good corpse” for “good.” The reason seems to be that any spatial kind of things (e.g., fleas, tables, rockets, planets, countries, or even more generic things such as animals) usually, if not always, have a standard of size.

We should begin by asking why we are so sure that we even understand expressions such as ‘a good state of affairs’ or ‘a good outcome’; for as Peter Geach pointed out years ago there are phrases with the word ‘good’ in them, as, e.g., ‘a good event’ that do *not* at least as they stand have a sense. Following this line one might suggest that philosophers are a bit hasty in using expressions such as ‘a better world’ (Foot 1985, pp. 199).

‘Good state of affairs’ and ‘good event’ are unlike ‘good corpse’: it is not that we have no idea what a speaker might mean by predicating them, but rather that there are too many possibilities—the speaker might mean that the thing is a state of affairs or event that is good for him and his friends, or for people generally, or that it is good for use as an example in a class in philosophy of history, and so on and on (Thomson 1997, pp. 279).

All three philosophers, based on similar reasons, conclude that it does not make sense to say that something is a good event or that an event has a quality of goodness, full stop. This is especially clear when we adopt a Davidsonian coarse-grained theory of what events are. An event is a particular entity that can be described in different ways. According to Davidson, there could be a single event that can be described as my moving my finger, my flipping the switch, my turning on the light, my alarming the burglar, etc. Since different descriptions of an event may provide or imply different standards of good to evaluate the event in question, an event can be good under one description while being bad under another.¹⁷ Also, an event can be good to a high degree under one description while being good to a low degree under another. As Geach remarks, “Caesar’s murder was a bad thing to happen to a living organism, a good fate for a man

¹⁷ Someone may object that “may” suggests that it is still possible for an event to have a particular moral quality. If so, it is not all events but only some events that cannot be what we are responsible for. However, it is questionable if there is any event that can only have one moral quality. Either it has only one description, or only one description among many gives a non-neutral moral quality and we regard neutral moral qualities under all other descriptions as non-moral at all. Both options are not obvious. The second, and more crucial, difficulty of this objection is that, even if as a matter of fact an event has only one morally relevant description, it is still meaningless to speak of responsibility for this event without specifying the description. If anything, it makes it worse in the semantic sense. If an event has multiple descriptions and each of them gives the event a determined moral quality, then it is merely semantically incomplete to speak of responsibility for this event without specifying which description: it can mean this, or this, or that, etc. But if an event has a moral quality under only one description and has no moral quality under others, then it is no longer incomplete to speak of responsibility for it: it can mean this, or nonsense, or nonsense, or nonsense, etc. This does not make the problem disappear.

who wanted divine worship for himself, and again a good or bad act on the part of his murderers; to ask whether it was a good or bad event would be senseless” (Geach 1956, pp. 41).

This view that it is not sensible to ascribe good or a degree of goodness to an event is obviously relevant to premise (2) that I want to defend in this section. However, it is also too quick and hasty to conclude that premise (2) is true. Several problems need to be addressed.

First, good simpliciter skepticism, although a reasonable point, does have critics (e.g., Zimmerman 1999, Smith 2010). It is not clear that we should simply trust Geach’s argument for his conclusion that the concept of good simpliciter or intrinsic value is not sensible. Second, Geach concludes that it makes no sense to speak of a good event, which does not necessarily entail that it makes no sense to speak of a morally good event or morally right event. Premise (2) is about the evaluation of an event from a particular moral perspective. Third, Geach’s conclusion, according to Foot (1985), applies not only to events but also objects, states of affairs, or outcomes. If we use this reason to reject events as the kind of things for which we can be responsible, shouldn’t we equally reject states of affairs or outcomes as the kind of things for which we can be responsible? That is a hard bullet to bite. In the rest of this section, I will try to respond to these three difficulties.

It is a more or less commonplace for philosophers that good simpliciter and final or intrinsic values exist. Even though it is debatable what things are just plain good, it is uncommon to hold the skeptical position that things can be good or bad only relatively. This casts serious doubt on the support for premise (2) that is made available by good simpliciter skepticism. In response, however, I argue that the support for premise (2) is actually not based on good simpliciter skepticism.

To see this, it is important to realize that even opponents of good simpliciter skepticism usually do not deny that “good” is sometimes used in the relative way suggested by good simpliciter skeptics. The fundamental difference between the two sides is whether the relative way of speaking of goodness is the only sensible way to do so. Many philosophers would argue that it is not: there is a list of things (e.g., pleasure, happiness, knowledge, friendship, etc.) that are simply good, full stop. And sensible ascriptions of such intrinsic good does not depend on any standard of evaluation that is provided or implied by whatever the good is relative to. How is this possible? Does it mean that ascriptions of intrinsic good do not need any standard of evaluation, or that it does not need any external standard of evaluation? No matter how we choose to answer this question, it remains unclear “what way of being good does being intrinsically good involve ... proponents of intrinsic value have not traditionally concerned themselves with specifying in any great detail the way of being good that being intrinsically good involves” (Zimmerman 1999, pp. 404).

With that being said, even if there are intrinsic values or good simpliciter, particular events are not the kind of things that can be intrinsically good. Take Hitler’s suicide as an example. Let us say that the relevant intrinsic values in this case are human life and justice. Then the event might still be both good and bad, depending on which standard of evaluation is adopted. And the selection of the standard is a matter of which intrinsic value is implied as relevant by how the event is described. When the event is described as “Hitler’s killing himself,” then the relevant intrinsic value is human life. Accordingly, the event would be evaluated as a bad one since an intrinsically valuable life was lost. Under the description of “the deserved ending of Hitler,” the event seems to be a good one according to the standard of justice.

In conclusion, since events are not intrinsically valuable, the debate between proponents and skeptics of good simpliciter does not change the fact that events must be evaluated according to a certain standard. Given that being an event is still too empty to provide or imply any particular standard of evaluation, it remains meaningless to say that x is a good (or bad) event.

So far, I have only explained why it does not make sense to say that x is a good event. But, to turn to the second difficulty mentioned above, this is not precisely what premise (2) is about, which is whether it makes sense to say that x is a morally good or morally right event. The discussion above has not taken into consideration the moral feature of the topic. Considering that the difficulty of speaking of a good event is due to the fact that “event” is too empty a word to single out a particular standard of evaluation, someone may object to me by arguing that evaluation from a moral point of view may actually constitute a unique standard so that the difficulty above can be resolved. This potential problem for premise (2) requires that there is only one fundamental moral value (or moral principle) so that something (e.g., an event) could be morally evaluated by whether and/or how much the value or principle is realized or followed.¹⁸

Many utilitarians are value monists in the sense that they believe that pleasure or happiness is the only fundamental moral value and all other values (e.g., knowledge and friendship) are only valuable in so far as they contribute to the fundamental value. Kant, as a deontologist, is a monist about moral principles. He argues that there is only one supreme or fundamental moral principle and all other principles are derived from it.

¹⁸ It is commonly held that being fundamental or non-derivative is a necessary but not sufficient condition for intrinsic values. Here, the relevant question is whether a moral point of view is constituted by only one fundamental moral. Even if the fundamental moral value is not an intrinsic one (e.g., a fundamental virtue which is not good for its own but only for people’s sake), it may still provide a particular standard of evaluation for events.

It is much more difficult to imagine how the objection could work with Kant's fundamental principle. His attempt to derive all moral duties from the supreme moral law has well-known difficulties (see Hegel's *Philosophy of Right*, Wood 1999). Even setting aside this problem, a person's action either respects or fails to respect a principle. Using a fundamental moral principle as a standard of evaluation entails that the moral quality does not allow for degrees. Maybe with Kant's Universal Law formulation of the fundamental principle, it involves two kinds of duties—perfect duties and imperfect duties—so that there are four instead of two possible moral qualities. But still, it is not enough to explain how we can, justifiably, have various moral reactions based on one's responsibility for actions as events. It is too counter-intuitive that failure to follow or respect the perfect duty of never lying should always be taken equally as failure to follow or respect the perfect duty of never committing murder.

