Agentially Controlled Action: Causal, not Counterfactual

Mere capacity views hold that agents who can intervene in an unfolding movement are performing an agentially controlled action, regardless of whether they do intervene. I introduce a simple argument to show that the noncausal explanation offered by mere capacity views fails to explain both control and action. In cases where bodily subsystems, rather than the agent, generate control over a movement, agents can often intervene to override non-agential control. Yet, contrary to what capacity views suggest, in these cases, this capacity to intervene does not amount to agential control or action. I illustrate this with a case study of how passive breathing, a mere behavior, is misclassified by mere capacity views. I end by revisiting the central alternative to mere capacity views: causal control views. Advances in our understanding of how agents exert control over unfolding movements indicate that the nature of control is characterized by ubiquitous, small-scale causal interventions.

Introduction

A hallmark of agency is agentially controlled movement. But what does that mean? Must one constantly influence an unfolding movement to be in control, or is it enough for the agent to be ready to intervene when necessary? This article argues against the noncausal *mere capacity views* of agentially controlled action, according to which agents control movements when they *could* (meaningfully) causally intervene in them. It then argues for the alternative *causal control views* of agentially controlled action, according to which agents only control movements when they do in fact intervene in them.

I argue that mere capacity views are wrong. When agents could — but do not — intervene in movements that are already under bodily control, mere capacity views fail to distinguish between bodily and agential control. Consequently, they wrongly ascribe these movements action status. I illustrate this with a case study of active and passive breathing, two modes of breathing that mere capacity views cannot differentiate. Active breathing is caused by the agent and, intuitively, agentially controlled action. Passive breathing is produced by a bodily subsystem and, intuitively, mere behavior. But agents have the capacity to meaningfully intervene not just in active breathing, but in passive breathing as well. Because of this capacity to intervene, mere capacity views wrongly classify passive breathing as agentially controlled action. Because mere capacity views generate such wrong verdicts, the mere capacity of an agent to intervene in a movement does not successfully explain control nor action.

This paper is divided into four parts. Part 1 situates work within the theory of action in terms of a threefold distinction involving mere behavior, action, and intentional action. Part 2 turns to the main counterexamples considered in this essay. I present cases in which bodily subsystems are the source of mere behavior, but the agent always stands poised to intervene. The case of breathing, one such movement, is considered in detail. Part 3 considers numerous objections to the counterexamples. Part 4 advances an initial sketch of an alternative picture of agential control that avoids the problems of mere capacity views.

1. Background of the Debate

Most philosophers care about control mainly because of its potential role in explaining action. With this idea in mind, a threefold distinction will be helpful to situate the explanatory aim of theories discussed in this paper. One central aim of action theory is finding necessary and sufficient conditions to specify when a movement is mere behavior, action, or intentional action. On the one hand, there are things that merely happen to agents, such as their digestion or the dilation of their pupils. These are often called *mere behavior*s, and typically generated by, and attributed to, bodily subsystems of the agent. On the other hand, there are things agents do, such as writing a play or fiddling with one's hair, which are attributed to the agent. Among these actions, some are *intentional actions*, things you intentionally do. For example, writing a play is something you intentionally do. Fiddling with your hair is something you do, but it might not always be something you do intentionally. If that is right, not all actions are intentional actions (Frankfurt 1978, Steward 2009, Levy 2013, Hyman, 2015). Others believe that all actions are intentional action (Davidson 1963). This paper remains neutral on whether all actions are intentional actions. Rather, I will argue that the prominent mere capacity views misclassify movements that are mere behavior as action, intentional or not. Before doing so, revisiting the debate that gave rise to mere capacity views is in order.

1.1. Classical Causal Views

Davidson was the most influential proponent of the classical causal theory, an influential account of (intentional) action to which the mere capacity views responded. Classical causal theory holds that any movement caused by an intention (or another appropriate mental state) is an (intentional) action (Davidson, 1963; Bishop, 1989; Enc, 2003). It thus focuses on the causal history of a movement in deciding its action status. But the famous objection from deviant causal chains demonstrates that there are cases in which an intention causes a movement that is nonetheless not (intentional) action, but rather a mere behavior (Peacocke, 1979; Mitchell, 1982; Davidson, 1980).

Climber: Tired of holding on, Carl intends to let go of the rope that keep his mountaineer friend from falling down a cliff. Realizing the horrific nature of this intention, he trembles in shock, accidentally letting go of the rope.

The intention to let go of the rope caused the letting go of the rope, but in a deviant way that intuitively does not amount to an action, much less an intentional action. Such cases illustrate that whether a movement is an action cannot be explained "ballistically" in terms of preceding mental states, such as intentions causing the movement. Many variations of such cases have been given, casting doubt on the possibility of explaining action by appealing only to the causal history of the movement.

Subsequently, many theorists have turned away from the causal history of a movement and towards the relation between the agent and the movement as it unfolds. Rather, they argue agents must *control* their movement for it to be an action. The rock climber is not in control of his letting go of the rope, and thereby his letting go is not considered an action. Rather, his trembling and subsequent letting go of the rope is something that merely happened to him. These control views, according to which action is agentially controlled movement, have made the challenge of explicating agential control a central issue in action theory (Frankfurt, 1978).

control a central issue in action theory (Frankfurt, 1978).

