3
Non-deviant Causation

3.1 Introduction

The last game of my high school basketball career, we almost completed an unexpected, improbable comeback against the defending state champions. After clawing back for two quarters, we had finally taken a one-point lead. We had the ball. There was little time left. A turnover occurred, and their forward dribbled madly towards an open lay-up. I, along with a teammate (I can't remember who), chased him down. He, surprised and nervous, jumped in the air with his back to the basket. I think that's what happened. Memory sometimes operates deviantly. But I recall seeing him in between us, ball down near his thigh, and breathing a sigh of relief. He had no chance of making the shot. I doubted he would even get a shot up. And yet. Perhaps in spite of his frantic state he managed an intention to make a shot—to win the game. An audacious intention, but winning a state championship, maybe, imbues you with audacity. What I know is that with a little flipping motion, he sent the ball up. It crawled onto the rim with a strange swirling spin that pulled it over and in.

Deviant causation ruins careers.

It also undermines intentional action. In cases of basic deviance, the problem occurs in between mental state and behavior. A philosopher wants to distract a commentator, and intends to knock over a glass. But this intention upsets him such that his hand shakes uncontrollably and accidentally knocks the glass over (Mele 1992: 182). Since the intention does not cause behavior in the right way, the philosopher does not intentionally knock the glass over.

Consequential deviance occurs after the behavior has begun, and before it is complete. Jones intends to kill K by shooting him. Jones misses by a mile, but the noise from his rifle scares a nearby herd of pigs that trample K (Davidson 1980: 78). The killing is unintentional because, again, it was deviantly caused.[1]

[1] A third type of deviance, tertiary deviance, concerns the history of the relevant mental state (see Mele 1987: 56). My account of non-deviant causation is not intended to address tertiary deviance.

The Shape of Agency: Control, Action, Skill, Knowledge. Joshua Shepherd, Oxford University Press (2021).
 DOI: 10.1093/oso/9780198866411.003.0003

I like to imagine the end of my high school basketball career as a case of double deviance. My opponent's intention was so audacious and unnerving that it led to the ugly-looking flip shot. And the spin was so wild, so out of control, that it corrected for the awful angle at which the ball hit the rim. The made shot was no intentional action.

Many have proposed solutions to causal deviance (e.g., Bishop 1989; Enç 2003; Schlosser 2007). None enjoy consensus. Is this a serious problem for causalists about intentional action? Some seem to think not. According to these theorists, deviant causation is a minor worry. According to others, the problem does not need to be solved (Tännsjö 2009).[2]

But for many, the problem is deep enough to motivate rejection of causalism. Ezio Di Nucci, for example, makes a characteristic inference in the following passage:

> Deviant causal chains are symptomatic of a fundamental problem with causalism . . . we should give up the fight to accommodate deviant cases and focus, rather, on developing an alternative to causal views of action . . .
>
> (Di Nucci 2013: 45)

Or, to take another example, Timothy Williamson writes of causalism's treatment of deviant causation in the past tense, as though the failure were complete:

> [The causal theory] succumbed to the problem of deviant causal chains: the causation from intention to success may be of the wrong kind for action . . . Attempts are still made to solve the problems for attempted analyses of action by adding epicycles, but none has worked.
>
> (Williamson 2017: 171–2)

In this chapter I offer a solution that works.

[2] John Hyman (2014) argues that deviant causation cases present a problem, not to causalist accounts of action, but to the Humean theory of causation. Once we reject the Humean theory in favor of a powers and dispositions-based ontology, the problem of deviant causation dissolves. I am not committing here to a particular view of causation, but Hyman's view, also expressed in the cited paper, that action explanation can be both causal and teleological, is congenial in some ways to remarks I make below.