Utilitarian value monists seem to offer a more promising solution for a singular standard of evaluation for events. This possibility is discussed and dismissed by Foot when she tries to show that consequentialism owes its appeal to the confusing thought that "it can never be right to prefer a worse state of affairs to a better" (Foot 1985, pp. 198) because it mistakenly assumes that we understand what "good states of affairs" or "better states of affairs" means. Clearly, her argument is not limited to states of affairs but applies to events as well.

Foot starts with showing that this form of utterances is speaker-relative—"depending on what the speakers and their group are *interested* in" (Foot 1985, pp. 201). So, changing the perspective or value that the speaker is interested in may change the meaning and the truth of such utterances. Then, Foot discusses a reply based on the thought that "what we should really be dealing with in this discussion is states of affairs which are good or bad, not simply, but *from the moral point of view*" (Foot 1985, pp. 202). She agrees that a clear sense can be given to the idea

of the best state of affairs from a moral point of view, but one problem with this understanding is that it assumes a utilitarian perspective which many people find implausible. Foot links this perspective to the idea that “a benevolent person must often aim at the good of others and call it ‘a good thing’ when for instance a faraway disaster turns out to have been less serious than was feared (Foot 1985, pp. 204)”. However, to think that this understanding of “good states of affairs” is an “inescapable idea” would require us to think not only that benevolence is a virtue but that benevolence is the only virtue by disregarding other virtues such as friendship and justice since both put limitations on the pursuit of maximum happiness or pleasure for most people.¹⁹

To be clear, Foot is not claiming that a utilitarian understanding of “good states of affairs” from a moral point of view does not make sense. Her point is that it assumes a moral perspective or theory that is seemingly problematic or not so compelling to most people. So, people who are not consequentialists should not feel attracted by the thought hidden in consequentialism, namely, that it can never be right to prefer a worse state of affairs to a better. I think that Foot’s point applies to talk about of “good events” from a moral point of view as well. If the only way to ascribe a particular moral quality to an event is to assume a utilitarian perspective or theory, then it becomes a severe limitation for responsibility for actions as events because the truth of utilitarianism turns out to be a necessary condition of sensibly speaking of this form of responsibilities.

The way I handle the second problem raises another worry, namely, that we cannot ascribe a particular moral quality to states of affairs and, thus, cannot sensibly speak of

¹⁹ Friendship requires us to give service to friends rather than to strangers or acquaintances even when doing so would not maximize the overall happiness or pleasure. And how justice limits the pursuit of maximum overall happiness or pleasure is well-known.

responsibility for outcomes of our actions (e.g., being responsible for the fact that someone is injured). It seems that our actions' consequences or outcomes can be either events or facts (i.e., states of affairs that obtain). If I hit someone, the consequence can either be the event of this person's getting injured or the fact that this person is injured. So, given that premise (1) is true, if premise (2) is also true, then due to the fact that the difficulty of ascribing a particular moral quality to events equally applies to states of affairs, it would be impossible for us to be responsible for such outcomes. Since outcomes or consequences are either events or facts (i.e., states of affairs that obtain), it would make responsibility for outcomes or consequences impossible as well. But it is widely accepted that we can be responsible for at least some consequences of our actions, which is not something that I am ready to give up so easily.

In response, I need to point out that Foot's skepticism about good simpliciter is less radical than Geach's. Foot never denies that we can sensibly speak of good states of affairs, but merely emphasizes that the meaningful use of such utterances depends on the value or whatever that the speaker is interested in, thus is speaker-relative. As she remarks,

The doubt is not about whether there is some way of using the words, but rather about the way they appear in the exposition of utilitarian and other consequentialist moral theories. It is important readily to accept the fact that we talk in a natural and familiar way about good states of affairs, and that there is nothing problematic about such usage. But it is also important to see how such expressions actually work in the contexts in which they are at home, and in particular to ask about the status of a good state of affairs. Is it something impersonal to be recognized (we hope) by all reasonable men? It seems, surprisingly, that this is not the case at least in many contexts of utterance of the relevant expressions (Foot 1985, pp. 199).

So, the real problem for Foot is that the utilitarian use of "good states of affairs" seems to assume that it is impersonal—*not* speaker-relative—and could be recognized by all reasonable men. Here, I think Foot makes a mistake. She is correct in arguing that the meaningful use of such utterances depends on the value that the speaker chooses. She is also correct in thinking that utilitarianism—the principle of maximization of overall happiness or pleasure—does not

represent an impersonal value that is expected to be recognized by all reasonable men. However, she is incorrect to conclude that it is impossible to find a value that is impersonal and recognizable by all reasonable men. How about the fact that an innocent kid, Mary, is saved by doctors after being hit by a car? Can we hope that all reasonable men would choose to evaluate this fact by appealing to the value of human life and reach the conclusion that it is a good fact? Of course, we can! There seems to be no other reasonable way to evaluate this fact.

To figure out where Foot goes wrong, let us take a look at her discussion of an example:

Suppose, for instance, that the supporters of different teams have gathered in the stadium and that the members of each group are discussing the game; or that two racegoers have backed different horses in a race. Remarking on the course of events one or the other may say that things are going well or badly, and when a certain situation has developed may say that it is a good or a bad state of affairs (Foot 1985, pp. 199-200).

Let us focus on the horse race when two people—John and Peter—bet on horse A and horse B respectively. When horse A turns out to win the race, John can sensibly say it is a good state of affairs, while Peter can sensibly say that it is a bad state of affairs. Does this prove or at least support Foot's view? I do not think so. To be clear, what does "it" refer to here? If "it" refers to the fact that John wins some money, then it is a good state of affairs. It is not only true for John but also true for Peter. Setting aside his frustration, Peter can genuinely congratulate John by saying "well done." This is not merely because Peter thinks that John has done well *for himself* since we do not say "well done" to a successful murderer in the same way. The difference is that there is something good about the fact that someone wins something that is desirable such as money. It is a good state of affairs that all reasonable people would recognize.

By the same token, if "it" refers to the fact that Peter loses some money, then, it is a bad state of affairs not only for Peter but for John and all of us. Of course, someone may insist, that it matters that it is Peter rather than John who loses the money. First, although the event of Peter's losing some money hurts the interest of Peter, the fact (or the state of affairs) that he loses cannot

be bad for Peter in this causal sense since, as most philosophers believe, facts do not cause things to happen. But then, how do we account for this difference if the state of affairs in question is equally bad for both John and Peter? The difference may not lie in the evaluation of the state of affairs, but in the justification for further moral reactions based on the same evaluation. Part of the reason is that we need to accommodate those cases in which a person is justified to morally react to something not so relevant to her own interests (e.g., an innocent stranger's death). Even if, for John and Peter, the fact that Peter loses some money is equally bad, it may still leave it open what further moral reactions become appropriate and justified for Peter and John respectively. I suggest that, unlike John's attitudes toward the fact the Peter loses some money, the fact matters to Peter in the sense that Peter desires the opposite of the fact while John does not.²⁰ I see this solution as hitting the balance between two unreasonable extremes: letting the special status of Peter defined by his desires and neglecting the fact that Peter has a special relation to the fact that he loses some money.

If "it" refers to the fact that horse A wins the race, I am inclined to say that it is neither good nor bad. Even though we can anticipate and understand that John wants to regard it as a good state of affairs and Peter wants to regard it as a bad one, this is so only because in the current context the fact that horse A wins the race is too close to the fact that John wins some money and Peter loses some. But strictly speaking, they are different facts or states of affairs. States of affairs, different from events, are extremely finely-grained. Unfortunately, Foot neglects this feature of states of affairs for some reason.

²⁰ Here, it would be misleading and inaccurate even to say that John desires that Peter loses money. What John really desires is that John wins some money, which is not identical to the fact that Peter loses money. This is another example of the finely-grained nature of states of affairs as a kind of things that we can be responsible for.