1.2. Capacity Views

Deviant causal chains show that the causal history of a movement does not settle its action status. In response, mere capacity views offer an alternative, noncausal explanation of agentially controlled action. On this view, what matters is not the actual intervention, but *whether an intervention could have taken place if necessary*. The ability to intervene in the movement, if it becomes required to achieve success, is sufficient to speak of control over a movement (Frankfurt, 1978; Zhu, 2004; Di Nucci, 2013; Levy, 2013). This eliminates the need for a causal story in the resultant picture of control, as Frankfurt explains in his discussion of purposive (i.e., controlled) behavior:

> "Behavior is purposive when its course is subject to adjustments which compensate for the effects of forces which would otherwise interfere with the course of the behaviour, and when the occurrence of these adjustments is not explainable by what explains the state of affairs that elicits them. The behaviour is in that case under the guidance of an independent causal mechanism, whose readiness to bring about compensatory adjustments tends to ensure that the behaviour is accomplished. […] The causal mechanisms which stand ready to affect the course of a bodily movement may never have occasion to do so; for no negative feedback of the sort that would trigger their compensatory activity may occur. The behavior is purposive not because results from causes of a certain kind, but because it would be affected by certain causes if the accomplishment of its course were to be jeopardized" (Frankfurt 1978, p.160)

Explaining whether a movement is an action requires more than just establishing that it was a controlled movement. Some mere behaviors are controlled by bodily subsystems, such as the dilation of your pupils. Action requires that it is the agent who controls the movement, as all control views agree. Mere capacity views take the existence of an agential mechanism that has the capacity to intervene to be sufficient to consider a movement in which it could intervene agentially controlled action. Some disagreement persists on the exact nature of the agential mechanism. Throughout my discussion of mere capacity views, I will hence remain agnostic on this. Instead, I rely on examples that invoke an agential intervention, thereby presupposing the activity of whatever agential mechanism is the correct one. Frankfurt presents us with such an example when he introduces a canonical case of agentially controlled action without causation in his discussion of the downhill driver (Frankfurt, 1978).

Downhill Driver: A driver is going downhill in her car. She is perfectly satisfied with the current speed and trajectory, which will ensure she arrives safely at her downhill destination, as intended. She decides not to intervene, and indeed does successfully arrive at her destination.

Is the downhill driver in control of her driving, and thereby acting? She is not causally intervening. In fact, one can suppose that hitting the gas or brakes, or steering the wheel, would have made her *less likely* to arrive at the destination. Whereas deviant causal chains such as the nervous climber show that the causal history of an event is not sufficient to establish that the event is an action, the downhill driver case goes further. It suggests that even as the movement unfolds, no causal

intervention may be needed to render the event an agentially controlled action. Frankfurt argues that what matters for both control and action is not a causal story, but a counterfactual one: the driver was ready to intervene if it became necessary to do so. Hence, he takes the driver to be acting: an intuitive result. The resulting picture of agentially controlled action is as follows:

Capacity to guide: A movement is agentially controlled action iff the agent has the capacity to intervene in the movement in a way that makes success more likely if the success of the movement is otherwise under threat of becoming jeopardized.

Zhu (2004) and Steward (2012) both develop such a view further. Zhu considers possible responses to the downhill driver case that try to reintroduce a causal story, particularly Mele's suggestion that a standing intention causally sustains the downhill drive as it unfolds (Mele 1997). Zhu rejects Mele's suggestion, noting that such standing intentions may ironically undermine the success of actions such as the downhill driver case. Instead, Zhu embraces a mere capacity view:

> "it is helpful to view actions as the guiding and controlling relation between an agent and his bodily movements. Moreover, we can understand the guiding and controlling functions in terms of certain underlying structures and the functions that these structures support, without invoking the problematic notions such as "agent causation" or "causal powers"" (Zhu, 2004, p.310)

Another interesting view is proposed by Helen Steward. Her metaphysically driven account of action focuses on "agential settling" (Steward, 2012). Steward is concerned with the possibility of freedom in a deterministic world. She argues forcefully that it is not just free action but rather any form of self-determined action that is incompatible with determinism. I cannot do justice to the many nuances of her view here. Rather, I will focus only on the notion of an agent settling a matter, which Steward takes to be the core concept of action explanation. By settling whether a given movement occurs or does not occur, an agent is involved in the complex causal web that unfolds in the world. Yet Steward does not take the settling to be a causal power the agent must exercise, thereby explaining how agency is not at odds with causal determinism. She further (and rightly) recognizes that many of our actions are causally advanced by subsystems that have little resemblance to the "top-down agent causation" we expect to encounter in action explanation. So, she, too, opts to explain agential settling in terms of a capacity of the agent, rather than an actual intervention.

> "It is not in virtue of the occurrence of any special sort of causal antecedent or component that the movement counts as the result of a moving by me. Instead, it is in virtue of my possession of an ongoing capacity to prevent altogether, stop in its tracks, reverse, alter, change the direction and speed of, or otherwise affect the motion in question. It is this power that means that my activity constitutes a settling by me of what in fact occurs with

> respect to my body as I dance, even when the movements are not the ones I should really have liked to be able to make, and even when the details of exactly how I move are not necessarily under direct, conscious supervision. The crucial point is that they could at any instant have come under that direct supervision" (Steward, 2012, p.52)

Notably, where Frankfurt focused his analysis of agentially controlled action on the agent's capacity to compensate for deviations in the movement if it threatens to become jeopardized, Steward expands her notion of agentially controlled action by highlighting the capacity to prevent and stop a

movement as it unfolds. This thought is also developed in Levy's intentional action-first account of voluntary action (2013). Levy takes full-fledged intentional action to be primitive. He then explains voluntary actions, such as absentmindedly walking or habitually brushing one's teeth, by reference to agential capacities. On his view, to be acting, the agent must have a capacity to intentionally act to stop or continue the movement. In doing so, he illustrates the appeal of a novel research program that takes intentional action to be the starting point, rather than the end goal, of explaining action. Both Levy and Steward thus develop a new form of the mere capacity view.