My point is that a fact *may* imply a particular standard of evaluation as relevant and appropriate as its measure. For example, the fact that a human life is saved suggests that this fact should be evaluated based on the value of human life. The fact that a person regains her freedom should be evaluated based on the value of freedom and liberty. The fact that a person is injured should be evaluated based on the disvalue of harm. It is conceptually impossible to evaluate, say, the fact that a human life is saved based on the value of freedom and liberty. I do not mean that all facts can be evaluated in such an uncontroversial manner which is recognized and agreed by all reasonable people. Maybe the fact is clearly unrelated to anything that is intrinsically valuable, e.g., the fact that $1+1=2$ or the fact that Earth is round. Or perhaps, it is not clear which value if any is the appropriate standard of evaluation for some other facts, e.g., the fact that a baby is born and the fact that someone is getting married.²¹ It is also possible to conjoin two evaluable facts into a fact that is not clearly evaluable, e.g., the fact that a person is killed and the fact that two other people are saved. This complex fact that a person is killed and two others are saved does not seem to have a clear moral quality or evaluation, at least not in the context of ascribing responsibility. Maybe it is a problematic idea for someone to be responsible for the complex fact that a person is killed and two other people are saved. Maybe the logic rule of conjunction does not apply to responsibility. Imagine another person who is a friend to both the victim and the two survivors: how is she supposed to react to this complex responsibility? Should she always treat them separately or figure out a way to combine her reactive attitudes towards the responsible person? If the former is conceptually the only valid option, then what

²¹ That life is intrinsically valuable is most obvious in the sense that sustaining a life or preventing loss of a life is valuable. It seems less clear that creating a life is intrinsically valuable as well even though I am inclined to agree it is. To see this, just imagine the asymmetry of many people's attitudes towards the time before their existence and after.

does it mean to say that she holds a person responsible for the complex fact and what does it mean to say that a person is responsible for the complex fact?

The question may seem more complicated than I suggested in another respect. Some may want me to evaluate the fact that bin Laden was killed, which is not a complex fact. Although most reasonable people who disagree with what he stands for would judge that it is a good state of affairs, there might be some reasonable followers of him who would take it to be a bad state of affairs. But this does not show that it is equally impossible to speak of a good state of affairs and a good event. There is still a fundamental difference between fine-grained states of affairs and coarse-grained events. That is, the different evaluation of the state of affairs or the fact that bin Laden was killed is not merely because the fact in question leaves room for different evaluations. Rather, it is due to the fact that different people disagree on some other facts—moral or non-moral—about bin Laden.²² So the problem for evaluating states of affairs is not semantic in nature. But for events, even if people agree on everything about an event, they may still feel puzzled how to make an evaluation. An event, without one of its many descriptions specified, would always leave room for different evaluations.

3. An Objection: Two Faces of Moral Responsibility

A challenge arises for premise (1)—*X* must have a particular moral quality in order for an agent to be morally responsible for *X*—when we start to consider Watson’s discussion of two

²² My view that the moral quality is not relative to a person’s beliefs about value does not assume the truth of value objectivism. My view is just that the moral quality depends on the value that is appropriate. If the truth of value (for a person) is relative to a person’s beliefs, then the moral quality is also relative to the person’s beliefs. Which person? Obviously, the person that is involved in the ascription of responsibility, especially the one who is in the position to hold someone responsible. It will make things easier and simpler if the one who is in the position to be held responsible shares the same value beliefs. But I leave it open whether the practice of responsibility is possible for two people when individual relativism about values is true and they do not have the same value beliefs. But I have a strong intuition that it cannot work in such a unilateral way.

faces of responsibility: attributability and accountability. In his seminal paper “Two Faces of Responsibility” (1996), Watson contends that there are two senses of responsibility and, when we take an agent to be responsible for something, it might have different meanings. The first sense of responsibility is what he calls “attributability.” When something, say, one’s action, is “attributable” to an agent, it is dependent upon and expressive of the agent’s underlying judgments and evaluative assessments. In virtue of this, Watson argues, the agent is open to “appraisal that is therefore appraisal of an individual as an adopter of ends” (Watson 1996, pp. 229), and such appraisals will concern “the agent’s excellences and faults—or virtues and vices—as manifested in thought and action” (Watson 1996, pp. 231). In this attributability sense, to be responsible for *X* is to be the appropriate target of what he calls “aretaic appraisals.” In contrast, the second sense of responsibility concerns the practices of subjecting an agent to some appropriate moral reward and punishment when a certain requirement of control (or avoidability) is satisfied, which makes this sense of responsibility close to the concept of responsibility depicted by Strawsonianism.

One way to understand Watson’s two faces of responsibility is to remind ourselves of the distinction between being responsible and being held responsible. Strawson dissolves the distinction and gives a unified account for both notions. Strawsonians endorse the distinction and give an account of the notion of being held responsible that is more or less similar to Strawson’s while offering various metaphysical accounts for the notion of being responsible. The ledger view consists in a credit/debit-based account for the notion of being held responsible and, like Strawsonians, leaves open what metaphysical account for the notion of being responsible they want to adopt. Watson also adopts the distinction between being responsible and being held responsible. What makes his view different from all of the above is that he gives two

independent pairs of accounts for both notions. A self-disclosure sense of attributability paired with moral appraisal represents both being responsible and being held responsible as one face on responsibility. A control-based notion of accountability paired with more or less Strawson's reactive attitudes represents both being responsible and being held responsible as another face of responsibility.

The challenge for premise (1) is that attributability/aretaic appraisal, if treated as a genuine kind of appraisal of responsibility, seems to be incompatible with premise (1). If to be appraised as responsible can be merely to have something attributed to a moral fault or excellence in the agent in the sense that it discloses or expresses a moral vice or virtue of the agent, there is no reason to require appraisal of responsible to relate to a particular moral quality. Although an object or an event may have more than one moral quality under different descriptions, it seems plausible to take the object or event itself to be the expression of a moral vice or virtue.

To see this, imagine a creative but unskillful art student who creates a piece of artwork that is brilliant in terms of originality but awful in terms of skills. The artwork itself, as an object, seems to express the creativity of the student. The same artwork also expresses the unskillfulness of the student. Therefore, although those two characteristics of the student—creativity and unskillfulness—can be expressed by two relevant aspects of the artwork, they can also be expressed by the same artwork itself. That an object is “multifaceted” does not prevent the object from being an expression of one of many related features or characteristics of the individual responsible for it. Of course, this is not a case of events as expressions of moral vices or virtues. But the lesson applies equally.

Consider a difference between Watson's view of attributability and the view based on the "ledger" metaphor. According to the ledger view, when an agent is regarded as responsible for *X*, it is like putting a moral credit or debit—which is expressed by just one aspect of *X*—on the agent's ledger. What makes Watson's view of attributability different from the ledger view is that, instead of attributing a moral credit or debit to the agent, *X* itself is "attributed" to the agent. This practice of moral responsibility is most compatible with premise (1), but it seems a possible one.

The difference between the ledger view and Watson's attributability view is not superficial. Where the ledger view bases the moral relevance and significance of one's responsibility on some moral credit or debit of the object of responsibility, Watson's view seems to put the emphasis on the moral evaluation of the agent or the agent's character or evaluative commitments which is more or less detached from that a single instance of responsibility concerns. Thus the ledger view fits best with the practice of case-by-case moral blame or praise where each instance of being responsible for something calls for a certain moral reaction in accordance with the moral quality of this thing. This can be generalized to include accountability/reactive attitudes as a version of the Strawsonian approach. In contrast, Watson's view of attributability seems to fit better with the practice of moral blame or praise that depends on an *overall* evaluation of the agent or her moral character.

However, as some philosophers (e.g., Wolf 1990, Levy 2005, Smith 2008) have already argued, there are good reasons to doubt that appraisal of "attributability" is a genuine kind of appraisal of responsibility. In her *Freedom within Reason*, Wolf raises a challenge to what she calls "real self views" by arguing that their central notion is "merely causal (and therefore ethically superficial)" and that the moral appraisals associated with self-disclosure "are simply

descriptions of a thing's qualities and differ in kind from moral blame in the strict (deep?) sense" (Watson 1996, pp. 232). Levy's objection focuses on two reasons against what he calls "attributionism": one is that attributionism fails to "accommodate the relatively stringent epistemic conditions that any adequate theory of moral responsibility must recognize," and also fails to "accommodate the intuitively powerful distinction between *bad* agents and *blameworthy* agents" (Levy 2005, pp. 2).