Capacity to Stop & Go: Necessarily, A is an action by agent S at time *t* iff S has the capacity and the opportunity at *t* to intentionally stop and continue making the bodily movements involved in A-ing. (Levy 2013, p. 715)

The proposals reviewed so far provide explanations of action. They do not aim to explain the narrower category of *intentional action*. Di Nucci proposes a mere capacity view that addresses the latter. He develops a Frankfurtian view in a series of papers criticizing classical causal theories (Di Nucci, 2008, 2011, 2013). On his view, when an epistemic constraint is added to the capacity to intervene in a movement, the resultant conditions pick out which actions are intentional actions. He provides an intuitive example of calling back a number that is displayed on one's phone. Contrast two such scenarios: in the first, I reach my friend under an unknown new number that was assigned to him yesterday. I can hardly be said to have intentionally called my friend, as I did not expect or know how to reach him. In the second, I recognize the number and call back. I intentionally called my friend. Even if my overt movements were identical, an important difference remains, as Di Nucci notes. Only the latter call was under my *rational control* in the sense of it being reasonable for me to expect that my dialling the number would result in my reaching of my friend (Di Nucci, 2011, p.196f.). Put formally, Di Nucci's proposal can be summarized as:

Capacity to guide + Epistemic constraint: A movement is an agentially controlled intentional action iff the agent has the capacity to intervene in the movement if the success of the movement otherwise threatens to become jeopardized, and the agent can reasonably be expected to know or find out about her movements under the relevant description.

What unites all these authors is that they turn away from a causal analysis of agential control in favour of a counterfactual one. They focus only on those possible situations in which the agent would have been required to intervene to ensure success. In providing a non-causal analysis of agency, mere capacity views overcome the problem of deviant causal chains that plagued "ballistic" classical causal theories.

Mere capacity views explain why the climber's letting go of the rope is not considered action, as the

climber had no capacity to intervene in the movement of letting the rope slip once he became nervous. Rather, the climber trembled when he realized his horrible intention, a trembling he could not make compensatory adjustments to, nor stop in its tracks. Hence, mere capacity views correctly classify his movement as an uncontrolled mere behaviour, rather than as an agentially controlled action.

Mere capacity views also explain why the downhill driver is acting despite not causally intervening. The downhill driver, we may stipulate, stands ready to intervene by hitting the breaks or taking the wheel. As such, she is in control over her movement, and thereby acting, regardless of whether she ends up actually intervening by grabbing the wheel. The fact that she would have intervened, had it become necessary, is sufficient to establish that she is performing an agentially controlled action. Thus, no causal story is invoked by the mere capacity theorist. By changing the analysis of control from a causal story to a counterfactual story (what the agent *could* do, if required), mere capacity

views deliver the correct verdict on both important problem cases, making genuine progress in action theory.

2. Argument and Illustration: Misattribution of Action as Shown by Passive and Active Breathing

Mere capacity views, however, encounter serious problems of their own, or so I will presently argue. Capacity views deem the capacity of an agent to intervene in a movement sufficient for agentially controlled action. But many movements in which agents can intervene are by default controlled non-agentially, and consequently not actions. Capacity views wrongly classify these movements as actions, rather than mere behavior, ignoring the de-facto control exerted by subpersonal mechanisms as the movement unfolds. To illustrate this, let us consider a class of examples in which mere capacity views misclassify movements.

As I am writing this, my swallowing, posture maintenance, and blinking are all flexibly adjusted in ways that allow for the proper functioning of my body. These passive processes are, intuitively, not my doing, nor are they under my control. Rather, bodily subsystems integrate information to generate flexible control over these passive movements. Yet as I reflect on this, I can, at any time, intervene in these passive processes. I can blink, swallow, or adjust my posture actively if I want to. But the mere capacity to do this at any time does not render these movements agentially controlled action *unless I causally intervene*. The possibility of active swallowing does not render passive swallowing an action. Let us look at one of these movements in more detail.

A case that vividly shows the failure of mere capacity views is that of breathing. There are two distinct ways in which humans breathe. Sometimes, humans *actively breathe*, controlling the depth and frequency of their breath, such as after exercise, during meditation, or before giving a presentation. Typically, though, breathing is a passive affair, with intricate bodily subsystems generating flexibly controlled breaths that ensure an adequate oxygen balance in one's blood. In physiology, this mode of *passive breathing* is called eupnoeic (John & Paton, 2003). Prima facie, active breathing is an agentially controlled action, whereas passive breathing is a mere behavior guided by bodily subsystems. Yet mere capacity views must classify *both* active and passive breathing as agentially controlled action. Let us look at the case of breathing in more detail.

Passive breathing is generated by a homeostatic mechanism (Prabakhar & Semenza, 2015), and is thus certainly a controlled behavior. Homeostatic mechanisms are usually divided into three components: a receptor or sensor, a control system, and an effector (Berridge 2004, chapter 2). Information is collected by the sensor. In the case of breathing, this is done primarily by chemoreceptors in the medulla oblongata, as well as the aortic and carotid bodies. This information is compared to a goal state. In the case of breathing, the goal state is maintaining an adequate supply of oxygen in the blood. When sensors detect a significant deviation from the goal state, typically by tracking CO_2 saturation, the relevant information is transmitted to the control system. In passive breathing, this system is the respiratory system in the pons and medulla oblongata (Richter & Smith 2014, Hudson et al. 2011). Based on the transmitted information, the respiratory system affects changes to the depth and rhythm of one's breath to ensure its continued success at supplying oxygen. Doing so requires the effectors, which are the muscles executing the movement. The respiratory system initiates movement in the diaphragm, as well as muscles in the abdomen and rib cage. In summary, passive breathing is a goal-directed, flexibly controlled movement caused by a bodily subsystem.