Here, I do not plan to rehearse, evaluate, or defend those objections to Watson's view that attributability constitutes a genuine face of responsibility. My response to the objection in this section is that, even if attributability is a genuine face of responsibility, it does not undermine my main argument against speaking of responsibilities for actions as events, or at least the spirit of my argument.

What is the point or underlying spirit of my main argument against the view that we can be responsible for actions as events? Basically, it is the thought that it is semantically incomplete to say that a person is responsible for something when the relevant moral quality of this thing is not singled out. We cannot understand what it means for someone to be responsible in such a case. Just as we cannot understand what it means for something to occupy a space without having a particular shape. It is not merely that we do not or cannot know its shape but that the thing is of the sort that cannot have a particular shape.

Even if Levy's and Smith's arguments fail to convince us that Watson's attributability is not a genuine face of responsibility and, for the sake of argument, we grant that it is a genuine responsibility, it seems that it is semantically incomplete to speak of responsibilities for actions as events. That is, we do not understand what it means to attribute something to a moral fault or excellence in the agent in the sense that it discloses or expresses a certain moral vice or virtue of

the agent. Just as an event under different descriptions may have different moral qualities, an event could disclose or express different moral vices or/and virtues of the agent as well.

Let us consider the event of Caesar's being murdered by Brutus. If we say that Brutus is responsible for this event in the sense of Watson's attributability, what do we have in mind? It is not clear. That stabbing a friend like that discloses or expresses the unfaithfulness of Brutus? That assassinating a dictator discloses or expresses Brutus's bravery and loyalty to democracy? That murdering a person discloses or expresses Brutus's disregard of the law? Obviously, there are more things we might have in mind.

To conclude, in this chapter, I defend two premises of the responsibility-based argument: (1) in order for an agent to be responsible for X , X must have a particular moral quality, and (2) events are not things that can have a particular moral quality. They entail that events are not things that we can be responsible for. Since we can be responsible for performing actions, actions are not events.

CHAPTER 5: A DAVIDSONIAN CHALLENGE

The conclusion I want to establish in this dissertation is that actions are not events, which serves as an ontological objection to the CTA. In chapters 2-4, I have argued that there are three undesirable implications of the assumption that actions are events: (1) we cannot perform actions intentionally; (2) we cannot have abilities to perform actions; (3) we cannot be morally responsible for performing actions. Since we can perform actions intentionally, we can have abilities to perform actions, and we can be morally responsible for performing actions, we must reject the assumption that actions are events.

In this chapter, I am going to discuss a challenge to my argument that could be raised against my conclusion, especially on behalf of Davidson. The challenge is based on Davidson's view—following Anscombe—that an action is done intentionally by its agent *only under a certain description*. This seems to resolve the first undesirable implication because, according to Davidson, it does not make sense to say that a person performs an action intentionally in the first place. Intentionality is not a two-place relation between the agent and her action, but a three-place relation between the agent, her action, and a certain description of the action (Davidson 2001, essay 3). If we accept the plausibility of this response, then it seems to apply to the second and the third undesirable implications as well. Davidson could and should argue that ability and moral responsibility are both three-place relations as well. Instead of saying that we have abilities to perform particular actions, what he needs to say is that we have abilities to perform particular actions under certain descriptions. Similarly, instead of saying that we are morally responsible for particular actions, what he needs to say is that we are morally responsible for particular actions under certain descriptions.

This usage of “under a description” represents a more unified and consistent position that Davidson *should* adopt for human agency in general, not just for intentional action in particular. This is because, assuming that actions are events, the difficulty of performing actions intentionally, of being able to perform actions, and of being morally responsible for actions can all be traced to the particularity of actions. If an action is a particular, as I have shown in chapters 2-4, we find it impossible to intend to perform this particular action, to be able to improve the success rate of performing this particular action (without equally improving the chances of others), or to assign it a particular moral quality which is required for moral responsibility.

Let me use the burglar case to illustrate. Suppose you returned home and turned on the light, without realizing that a burglar in your home was also alerted by your turning on the light. Davidson claims that your action was intentional under the description “turning on the light” but unintentional under the description “alerting the burglar.” It does not make sense to say that your action was intentional, period. Likewise, it is the case that you were able to perform the action under the description “turning on the light.” It may be the case that you were also able to perform the action under the description “alerting the burglar” depending on whether an epistemic condition is built into the notion of ability. However, your action—assuming it is an event—will also acquire various descriptions mainly by sheer luck, such as “causing a unique pattern of air vibration,” but you were not able to perform the action under the description “causing a unique pattern of air vibration.” Also, you are morally responsible for the action under the description “turning on the light,” but not under the description “altering the burglar” or “causing a unique pattern of air vibration.”

To be clear, suggested general use of “under a description” language is only a proposal suggested by me for the benefit of those who wish to defend the view that actions are events. If

proponents of such a view are reluctant to adopt it, there seems to be no better solution to the difficulties regarding notions of abilities and moral responsibilities that are revealed in chapters 3-4. So, I plan not to say too much about why followers of Davidson should accept my proposal of using “under a description” more consistently. I will focus rather on why this solution only appears to be successful.

By claiming that the solution appears to be successful, what I mean is that a unified and consistent use of “under a description” language may seem to neutralize my attack on the view that actions are events by making it just another manifestation of the perennial dispute between nominalists and realists. To simplify the issue a little, realists tend to accept the existence of universals, while nominalists tend to reject the existence of universals but accept the existence of particulars. If the difference between Davidson and myself is merely that I want to introduce act-universals into our ontology while Davidson, as a nominalist, prefers not to allow universals but rather to speak of events under certain descriptions, then it appears that there can hardly be anything interesting and original about my ontological objection.¹ Without finally settling the nominalist-realist debate, there is no hope to convince a nominalist like Davidson to use the notion of act-universals instead of the notion of events under certain descriptions to account for the phenomenon of human agency.

Of course, what Davidson and proponents of the CTA claim goes beyond merely adopting the nominalist ontology of events under descriptions. They want to say that actions are events. I take for granted a few fundamental assumptions about the concept of actions. They are:

¹ I do not mean to deny that the debate between nominalists and realists might be a real one, not just a choice based on different philosophical “tastes.”

we can perform actions intentionally; we can have abilities to perform actions; and we can be morally responsible for performing actions. Without begging the question against Davidson and his followers, it seems equally plausible for a nominalist to insist that we can perform actions intentionally under certain descriptions; we can have abilities to perform actions under certain descriptions; and that we can be morally responsible for actions under certain descriptions. So, if the lesson that we learn from chapters 2-4—for a realist—is that actions are universals instead of particulars because only universals can play the role that the notion of actions is supposed to play and relates to other concepts—such as intentionality, ability, moral responsibility—that are essential to a theory of agency, for a nominalist, the lesson becomes that actions are essentially sensitive to descriptions of events because only under certain appropriate descriptions, can action-related events relate to intentionality, ability, and moral responsibilities and thus reveal the phenomenon of human agency. In short, events become things that qualify as actions only under certain descriptions.

It is a less serious matter if Davidson still wants to insist that actions are events. That seems to be a mere terminological difference. What really matters is the lesson that attributing something to one's agency is an intensional context that essentially depends on this thing's descriptions. Davidson does not miss this crucial point entirely. As he puts it, “the *criterion* of agency is, in the semantic sense, intensional” because “a person is the agent of an event if and only if there is a description of what he did that makes true a sentence that says he did it intentionally” (Davidson 2001, pp. 47-8). That explains why, to him and his followers, a causal analysis of actions is just a causal analysis of intentional actions. But here we encounter an important question about the whole project of providing a causal account or analysis of the notion of actions. Since the causal relation or event-causal relation is well known to be

extensional rather than intensional, is its success even possible? To put it another way, if an appropriate description is essential for something to be counted as an action, how can we have a causal analysis of actions when causal relations are not dependent on descriptions at all?²

My questions above thus suggest that the difference between Davidson and myself may go deeper than a preference between nominalist and realist ontology based on philosophical tastes. Even if we accept a nominalist way of speaking of actions and go easy on whether actions are events or events under certain descriptions, the CTA still faces a potential difficulty of analyzing a description-sensitive notion by appealing to causal relations that are not description-sensitive. In my view, this difficulty is a manifestation of the initial lesson—when a realist ontology is adopted—that actions cannot be particulars but universals.