Active breathing is different. One can take a breath to achieve many aims, such as to communicate impatience, to make a rhetorically dramatic pause in a speech, to start meditating, or to "catch one's breath" after exercise. When agents catch their breath, they are trying to get more oxygen, as exercise has lowered the oxygen level in our blood. The agent may not know the physiological details causing the sensation. Yet the information that is collected by the same sensor that enables passive breathing becomes known to her via the sensation of being out of breath. She can then, on a personal level, initiate active breathing by deploying motor commands via the motor cortex, adjusting the movements of our diaphragm, as well as muscles in the abdomen and rib cage. Both active and passive breathing generally use the same effectors (i.e. muscles), which may be activated to different degrees, such as when voluntarily hyperventilating (Hudson et al. 2020). When active breaths occur in response to high CO_2/low O_2 levels in one's blood, they use the same informational signals, collected by the same sensory mechanisms. The effect is that when the goal of maintaining an adequate supply of oxygen in one's blood becomes jeopardized, such as after exercise, one becomes aware of this state of depletion and initiates changes in the movement that typically ensure success. A crucial difference between active and passive breathing lies in the motor signals that cause them. Passive breathing emanates from the respiratory centre in the medulla oblangata and pons (Richter & Smith 2014), while active breathing also recruits different motor command neurons in motor cortex and the cerebellum, neural structures associated with voluntary movement and motor control (Evans et al. 1999, McKay et al 2003, Hudson et al 2011).

Of course, exercise is not the only example of this active breathing intervening in eupnoeic breathing. Every time oxygen saturation deviates sufficiently from its goal state, an agential response is prompted. Moreover, as we have seen, capacity views understand agential control to be the capacity of an agent to intervene in a movement, in a way that makes success more likely if the success of the movement otherwise threatens to become jeopardized. In the case of passive breathing, the goal is to achieve an adequate oxygen balance. The bodily mechanism to achieve this goal state uses information gathered by receptors to bring about muscle movements, leading to eupnoeic breathing patterns. But when the goal of achieving the oxygen balance is sufficiently threatened, this information is broadcasted to the agent, who stands ready to intervene by actively causing deeper breaths: active breathing. Active breathing thus allows the agent to intervene in one's passive breathing if passive breathing's success threatens to be jeopardized. But this is highly problematic for capacity views, as it forces them to classify passive breaths as agentially controlled action. This is

because they take the *capacity* for intervention, rather than the intervention itself, to explain a movement's status as agentially controlled action. On these views, the agent is agentially acting whenever they breathe, simply because they could intervene if required – much like the downhill driver.

Newer variations of the mere capacity view do not solve this problem, either. Consider the epistemic constraint condition, which requires an agent to be reasonably expected to know or find out about her movements as they unfold. Agents can reasonably be expected to know both when they breathe and when their breathing has failed to achieve its aim, simply because such failures are automatically raised to the agent's attention. Since early childhood, the agent will know that the adequate response to this situation is to start actively breathing. The stop-and-go variation of capacity views, too, fails to distinguish passive and active breathing: agents can stop and continue their breathing at will, having both ample capacity and opportunity to do so either by inhibiting the passive breath or by substituting it with active breaths. In brief, for all capacity accounts considered, *agents have the require capacity for intervention over their passive breathing*, rendering passive breathing an agentially controlled action.

Classifying passive breathing as agentially controlled action is misguided: it is, after all, typically cited as a standard example of mere behavior (di Nucci, 2008 p.17; di Nucci, 2011, p.180; Mele, 1997, p.142). This is, of course, because control over passive breathing is generated by bodily subsystems, rather than the agent, as our discussion of breathing's homeostatic mechanism illustrated. Hence, we

do not attribute passive breathing to the agent in the same way we attribute actions to the agent. We certainly would not classify it as an instance of the narrower classes of *voluntary* or even *intentional* actions. Rather, we classify passive breathing as a mere behavior, produced by the agent's body, the same way we classify one's digestion as a mere behavior.

The classical causal theory does not face this problem, as it focuses on the causal history of a movement. In passive breathing, changes are caused by the respiratory system in the pons and medulla oblongata, both part of the brainstem; hardly where one would locate complex cognition by the whole individual, but rather a bodily subsystem. Active breathing emanates from the motor cortex, an area associated with planning, control and execution of voluntary action. The classical causal theorist can exploit this difference in the causal history of active breathing to correctly classify active breathing as an action and passive breathing as mere behavior controlled by a bodily subsystem. The mere capacity view cannot help itself to this solution, as it aims to provide a non-causal explanation of action.

In summary, I have argued that mere capacity views wrongly categorize movements where control is generated by a bodily subsystem, but agents retain a capacity for intervention, such as in swallowing, blinking and body posture. I illustrated this with a case study of breathing. Causal control for passive breathing is generated by a bodily subsystem. Agents have a capacity of control over passive breathing, yet the mere capacity does not render passive breathing an agentially controlled action. What matters is not what agents counterfactually *would* do to control a movement if it threatens to become jeopardized. What matters is the causal differences between passive and active breathing.

3. Objections

How could proponents of a mere capacity account respond to the challenge from breathing, blinking, and swallowing? They may be tempted to bite the bullet and classify these movements as actions, but this is unattractive. Passive breathing, blinking, and swallowing are standard examples of mere

behavior and are classified as such by both proponents of capacity and causal views (e.g., di Nucci, 2008 p.17, di Nucci, 2011, p.180; Levy, 2013 p. 174; Mele 1997, p.142).

Homeostatic processes are permanent and required for proper life function. One does not cease to breathe, blink or swallow unless one dies. If mere capacity views were right, agents would be permanently controlling and acting as these processes unfold. When sleeping, one sometimes wakes up due to breathing problems and exercises control via active breathing — does the capacity to do this imply that one acts even when one is sound asleep? This would be a highly counterintuitive result.

Sometimes agents lose the capacity to regulate their breathing. When having nearly drowned, the capacity to override our passive breathing pattern is temporarily lost. Is the agent's permanent action of breathing thereby paused only after nearly drowning? Does the eupnoeic breathing pattern gain or lose action status based on momentary analyses of whether an agent currently can intervene?

Biting this bullet does not only bring counterintuitive results; it also ignores the intricate bodily control structures involved in homeostatic processes. Passive homeostatic processes like passive breathing do not require an agent. They can be generated by the brainstem and spinal cord alone (John & Paton, 2003). Of course, healthy agents, but not the brainstem and spinal cord, can control their breathing actively. However, the point is that the brainstem's and spinal cord's outputs do in fact constitute agentially controlled action on a mere capacity view once a (possibly inactive) agential mechanism is added. Regardless of what agential mechanism a specific capacity view will rest on, it would render such behaviour active even if the agent never used that agential mechanism during their life. The control generated by the passive homeostatic process should not be ignored: as long as

the agential mechanism does not become active, the resultant movements are not attributable to the whole individual, but rather to dedicated, well-understood bodily subsystems. An analysis of control that does not recognize the role of these subsystems fails to deliver an accurate analysis not just of action, but also of control.

Di Nucci and Levy both recognize that agent-responsive passive homeostatic movements may pose a challenge to their view. Di Nucci (2008) considers the life-sustaining function of agent-responsive homeostatic movements as a reason for why agents cannot have guidance over them.

> "Even though we can avoid blinking for a while, 'we can't avoid blinking for good, and we can't always avoid blinking; the same way in which we cannot avoid breathing for good (assuming that killing ourselves does not count as a way of controlling our breathing patterns)' (Di Nucci, 2008, p. 26)."

Levy, who considers the case of blinking as well, considers a similar response when writing:

> "Indeed, the kind of intentional control one has over one's blinking is generally very limited and unstable: even when the opportunity is present, the capacity to refrain from blinking is only temporary. Swallowing represents a similarly intermediate case of less than full-blooded intentional control." (Levy 2013, p.715)

Many actions offer less than full-blooded intentional control. For example, I cannot indefinitely raise my arm: my muscles will give in after a certain time. Other actions, such as eating, cannot be permanently avoided without risking death. But up until death, a strongly willed agent can avoid them, such as when an individual starves themselves to death. In contrast, one's breathing or blinking can only be inhibited for a certain time. When one's oxygen balance is sufficiently endangered, bodily subsystems override agential control to maintain the homeostatic balance. I take Di Nucci and Levy to point to this difference to argue that the control over these movements is importantly limited. In the case of breathing and blinking, a powerful reflex overrides agential control after a certain threshold is reached.

Not all passive movements in which agents can intervene are like this. For example, the protective swallowing reflex can indefinitely be inhibited unless foreign objects are in one's mouth (Ertekin et al., 2001, Nishino et al., 2011). You can see this clearly by simply concentrating on not swallowing for a few minutes, which is uncomfortable and requires focus, but is not that difficult. Indeed, while rare, clinical cases of agents inhibiting their swallowing for long stretches of time exist. Patients that develop Phagophobia (fear of swallowing), often after a traumatic event, typically avoid swallowing for long stretches of time (Franko et al., 1997). So do agents with psychiatric cases of dysphagia (swallowing difficulties, Zald & Pardo, 1999). Hence, it seems plausible to me that an agent may just as well permanently avoid swallowing as they may avoid eating. An even clearer example is found in posture maintenance (Ivanenko & Gurfinkel, 2018). Is there any reason to believe that one's passive posture maintenance could not be permanently overridden by a determined agent manually controlling his posture? Certainly, the agent will be miserable, but so will the person starving themselves to death. This illustrates that there are some agent-responsive passive homeostatic movements we can avoid for good. Hence, the appeal to limited control over such movements fails to avoid the objection presented by agent-responsive passive homeostatic movements.

The appeal to limited control does not resolve the counterexamples here considered. Can we amend capacity views to find alternative ways of providing the distinction between passive homeostatic

processes and active agential interventions?

One may seek to individuate the movements of passive and active breathing in such a way that active breathing can no longer be considered an intervention into the passive breathing pattern. But it is unclear what then would ground individuating the movements. The goal of the movement, e.g. the maintaining of an adequate oxygen balance, can and often will coincide between passive and active breathing. Neither can it be the information provided, which stems from the same mechanism. Nor can direct reference to the activity of the agential mechanism be used. Mere capacity views hold the possibility of the intervention by an agential mechanism to be sufficient to consider a movement agentially controlled action, regardless of what other mechanism produced the movement in the first place, which is what enables them to explain cases like the downhill driver.

The crucial difference in the movements we considered is the cause of the movement. While passive breathing is caused by the respiratory system, active breathing involves cortical regions such as the superior motor cortex and premotor cortex to initiate the compensatory adjustments of the breathing mechanism (Evans et al., 1999). Yet using this difference to individuate active and passive breathing as incommensurate processes amounts to analysing breathing in terms of its causal history. This runs counter to the purported aim of mere capacity views to provide an explanation of agential control that does not rest on the causal history of the movement. Action is explained by mere capacity views "regardless of what features of his prior causal history account for the fact that it is occurring" (Frankfurt, 1978, p.159).

Another difference could be found in the attention the agent devotes to their respective capacity to intervene in the movement. After all, I may be aware that I can intervene in my breathing when I actively breathe, but I do not attend to this capacity if I passively breathe. Two problems arise. First, it raises a dilemma. Either such awareness must be accessed in a present mental state of the agent (explicit awareness), or, while it need not actually be accessed, it must be accessible to the agent if required (tacit awareness). If explicit awareness is chosen, the capacity view of agential control seems much too restrictive. While driving, I shift my gears. In doing so, I am typically not aware of how I could exert control if shifting gears went wrong, yet I can do so if necessary. However, shifting gears is an action I control. Consequently, the awareness that is necessary for agential control cannot be explicit awareness. If tacit awareness is chosen, the objection no longer rules out cases of passive breathing. I am tacitly aware that I can breathe at any given moment, and this information is accessed when required. Regardless of whether explicit or tacit awareness is chosen, the mere capacity view incurs misclassifications.