In response to my argument, Davidson might point out that a causal analysis of actions can accommodate the fact that talk about action is description-sensitive and thus in an intensional context. Although a causal analysis of actions could be formulated differently, such an analysis would contain two conditions that, together, constitute a sufficient and necessary condition of acting intentionally. First, there are some mental states or events of the agent that justify the agent's alleged action (at least from the perspective of the agent)—the rationalization condition. Second, the alleged action is caused by those mental states or events. Davidson may argue that the relation of rationalization is sensitive to descriptions—the causation condition. To see this, let us use Davidson's explication of his notion of primary reason as an example:

R is a primary reason why an agent performed the action *A* under the description *d* only if *R* consists of a pro attitude of the agent towards actions with a certain property, and a belief of the agent that *A*, under the description *d*, has that property (Davidson 2001, pp. 5).

² Of course, as Davidson emphasizes in many places, causal explanations are intensional and dependent on descriptions (e.g., of facts). But there is a difference between causal explanations and causal relations.

Clearly, a primary reason can rationalize one's action only under a certain appropriate description. As Davidson says:

Since reasons may rationalize what someone does when it is described in one way and not when it is described in another, we cannot treat what was done simply as a term in sentences like 'My reason for flipping the switch was that I wanted to turn on the light'; otherwise we would be forced to conclude from the fact that flipping the switch was identical with alerting the prowler, that my reason for altering the prowler was that I wanted to turn on the light (Davidson 2001, pp. 5).

Since the rationalization condition is sensitive to descriptions (of the alleged action), it seems to handle my challenge. This could explain why Davidson himself does not find it necessary to discuss the intensional-extensional gap that appears to exist in such a causal analysis. But I believe the rationalization condition being description-sensitive cannot bridge the intensional-extensional gap inherent to any causal account of actions. Two reasons will be provided to support my belief.

First, as noted at the beginning, there is a well-known problem with the attempt to causally analyze the notion of actions—deviant causal chains or wayward causal chains—the intractability of which seems to corroborate the existence of the intensional-extensional gap. A classic example of causal deviance is presented by Davidson in his paper "Freedom to Act":

A climber might want to rid himself of the weight and danger of holding another man on a rope and he might know that by loosening his hold on the rope he could rid himself of the weight and danger. This belief and want might so unnerve him as to cause him to loosen his hold, and yet it might be the case that he never *chose* to loosen his hold, nor did he do it intentionally (Davidson 2001, pp. 79).

This example shows that being caused and rationalized by some mental states or events does not suffice for action or intentional action. Many efforts have been made to modify the causal analysis in order to save the causal approach to agency (e.g., Bishop 1989, Mele 1992, Schlosser 2007). But none of them has convincingly solved the difficulty of deviant causal chains. Maybe, it is just like what Wilson concludes in his *The Intentionality of Human Action*:

“that the evidence points to more than infelicity or incompleteness in the various causalist proposals – it points, that is, to a global breakdown in the whole project of reduction” (Wilson 1989, pp. 258).³

Let us assume for the sake of argument that we are justified in believing in the intractability of deviant causal chains. Why do I claim that the problem corroborates the existence of the intensional-extensional gap that I take to be inherent to the causal analysis project? To see this, consider the remedy proposed by Schlosser (2007). Schlosser’s solution fits into the standard way of saving the causal analysis by adding “in the right way” to the causation condition and specifying what the right way means. According to him, the problem of deviant causal chains can be solved if the CTA requires that the relevant mental states or events cause the agent’s behavioral events *in virtue of the intentional contents* of those mental states or events. The climber example apparent fails to satisfies this new condition.

The reason states cause the movement, because they cause the nervousness, which causes the movement. But the reason states do not cause the movement in virtue of their contents, because the event of nervousness, trivially, does not cause the movement in virtue of content (Schlosser 2007, pp. 192).

This simple modification of the causal condition seems straightforward. However, it depends on a crucial assumption about the nature of causality: intentional content is causally relevant and explanatory. This means that causal relations (or the truth of those sentences reporting them), instead of being entirely language-independent, are affected by how the mental

³ Speaking of “reduction,” Davidson once gave a response to the difficulty of deviant causal chain by first admitting the intractability of the difficulty and then claiming that the difficulty does not undermine the CTA since “the reduction is not definitional but ontological” and “the ontological reduction, if it succeeds, is enough to answer many puzzles about the relation between the mind and the body, and to explain the possibility of autonomous action in a world of causality” (Davidson 2001, pp. 88). But here, I agree with Bishop’s criticism that merely clarifying the ontological nature of the reduction is not enough to dismiss the deviant causal chain problem because “If it is impossible to state event-causal conditions that are sufficient for action then it is impossible to provide a CTA analysis to back up CTA’s core ontological thesis” (Bishop 1989, pp. 104).

events and the behavioral events are described. Setting aside the fact that this assumption would probably be rejected by many philosophers—for example, and especially Davidson—it at least shows that the deviant causal chain can be solved only if we do not take causal relation to be purely extensional.

Let us see how Davidson understands the deviant causal chains. In his “Psychology as Philosophy,” he suggests that the source of the problem is that rationality enters into not only the rationalization condition but also into the causation condition of the causal analysis:

Two ideas are built into the concept of acting on a reason (and hence, the concept of behaviour generally): the idea of cause and the idea of rationality. A reason is a rational cause. One way rationality is built in is transparent: the cause must be a belief and a desire in the light of which the action is reasonable. But rationality also enters more subtly, since the way desire and belief work to cause the action must meet further, and unspecified, conditions (Davidson 2001, pp. 233).

Rationality implies description-sensitivity. There is no question of rationality in an extensional context such as causal relations. If, as Davidson says, rationality enters not only into the rationalization part but also the causation part of the notion of action—which means a right way of causing must be a rational way—then surely it is not good enough to bridge the intensional-extensional gap by putting together a purely intensional rationalization condition and a purely extensional causation condition.

The first reason—based on the deviant causal chains problem—has its weakness. Although illuminating, it is mere corroborating evidence and relies upon the apparent plausibility of the diagnosis of Schlosser or Davidson on the issue. Here, I would like to point to another reason why the rationalization condition being description-sensitive cannot bridge the intensional-extensional gap. It seems to me that the rationalization condition should not be a part of the causal analysis at all. Adding this condition is a mistake due to the misleading way of how Davidson raises his famous challenge against anti-causalists about action explanation. Anti-

causalists tend to think that action explanations are not causal explanations. The basic idea of Davidson's challenge is how to provide an account of the notion of acting for a reason without appealing to a causal relation between one's relevant mental states or events and one's relevant behavior. The challenge seems a difficult one considering that a person can have more than one reason that rationalizes her behavior but the behavior is done only for one of them. Here is an example offered by Mele to illustrate such a situation:

Al has a pair of reasons for mowing his lawn this morning. First, he wants to mow it this week and he believes that this morning is the most convenient time. Second, Al has an urge to repay his neighbor for the rude awakening he suffered recently when she turned on her mower at the crack of dawn and he believes that his mowing his lawn this morning would constitute suitable repayment. As it happens, Al mows his lawn this morning only for one of these reasons. In virtue of what is it true that he mowed his lawn for this reason and not the other, if not that this reason (or his having it), and not the other, played a suitable causal role in his mowing his lawn? (Mele 1997a, pp. 240).