I have argued that homeostatic movements such as passive breathing, blinking, swallowing and posture maintenance pose an insurmountable problem to capacity views. Distinguishing these movements from their active counterparts requires a causal analysis. Yet classical causal theories also suffer from problems: deviant causal chains and cases such as the downhill driver suggest they are misguided. Rather than focusing on the causal history of a movement, perhaps we should focus on the causal picture that emerges as the movement unfolds. The last section describes how such causal control views can solve the problems this article considers by summarizing recent advances in our understanding of how agents exert control over their movements as they unfold. The picture that emerges is one in which agentially controlled action is a thoroughly causal phenomenon.

4. An Alternative: Causal Control Views

This section sketches an alternative view: causal control views. Like mere capacity views, they hold

that agents only act when they agentially control the movement as it unfolds. Unlike mere capacity views, they believe that agents continuously make small, causal adjustments as movements unfold. I review the work of three authors that have used advances in our empirical understanding of agential interventions in action to argue for the existence of such small, ubiquitous causal adjustments by agents in their actions. I then resolve the problem cases considered in this paper to show that in each case of agentially controlled actions, agents cause changes as the movement unfolds, whereas in cases of mere behavior, agents do not cause changes as the movement unfolds.

In a series of recent papers, Wu has suggested that attending is not a passive matter: for Wu, attention is necessary for agentially controlled action (Wu, 2011, 2016, 2019, 2023). Whenever agents perform an action, Wu argues, many options present themselves: I could pick up the glass of water with my left or right hand, using all kinds of different grip variations, for example. Yet acting requires selecting on one such course of action, which is enabled by attention. Hence, attention is necessary for action selection. Wu expands upon this picture when discussing intentional action. In intentional action, the standing intention serves as a top-down influence on one's attention, enabling an action-relevant coupling between what the stimulus and the agent's response. Such attentional control is necessary to sustain the action, as even during the unfolding of the movement, many

different ways of continuing the movement present themselves, requiring decision. Wu believes attentional control to be a necessary, and central, part of agentially controlled action.

In a series of papers, Buehler (2019, 2021, 2022) further explores the role of attention in action. Research in attention typically distinguishes between exogenous attention and endogenous attention. In exogenous attention, attention is captured by an external stimulus, such as a loud yell or a bright light. In endogenous attention, and individual's goals and intentions steer the deployment of attention, such as when focusing on a detail in a painting. Buehler argues that an agent's ability to deploy attention implements the exercise of flexible, occurrent control (2019). On this view, an agent attending to an aspect of a task exercises control over a task. Buehler is careful to note that attention is only one form of control and does not take attentional control to be necessary or sufficient for agentially controlled action. Nonetheless, his research similarly illustrates the active, agential character of specific forms of attention. In two later papers (2021, 2022), Buehler fleshes out this picture in two important ways. He elaborates on the mechanism that constitutes the ability for attentional control, which is the executive system. He also explains why this mechanism can be attributed to the agent, despite the executive system being a psychological subsystem. As opposed to e.g. one's respiratory system, the executive system flexibly integrates and coordinates the workings of other subsystems in a characteristically agential way: in fact, it is unclear how agency could arise without such integration by the executive system. Buehler thus identifies an agential mechanism and uses it to explain when and why control provided by a mechanism should be attributed to the agent.

Sripada's work on self-control presents us with an empirically informed account of occurrent control that fleshes out the way in which this executive mechanism allows agents to exert control: via decisions. Sripada argues that when agents exert self-control, they do so by a continuous deployment of atomic, intrapsychic actions called cognitive control actions (Sripada 2019). For example, both ignoring a salient stimulus and refraining from grabbing a tasty treat require the agent to exert self-control by using cognitive control. More recent work by Sripada (forthcoming) further develops this decisionist framework, arguing that agents exert control over their actions by continuously choosing between different options that present themselves before and during the execution of an action. These "value-based decisions" unfold rapidly, over a few hundred milliseconds. Nevertheless, they are personal-level events and are often accompanied by a phenomenology of choice. Sripada takes such decisions to form the basis of agentially controlled action: agents control their actions in virtue of continuously deciding on the course of the movement as it unfolds.

In summary, despite important differences, all three authors shed light on small, causal activity by

the agent as a movement unfolds. Regardless of whether one takes the central form of such activity to be top-down attention, executive control, or decision, this ongoing "micro-activity" monitors and integrates information to enable agential control. But, as all these authors agree, such control is not counterfactual. Rather, agents continuously micromanage movements by actually deploying attention and making decisions that guide the unfolding of the movement.

I believe that these insights on the nature of attention and decision resolve the problem cases we encountered. Recall the downhill driver, who is introduced as a passive bystander to her action of rolling downhill. Is this driver really passive? As described, we imagine that she attends to the speed and direction of her car, as well as her means of intervention as the car goes downhill. After all, it is only by accurately predicting the trajectory and speed of her downhill drive that she knows there is no need to hit the brakes or grab the wheel. No overt movement occurs. Yet she is active in the movement as it unfolds, as her continued attention and decisions constitute small exercises of causal

control. In fact, this is necessary for us to consider the case an agentially controlled action. If the downhill driver were just to let her thoughts wander freely to revisit a heated argument at work, or maybe even close her eyes, thereby no longer causally directing her attention to the task, she would no longer be in control of her movement, and would not be acting. This resolves the mystery of the downhill driver: it is not her counterfactual capacity to intervene, but her small-scale actual causal interventions as the movement unfolds, that render the movement agentially controlled action. Causal control views take such small-scale interventions to be the basic building blocks from which more complex actions, such as driving, are generated.

Deviant causal chains, too, pose no problem for causal control theories. When the climber trembles and loses grip of the rope, he has no way of deciding upon the course of the movement. Rather, the emotional realization that he intended to kill his friend triggers an automatic, involuntary response which crowds out the climber's agential control over the process. The agent's loss of causally efficacious choice characterizes the deviancy in such deviant causal chains. The agent cannot decide to keep holding the rope. As this is prevented by mandatory processing overriding decisional responses, the causal control theorist can rightfully claim the ensuing movement is not an action. No decision to let go of the rope becomes causally efficacious, nor does the agent have any possibility of suppressing or guiding the trembling. This is because the causal control loop that is required for agential control and action status is broken as the movement unfolds – the agent has lost control. Causal control views differ from their predecessor, the classical causal view, in that they focus on causal control throughout the movement, rather than the cause that precedes the movement. Hence, they are not threatened by the deviant causal chains, which exploit the ballistic nature of classical causal views.

Lastly, turn back to the case of active and passive breathing. As we saw earlier, passive and active breathing can rely on the same information, collected by the same receptor, to activate the same muscles in regulating oxygen saturation of the blood. Yet, causal differences remain. Active breathing is caused by, and causally sustained by, ongoing person-level decisions implemented in motor cortex, whereas passive breathing relies on the respiratory subsystem. It is thus clear that a causal control theory can distinguish both movements by their causal profile. After all, cortical activation is detectable throughout the process of breathing in, not just at the onset of it. This matches the phenomenological profile of active breathing. As I am writing this, I start a voluntary breath: I can flexibly control its unfolding as it occurs, attending to the movement of my muscles and the air I take in. I can prolong the movement as it unfolds to take a deep breath up until my lungs are filled, or abort after a quick whiff. Active breathing is thoroughly under my occurrent causal control, as both its neural signature and its phenomenological profile attest. Passive breathing is not, as it involves no executive brain networks. In fact, merely deploying attention to my breathing forces me into agentially controlled, active breathing. It is impossible to passively breathe if one exerts ongoing agential control over one's breath. Both the causal and phenomenological profile of breathing ground a principled distinction between active and passive breathing that rests on ongoing

micromanagement through active decisions that occurrently control active, but not passive, breathing.

Causal control views thus explain problematic cases of agentially controlled action. More detail must be added to causal control views, on which decisions of the agent, such as regarding the deployment of attention and cognitive control, act as the basic building blocks of agential control. Such work is currently underway (Wu, 2023; Buehler, 2023; Sripada, forthcoming; Wong, forthcoming). But already it seems very plausible that recent developments tying together the ubiquitous, micro-level decisions agents make as movements unfold, the nature of attention in action, and the more

philosophically-laden notion of agential occurrent control are promising antidotes to problems that plagued the kind of "ballistic" classic causal views which marked the philosophy of action of the last century. In understanding agential control as small-scale interventions that occur as the movement unfolds, contemporary causal control theorists highlight the many small ways that agents contribute to movements after they are initiated.

Conclusion

Mere capacity views promise an account of agential control, and thereby action, that does not rely on a causal analysis. I have argued that these views fail. Agents have the capacity to intervene in many mere behaviors. Capacity views misclassify these mere behaviors by attributing their control to agents and granting them action status. Causal control views, on the other hand, correctly classify these movements by taking only those cases in which agents actually intervene to be agentially controlled action. Agentially controlled action is a causal phenomenon, not a mere capacity.

Acknowledgments

I am thankful for having received a tremendous amount of support in developing this project. I thank Hong-Yu Wong, Krisztina Orban, Chandra Sripada, Daniel Gregory, Denis Buehler, Gregor Hochstetter, Wesley Sauret, Alfred Mele, Stephen Butterfill, Elisabeth Pacherie, Sarah Moss, Ishani Maitra, Sarah Buss, Michelle Dyke, Marcela Herdova, Julian Saccone, Florian Teichmann, Selina Guter, Klaus Corcilius, Michael Wacker, Chiara Brozzo, Roberta Locatelli, Ann Whittle, Gabrielle Kerbel, Margot Witte, Mica Rapstine, Sean Costello, Aaron Glasser, Paul de Font-Reaulx, Elizabeth Beckman, Sarah Sculco, Lorenzo Manuali, Joxelyn Wang, as well as audiences at SSPP, ESPP, GSWG, the UMICH ME-WIP, the PONS forum and the APA central. Finally, I thank three anonymous reviewers. Their generous and helpful comments helped improve the paper significantly.

Anscombe, G. E. M. (2000). *Intention*. Harvard University Press.

Bermúdez, J. P., & Felletti, F. (2021). Introduction: Habitual Action, Automaticity, and Control. *Topoi*, *40*(3), 587-595.

Bermúdez, J. P. (2017). Do we reflect while performing skillful actions? Automaticity, control, and the perils of distraction. *Philosophical Psychology*, *30*(7), 896-924.

Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience. *Physiology & behavior*, *81*(2), 179-209.

Bishop, J. C., & Bishop, J. (1989). *Natural agency: An essay on the causal theory of action*. Cambridge University Press.

Buehler, D. (2018). The central executive system. *Synthese*, *195*(5), 1969-1991.

Buehler, D. (2019). Flexible occurrent control. *Philosophical Studies*, *176*(8), 2119-2137.

Buehler, D. (2022). Agential capacities: a capacity to guide. *Philosophical Studies*, *179*(1), 21-47

Buehler, D. (2023). Explicating agency: the case of visual attention. *The Philosophical Quarterly*, *73*(2), 379-413.

Donald Davidson (1963). Actions, reasons, and causes. *The journal of philosophy*, 60(23):685–700.

Davidson, D. (1980). Freedom to act. In *Essays on Actions and Events*. Routledge.

Davidson, D. (2016). Agency. In *Agent, action, and reason* (pp. 1-37). University of Toronto Press.