As Mele puts it, the gist of the question is this: "in virtue of what is it true that a person acted in pursuit of a particular goal?"⁴ The way this challenge is raised suggests a common ground between causalists and anti-causalists. That is, the rationalization condition should be part of an account or analysis of actions. This is partly shown by Davidson's approving remark—inspired by the later Wittgenstein—on how to make an action seem rationalized and reasonable:

And there is no denying that this is true: when we explain an action by giving the reason, we do redescribe the action; redescribing the action gives the action a place in a pattern, and in this way the action is explained (Davidson 2001, pp. 10).

Since, in Davidson and his followers' opinion, this challenge cannot be met without appealing to a causal relation, the causation condition—no matter how it should be specified—should be added to an account or analysis of actions. Naturally, they would think that this should be done while keeping the original rationalization condition in place. But I want to argue that the

⁴ Mele 2003b, p. 39.

rationalization condition should not be part of a causal analysis of action. To see this, consider this question: Is it possible for some mental events or states that, under any description, cannot rationalize one's behavior to cause the behavior in the right way? Take Mele's mowing case as an example. It is clear that both reasons could separately cause the action of mowing in the right way, even without specifying what "right way" really means beyond our pre-theoretical intuition. Now, let us suppose that Al also desires or intends to stay in bed that morning, can such a mental state (or Al's having it) causes him to mow the lawn that morning? Maybe. But, however it might cause Al's mowing the lawn, I cannot see that it would then be done *in the right way*. I suggest that this is because non-rationalizing mental reasons simply cannot cause one's behavior in the right way. Thus the rationalization condition is already inherent and essential to the right way of causing. Therefore, there seems no need to keep the rationalization condition—separate from the causation condition—that seems redundant and extrinsic. Of course, it is probably unharmed to keep it. But if Davidson is correct in thinking that rationality—thus description-sensitivity—should be subtly *built into* the causation condition, merely adding a rationalization condition as a "patch" on the causal analysis does not suffice for bridging the intensional-extensional gap.

To conclude, even if a nominalist ontology and the language of "under a description" are adopted by Davidson as a response to my chapters 2-4, the crucial lesson we learn from those chapters is still valid. For a realist, the lesson is an ontological preference to understand actions as universals rather than particulars. And since act-universals are not something that can be caused in a world of event-causality, the CTA fails necessarily. For a nominalist, essentially the same lesson gets readdressed and becomes a comment on the description-sensitivity of agency,

which makes agency unfit to be analyzed by purely event-causal relations. Thus, the CTA still fails necessarily.

CHAPTER 6: CONCLUSION

In this dissertation, I raised an ontological objection to the event-causal theory of action that actions are not events. Since actions are not events, any attempt to account for the notion of action by appealing to how actions enter into the event-causal relations necessarily fails. The ontological objection consists in three separate arguments sharing the same spirit, each of which gives the same conclusion. They are defended in chapters 2-4:

The intention-based argument (chapter 2):

- (1) We can perform actions intentionally.
- (2) Events are not things that we can bring about intentionally.
- (3) Therefore, actions are not events.

The ability-based argument (chapter 3):

- (1) We can have abilities to perform actions.
- (2) Events are not things that we can have abilities to bring about.
- (3) Therefore, actions are not events.

The responsibility-based argument (chapter 4):

- (1) We can be morally responsible for performing actions.
- (2) Events are not things that we can be morally responsible for.
- (3) Therefore, actions are not events.

For each argument, premise (1) is a basic assumption I take for granted. So, the focus of each argument is how to show the truth of its corresponding premise (2). Hopefully, I have successfully shown that, due to the particularity of events, it is impossible for human beings to bring about events intentionally, to have abilities to bring about events, or to be morally responsible for events.

In Chapter 5, I anticipate and reject a Davidsonian response that appears to neutralize the force of my ontological objection to the CTA. The response is that the difference between followers of the CTA and me is just another manifestation of the perennial dispute between nominalists and realists. I argue it is not. Even if nominalists and realists have different expressions of what they can learn from my ontological objection, the message is the same in spirit.

What is the message or lesson of my ontological objection, then? Do I deny that actions can be performed in a world of event-causality where the only kind of causation is events causing events? Not necessarily. I believe actions can be performed and they do not involve anything foreign to the ontology of natural science (e.g., irreducible agent-causation). On this point, I am not different at all from Davidson and followers of his spirit (e.g., Bishop 1989) when their theories can be seen as various attempts to explain “the possibility of autonomous action in a world of causality” (Davidson 2001, pp. 88) or “the possibility of accommodating *actions* within the natural universe” (Bishop 1989, pp. 2). This project of reconciliation of (free) agency and scientific causality can at least be traced back to Kant’s first and second *Critique*.

I believe a causal account of action misunderstands the nature of human agency or action. But Davidson challenges: How can we find a non-causal account of human action without appealing to causal relations? Some philosophers (e.g., Scott Sehon, George Wilson) try to develop a teleological account of actions. I find them unsuccessful in replying to Davidson’s challenge (see Mele 2010). An adequate reply to Davidson’s challenge needs to explicit a non-causal relation of “production” between one’s reason and one’s action. Davidson’s mistake is that he assumes that the only naturalistic kind of relation of “production” is event-causation.

Teleological theorists' mistake is that they neglect the essential role played by such a relation of "production" in human agency.

So, what we should look for is a kind of "non-causal bringing about" (Taylor 1964). Inspired by Frankfurt's 1978 paper "The Problem of Action," as a final remark, I would like to offer a brief explanation of what such a "non-causal bringing about" might look like. How can something belong to an agent in terms of "production" without being caused by this agent's reasons or this agent herself? Consider an example offered by Frankfurt:

A driver whose automobile is coasting downhill in virtue of gravitational forces alone may be entirely satisfied with its speed and direction, and so he may never intervene to adjust its movement in any way, this would not show that the movement of the automobile did not occur under his guidance (Frankfurt 1978, pp. 160).

In this case, the movement of the car (or its reaching the desired destination) is guided or "produced" by the driver even though the driver (either her body or her mental states) does not cause the car's movement.

Frankfurt's case may seem unsatisfactory in several ways. First, this driver is extremely lucky to find it unnecessary to intervene in order to guide her car to move as she desires or intends. But we are not so lucky in most situations. In reply, consider this situation. I desire to draw a dinosaur on a wall based on my mental image of a dinosaur. And a person, for some reason, is ejecting millions of small ink particles towards the wall in a random pattern. Based on an advanced control mechanism, I am able to do two things: first, I quickly tell which ink particle will land on the right spot for the sake of drawing the dinosaur and which ink particle will not; second, I accurately intercept those "wrong" particles but let those "right" particles hit the wall. As a result, a drawing of the dinosaur I have in my mind is produced on the wall. The drawing is undoubtedly produced by me, although I have no causal relation to either the wall or those ink particles that constitute the drawing. Unlike Frankfurt's driver, my non-causal

production of the drawing is not due to any substantial luck. It is rather based on relevant skills and abilities.

Frankfurt's case can also seem unsatisfactory because the car's movement is not the driver's action while the driver's bodily movement—sitting still—is arguably caused by the driver's desire or intention (Mele 1997b, pp. 138). So, the complaint is that the basic action—the driver's sitting still—is still causally explained by her relevant mental event or state and the car's movement is merely the consequence of her action. Since actions seem to always consist in agents' basic bodily movements (Davidson 2001, essay 3) and those bodily movements—as basic actions—must be caused by something inside agents' bodies, actions must also be caused by something inside agents' bodies (most likely, mental events or states).

This complaint misses the point of the original question: How can *something* belong to an agent in terms of “production” without being caused by this agent's reasons or this agent herself? The car's movement belongs to the driver in the sense that the driver is the agent of the car's movement. Whether the car's movement is an action of the driver or not is less important. As long as happenings in the world can be attributed to human agency, it seems harmless even if nothing in the world can be called “action.” Also, even if there must be a basic bodily movement each time we act or guide something to happen, it is not true that this basic bodily movement is caused by the agent's mental event or state. One's bodily movement, no matter how basic it seems to be, must be guided by Frankfurt's causal mechanism rather than simply be caused by mental events. Take my lifting a coffee mug as an example. If my bodily movement is purely caused by my prior intention to lift it, then I probably will lift it in a wrong way, maybe using too little force or too much because, most importantly, I have not felt the weight of the mug until I

am in the middle of the process of lifting it. I need to adjust my bodily movement and guide it to move in such a way that the mug is lifted appropriately.

Is there anything that is caused by my intention to lift the coffee mug? Yes, and too many. My intention causes the first phase of my hand's movement and many small adjustments—particular events—during the whole process. That is why it is called by Frankfurt a *causal* mechanism. Sometimes, when we are extremely lucky as Frankfurt's driver, no adjustment or even the first phase of the movement needs to be caused in order to act. But it does not mean that no agency is working in such a lucky situation. The agency lies in the guidance of the car. Likewise, when there are many adjustments in most situations, the agency still lies in the guidance rather than causations. That explains why a non-causal account of agency is not necessarily without the existence of causality. To put it another way, many small adjustments are caused by my intention. But each of them is not yet an action or even a bodily movement. It is puzzling how they can constitute an action when they are put together into a bodily movement.

A virtue of this picture of human agency is that it easily accommodates that actions are not particulars but universals. By engaging in the causal mechanism, I can guide my hand and the coffee mug into a specified—not a particular—status: the mug being lifted by my hand appropriately (not too fast, not too slow, not too high, not too low). This can be realized in many slightly different ways. In other words, what is *guided* to happen is not any particular event or bodily movement but the exemplification of an act-universal. There are some particular events that are caused to happen, but they are not *guided* to happen. Thus, they are beyond the scope of our agency and control.

For the question of how human agency works in making choices, this non-causal modal may have its most interesting application. It is difficult to explain how a person exercises

substantial control over a choice to be made between two incompatible things that she both desire. The difficulty is *not* that, if one desire is overwhelmingly stronger than the other, the person has no control over the result. The difficulty is rather that, when one desire is not significantly stronger than the other, it is still puzzling how the choice is up to the agent. It seems the result—forming an intention to pursue one of them—is caused and determined solely by those two desires, either deterministically or indeterministically. Merely saying that she is identified with one of the desires by forming a higher-order desire to will it (Frankfurt 1971, 1987) does not alleviate the worry that agency disappears in the process of making such a choice. What we look for is a metaphysical connection between the agent and her choice, not a psychological one or an ethical one.

What if the magic moment of exercising agency is not when the intention is caused to form? What if an agent exercises control over her choice (forming an intention) not by causing it to happen, but by allowing or disallowing it to continue to exist until being executed? There is often, if not always, a certain interval between one's forming an intention and one's acting on the intention, which gives the agent some time to refrain from acting and forms an intention again.¹ This second intention-forming will still be caused by the desires, not by the agent in any significant sense. Whether the same (kind of) intention is formed or not is not up to the agent. What is up to the agent is whether she allows or disallows the second intention to exist and be executed. By engaging in this mechanism of guidance, the agent obviously has substantial control over the *final* intention that makes an influence in the outside world. In this picture, one's agency does not work at the moment of one's forming an intention. It works after the intention is

¹ This could explain why a person should be less responsible for passion killing. The reason is not that the passion is too strong to overcome. The reason is that passion killing happens too quickly and the interval between forming the intention and acting on the intention is too short for the agent to overcome the passion.

formed. Or more precisely, it works during a period of time that may contain more than one instances of intention-forming and the intervals in between.

For such a mechanism to work, the agent needs to identify with one of the first-order desires so that she has some idea about which intention is to be rejected and which intention is to be welcomed. In this sense, Frankfurt's idea reveals its metaphysical significance in a theory of how an agent exercises control. Identifying with one desire is not merely judging it to be good. More importantly, it consists in one's deep inclination to prefer such a desire in the mechanism illustrated above. And this modal of control is not about how the intention is caused but about how the causation of intention is regulated and guided to the direction that the agent is identified with. Because I take one's allowing or disallowing a certain intention as a result of her inclination rather than a decision, I do not think this modal of agency can explain *free* agency in the libertarian sense. With enough times of intention-forming, even a morally corrupted person could have decided to do the right thing. But it is a matter of habit (and luck) how hard a person would struggle with herself for the sake of morality. So, this modal of guidance control illustrates how agency, although not free agency in a really robust sense, works in making choices between competing desires.

REFERENCES

- Adams, F. (1986). Intention and Intentional Action: The Simple View. *Mind & Language*, 1(4), 281-301.
- Adams, F. (2007). Trying with the Hope. In M. Timmons, J. Greco, & A. Mele (Eds.), *Rationality and the Good*. Oxford: Oxford University Press.
- Adams, F., & Mele, A. (1989). The Role of Intention in Intentional Action. *Canadian Journal of Philosophy*, 19(4), 511-531.
- Aguilar, J. H., & Buckareff, A. A. (Eds.). (2010). *Causing Human Actions: New Perspectives on the Causal Theory of Action*. Cambridge, MA: MIT Press.
- Alvarez, M. (1999). Actions and Events: some semantical considerations. *Ratio*, 12(3), 213-239.
- Anscombe, G. E. (2000). *Intention*. Cambridge, MA: Harvard University Press.
- Aune, B. (1977). *Reason and Action*. Dordrecht, Holland: D. Reidel Pub.
- Bach, K. (1980). Actions Are Not Events. *Mind*, 89(353), 114-120.
- Bennett, J. (1988). *Events and Their Names*. Indianapolis: Hackett Pub. Co.,.
- Bennett, J. (1991). Precis of Events and Their Names. *Philosophy and Phenomenological Research*, 51(3), 625.
- Bishop, J. (1983). Agent-Causation. *Mind*, 92(365), 61-79.
- Bishop, J. (1989). *Natural Agency: An Essay on the Causal Theory of Action*. Cambridge: Cambridge University Press.
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Bratman, M. E. (1999). *Intention, Plans, and Practical Reason*. Miami, FL, United States: Centre for the Study of Language & Information.
- Chang, R. (1997). Introduction. In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason*. Cambridge, Mass: Harvard University Press.
- Chisholm, R. (1970). Events and Propositions. *Noûs*, 4(1), 15.
- Costa, M. J. (1986). Acting Intentionally and Minimal Abilities. *Analysis*, 46(3), 144-147.
- Davidson, D. (1963). Actions, Reasons, and Causes. *The Journal of Philosophy*, 60(23), 685-700.
- Davidson, D. (1985). Reply to Quine on Events. In E. LePore & B. P. McLaughlin (Eds.), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Oxford: B. Blackwell.

- Davidson, D. (2001). *Essays on Actions and Events*. Oxford: Clarendon.
- Donnellan, K. (1968). Putting Humpty Dumpty Together Again. *The Philosophical Review*, 77(2), 203-215.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, H. (1971). Freedom of the Will and the Concept of a Person. *Journal of Philosophy*, 68, pp. 5-20.
- Frankfurt, H. (1978). The Problem of Action. *American Philosophical Quarterly*, 15(2), 157-162.
- Frankfurt, H. (1987). Identification and Wholeheartedness. In F. Schoeman (Eds.), *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*. Cambridge: Cambridge University Press.
- Foot, P. (1985). Utilitarianism and the Virtues. *Mind*, XCIV(374), 196-209.
- Geach, P. T. (1956). Good and Evil. *Analysis*, 17(2), 33.
- Ginet, C. (2014). Can an Indeterministic Cause Leave a Choice Up to the Agent. In D. Palmer (Eds.), *Libertarian Free Will: Contemporary debates*. New York: Oxford University Press.
- Ginet, C. (2016). Reasons Explanation: Further Defense of a Non-Causal Account. *The Journal of Ethics*, 20(1-3), 219–228.
- Goldman, A. I. (1970). *A Theory of Human Action*. Englewood Cliffs, NJ: Prentice-Hall.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Harman, G. (1976). *Practical Reasoning*. *The Review of Metaphysics*, 29(3), 431-463.
- Harman, G. (1986). *Change in View*. Cambridge, Mass: MIT Press.
- Hegel, G. W., & White, A. (2002). *The Philosophy of Right*. Newburyport, MA: Focus Pub./R. Pullins.
- Honoré, A. M. (1964). Can and Can't. *Mind*, LXXIII(292), 463-479.
- Horgan, T. (1978). The Case Against Events. *The Philosophical Review*, 87(1), 28-47.
- Hornsby, J. (1980). *Actions*. London: Routledge and Kegan Paul.
- Hornsby, J. (2004). Agency and Actions. *Royal Institute of Philosophy Supplement*, 55, 1-23.