Derchi, C. C., Mikulan, E., Mazza, A., Casarotto, S., Comanducci, A., Fecchio, M., ... & Sinigaglia, C. (2023). Distinguishing intentional from nonintentional actions through eeg and kinematic markers. *Scientific Reports*, *13*(1), 8496.

Di Nucci, E. (2008). Mind Out of Action: The Intentionality of Automatic Actions.

Di Nucci, E. (2011). Automatic actions: Challenging causalism. *Rationality Markets and Morals*, *Forthcoming*.

Di Nucci, E. (2013). Action, deviance, and guidance. *Abstracta*.

Enç, B. (2003). How we act: Causes, reasons, and intentions.

Ertekin, C. (2001). Voluntary versus spontaneous swallowing in man. *Dysphagia*, *26*(2), 183-192.

Evans, K. C., Shea, S. A., & Saykin, A. J. (1999). Functional MRI localisation of central nervous system regions associated with volitional inspiration in humans. *The Journal of Physiology*, *520*(2), 383-392.

Frankfurt, H. G. (1978). The problem of action. *American philosophical quarterly*, *15*(2), 157-162.

Franko, D. L., Shapiro, J., & Gagne, A. (1997). Phagophobia: A form of psychogenic dysphagia a new entity. *Annals of Otology, Rhinology & Laryngology*, *106*(4), 286-290.

Goldman, A. I. (2015). *Theory of human action*. Princeton University Press.

Happee, R., De Vlugt, E., & Schouten, A. C. (2008). Posture maintenance of the human upper extremity; identification of intrinsic and reflex based contributions. *SAE International Journal of Passenger Cars-Mechanical Systems*, *1*(2008-01-1888), 1125-1135.

Hudson, Anna L., Simon C. Gandevia, and Jane E. Butler. "Common rostrocaudal gradient of output from human intercostal motoneurones during voluntary and automatic breathing." *Respiratory physiology & neurobiology* 175, no. 1 (2011): 20-28.

Hudson, Anna L., Lee D. Walsh, Simon C. Gandevia, and Jane E. Butler. "Respiratory muscle activity in voluntary breathing tracking tasks: implications for the assessment of respiratory motor control." *Respiratory physiology & neurobiology* 274 (2020): 103353.

Hyman, J. (2015). *Action, knowledge, and will*. Oup Oxford.

Irving, Z. C., & Glasser, A. (2020). Mind-wandering: A philosophical guide. *Philosophy Compass*, *15*(1), e12644.

Ivanenko, Y., & Gurfinkel, V. S. (2018). Human postural control. *Frontiers in neuroscience*, *12*, 171.

Levy, Y. (2013). Intentional action first. *Australasian Journal of Philosophy*, *91*(4), 705-718.

Mckay, Leanne C., Karl C. Evans, Richard SJ Frackowiak, and Douglas R. Corfield. "Neural correlates of voluntary breathing in humans." *Journal of applied physiology* 95, no. 3 (2003): 1170-1178.

Mele, A. R., & William, H. (1992). *Springs of action: Understanding intentional behavior*. Oxford University Press on Demand.

Mele, A. R. (1997). Passive action. In *Contemporary Action Theory Volume 1: Individual Action* (pp. 135-143).

Springer, Dordrecht.

Mitchell, D. (1982). Deviant causal chains. *American Philosophical Quarterly*, *19*(4), 351-353.

Nishino, T. (2013). The swallowing reflex and its significance as an airway defensive reflex. *Frontiers in physiology*, *3*, 489.

Pacherie, E., & Mylopoulos, M. (2021). Beyond automaticity: The psychological complexity of skill. *Topoi*, *40*(3), 649-662.

Peacocke, C. (1979). Deviant causal chains. *Midwest Studies in Philosophy*, *4*, 123-155.

Prabhakar, N. R., & Semenza, G. L. (2015). Oxygen sensing and homeostasis. *Physiology*, *30*(5), 340-348.

John, W. M. S., & Paton, J. F. (2003). Defining eupnea. *Respiratory physiology & neurobiology*, *139*(1), 97-103.

Richter, Diethelm W., and Smith, Jeffrey C. "Respiratory rhythm generation in vivo." *Physiology* 29, no. 1 (2014): 58-71.

Shepherd, J. (2014). The contours of control. *Philosophical Studies*, *170*(3), 395-411.

Shepherd, J. (2021). *The shape of agency: Control, action, skill, knowledge*. Oxford University Press.

Sripada, C. (2021). The atoms of self-control. *Noûs*, *55*(4), 800-824.

Sripada (Manuscript under review). Agency Computationalized: A Valuationist Model of Human Agent Architecture

Steward, H. (2009). Sub-intentional actions and the over-mentalization of agency. *New essays on the explanation of action*, 295-312.

Steward, H. (2012). *A metaphysics for freedom*. Oxford University Press.

Steward, H. (2020). Agency as a two-way power: A defence. *The Monist*, *103*(3), 342-355.

Steward, H. (2022). Frankfurt cases, alternative possibilities and agency as a two-way power. *Inquiry*, *65*(9), 1167-1184.

Whittle, A. (2021). Control and abilities to do otherwise. *Inquiry*, 1-21.

Wong, H.Y. (Manuscript). On a motor theory of bodily action.

Wu, W. (2011). Confronting Many-Many problems: Attention and agentive control. *Noûs*, *45*(1), 50-76.

Wu, W. (2016). Experts and deviants: The story of agentive control. *Philosophy and Phenomenological Research*, *93*(1), 101-126.

Wu, W. (2019). Action always involves attention. *Analysis*, *79*(4), 693-703.

Wu, W. (2023). *Movements of the Mind: A Theory of Attention, Intention and Action*. Oxford University Press.

Zald, D. H., & Pardo, J. V. (1999). The functional neuroanatomy of voluntary swallowing. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, *46*(3), 281-286.

Zhu, J. (2004). Passive action and causalism. *Philosophical Studies*, *119*(3), 295-314.