- Hornsby, J. (2010). Trying to Act. In O'Connor, T., & Sandis, C. (Eds.). *A Companion to the Philosophy of Action*. John Wiley & Sons.
- Hornsby, J. (2012). Actions and Activity. *Philosophical Issues*, 22(1), 233-245.
- Katzoff, C. (1989). Intentional Action—Sometimes a Matter of Luck. *Philosophical Investigations*, 12(3), 234-242.
- Kim, J. (1976). Events as Property Exemplifications. *Action Theory*, 159-177.
- Levy, N. (2005). The Good, the Bad, and the Blameworthy. *Journal of Ethics and Social Philosophy*, 1(2), 1-16.
- Lombard, L. B. (1982). Events and the Essentiality of Time. *Canadian Journal of Philosophy*, 12(1), 1-17.
- Maier, J. (2013). The Agentive Modalities. *Philosophy and Phenomenological Research*, 90(1), 113-134.
- McCann, H. (1987). Rationality and the Range of Intention. *Midwest Studies in Philosophy*, 10(1), 191-211.
- McDowell, J. (2011). Some Remarks on Intention in Action. *The Amherst Lecture in Philosophy*, 6, 1-18.
- McIntyre, A. (2014). Doctrine of Double Effect. In Edward N. Zalta (Eds.). *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition).
- Melden, A. I. (1961). *Free Action*. London: Routledge & Paul.
- Mele, A. R. (1989). She Intends to Try. *Philosophical Studies*, 55(1), 101–106.
- Mele, A. R. (1992). *Springs of Action: Understanding Intentional Behavior*. New York: Oxford University Press.
- Mele, A. R. (1997a). Agency and Mental Action, *Philosophical Perspectives*, 11, 231-49.
- Mele, A. R. (1997b). Passive Action. In G. Holmstrom-Hintikka & R. Tuomela (Eds.), *Contemporary Action Theory*, Dordrecht: Kluwer.
- Mele, A. R. (2003a). Agents' abilities. *Noûs*, 37(3), 447-470.
- Mele, A. R. (2003b). *Motivation and Agency*. Oxford: Oxford University Press.
- Mele, A. R. (2010). Teleological Explanations of Actions: Anticausalism versus Causalism. In J. Aguilar & A. Buckareff (Eds.), *Causing Human Actions: New Perspectives on the Causal Theory of Action*. Cambridge, MA: MIT Press.
- Mele, A. R., & Moser, P. K. (1994). Intentional Action. *Noûs*, 28(1), 39.

- Mele, A., & Sverdlik, S. (1996). Intention, Intentional Action, and Moral Responsibility. *Philosophical Studies*, 82(3), 265-287.
- Montague, R. (1969). On the Nature of Certain Philosophical Entities. *Monist*, 53(2), 159-194.
- Moore, M. (2010). Renewed Questions about the Causal Theory of Action. In J. Aguilar & A. Buckareff (Eds.), *Causing Human Actions: New Perspectives on the Causal Theory of Action*. Cambridge, MA: MIT Press.
- Nagel, T. (1986). *The View from Nowhere*. New York: Oxford University Press.
- O'Shaughnessy, B. (1980). *The Will, ii*. Cambridge: Cambridge University Press.
- Oakley, J. (1992). *Morality and the Emotions*. London: Routledge.
- Peacocke, C. (1979). Deviant Causal Chains. *Midwest Studies in Philosophy*, 4(1), 123-155.
- Peacocke, C. (1985). Intention and Akrasia. In B. Vermazen & M. Hintikka (Eds.), *Essays on Davidson*, Oxford: Clarendon Press.
- Pereboom, D. (2001). *Living Without Free Will*. Cambridge: Cambridge University Press.
- Quine, W. (1985). Events and reification. In E. LePore & B. P. McLaughlin (Eds.), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Oxford: B. Blackwell.
- Schlosser, M. E. (2007). Basic Deviance Reconsidered. *Analysis*, 67(3), 186-194.
- Schroeder, S. (2001). The Concept of Trying. *Philosophical Investigations*, 24(3), 213-227.
- Smith, A. M. (2007). On Being Responsible and Holding Responsible. *The Journal of Ethics*, 11(4), 465-484.
- Smith, A. M. (2008). Control, Responsibility, and Moral Assessment. *Philosophical Studies*, 138(3), 367-392.
- Smith, M. (2010). On Normativity. *Analysis*, 70(4), 715-731.
- Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language* (Vol. 626). Cambridge university press.
- Steward, H. (1997). *The Ontology of Mind: Events, Processes, and States*. Oxford: Clarendon Press.
- Steward, H. (2012). Actions As Processes. *Philosophical Perspectives*, 26(1), 373-388.
- Strawson, P. (1962/2003). Freedom and Resentment. *Proceedings of the British Academy*. Reprinted in G. Watson (Eds.). *Oxford Readings in Philosophy: Free Will*. Oxford: Oxford University Press.

- Sverdlik, S. (1996). Consistency among Intentions and the 'Simple View'. *Canadian Journal of Philosophy*, 26(4), 515–522.
- Taylor, C. (1964). *The Explanation of Behaviour*. London: Routledge and Kegan.
- Thompson, M. (2008). Naïve Action Theory. In M. Thompson, *Life and Action: Elementary Structures of Practice and Practical Thought*, Cambridge, MA: Harvard University Press, 85–146.
- Thomson, J. J. (1977). *Acts and Other Events*. Ithaca, NY: Cornell University Press.
- Thomson, J. J. (1997). The Right and the Good. *The Journal of Philosophy*, 94(6), 273.
- Trenholme, R. (1978). Doing Without Events. *Canadian Journal of Philosophy*, 8(1), 173–185.
- Vance, C. (2016). The Recycling Problem for Event Individuation. *Erkenntnis*, 81(1), 1–16.
- Velleman, J. D. (1992). What Happens When Someone Acts? *Mind*, 101(403), 461–481.
- Vihvelin, K. (2013). *Causes, Laws, and Free Will: Why Determinism Doesn't Matter*. New York: Oxford University Press.
- Wallace, R. J. (1994). *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Watson, G. (1996). Two Faces of Responsibility. *Philosophical Topics*, 24(2), 227–248.
- Wilson, G. (1989). *The Intentionality of Human Action*. Stanford: Stanford University Press.
- Wittgenstein, L., & Anscombe, G. E. M. (1973). *Philosophical investigations* (3rd ed.). Oxford: Blackwell Publishers.
- Wolf, S. R. (1990). *Freedom Within Reason*. Oxford University Press.
- Wood, A. W. (1999). *Kant's Ethical Thought*. Cambridge: Cambridge University Press.
- Zhu, J. (2010). On the Principle of Intention Agglomeration. *Synthese*, 175(1), 89–99.
- Zimmerman, M. J. (1988). *An Essay on Moral Responsibility*. Totowa, NJ: Rowman & Littlefield.
- Zimmerman, M. J. (1999). In Defense of the Concept of Intrinsic Value. *Canadian Journal of Philosophy*, 29(3), 389–409.

VITA

Jiajun Hu received a Bachelor of Engineering in Mechanical Engineering and Automation from Shanghai Jiao Tong University in 2005. He received a Master of Arts in Philosophy from Fudan University in 2012. He is currently finishing his PhD in philosophy at the University of Tennessee, Knoxville. After its completion, he will study business analytics at Arizona State University.