3.2 Non-deviant Causation Explained

What we want is to explain the difference between the causal pathway(s) operative in two kinds of case.[3] The account of non-deviant causation I offer here draws on the account of control's possession developed in chapter two.[4]

Recall the shape of that account. The possession of some amount of control with respect to a plan-state is the possession of a capacity (or package of dispositions) to repeatedly and flexibly bring behavior to match the content of that plan-state, to some degree, across a well-selected set of circumstances. What we need in order to move from the possession to the exercise of this capacity is a sense of the causal routes taken in uncontroversially successful instances of the capacity's exercise.

Regarding control's possession, we were somewhat relaxed about the selection of sets of circumstances. Non-deviant causation happens in a particular circumstance. So we need a way to say, of some particular happening, that it was non-deviant. We need to move from a particular happening to the modal space surrounding that happening, against which the particular happening can be revealed as falling into a certain pattern.

We need to deploy care. Consider the following kind of issue.

Tormund is anxious at t. Fix his current level of anxiety as part of a well-selected set of similar circumstances, and he does poorly, via particular causal routes that involve scattered attention. Tormund is also well-rested at t. Fix his current level of energetic resources as part of a well-selected set of

[3] When I first developed this account, I used the notion of causal pathways without much understanding of how some scientists deploy it. Then I read Winning's (2019) excellent discussion of control, and Ross's (2018) very interesting work on differences in scientific usage of mechanisms and pathways as constructs for causal explanation. For those interested, the account of non-deviance I offer could be taken to apply to systems conceived as mechanisms, or as (sometimes ad hoc) packages of mechanisms, or to pathways. The notions of deviance and non-deviance seem to me to apply more abstractly, both to the operations of well-specified mechanisms, and to the modelling of various pathways in whatever causal space one wishes to model.

[4] My proposal is similar in spirit to Aguilar's (2012). Aguilar appeals to the notion of reliability in developing a sufficient condition for non-deviant causation: according to Aguilar, the mechanism that produces an intentional action must be reliable (2012: 10). In my view this is on the right track. One might thus see what follows here as a way to flesh out the nature of the reliability Aguilar discusses. I do not know if Aguilar would agree with this. The view I offer is, after all, very closely tied to the account of control developed above. Further, Aguilar applies his discussion of reliability only to overt basic actions—the account developed here is developed independently of an account of intentional action, but is intended to apply to all intentional action-types (mental and overt). Finally, in my view Aguilar's discussion does not place sufficient weight on the importance of circumstance to a proper understanding of reliability. Even so, I suspect that Aguilar and I are in substantial agreement.

similar circumstances, and he does well, via particular causal routes that involve more focused attention. Two well-selected sets give conflicting verdicts about how much control Tormund might be thought to possess in the actual circumstance he is in at t.

Say, now, that Tormund behaves at t in a certain way—this way is unusual given his anxiety, but fairly normal given his level of rest. Is this a deviant success, as the set that fixes his anxiety level would suggest? Or is this a non-deviant success, as the set fixing his energetic resources would suggest?

Neither. The problem is that if we wish to understand control in the actual happening, we need to give preference to the actual happening in how we understand the modal space surrounding it. Not just any old well-selected set of circumstances will do. We have to treat the actual happening as baptismal regarding the modal space surrounding it. We need to start from Tormund's actual circumstance and specify all of the relevant causal parameters—the anxiety and the energetic resources, and much else besides—into the relevant set of circumstances. The actual circumstances, then, play a constraining role on the system one is modelling by selecting a set. They do so in at least two ways.

First, they force greater specificity and detail regarding the dispositional condition of the agent in the circumstances. Agents vary over time. Here we need to get a snapshot of the agent in order to proceed.

Second, the actual circumstances have to be typed in some way. We need to understand how the circumstances may vary. This will likely come in the form of probability distributions regarding certain ways things might unfold. This will present epistemic difficulties in many cases. These problems are familiar from attempts to build causal models of complex systems. These models require judgments about causal variables and causal patterns in the flux of events. When we make such judgments, we identify the type of circumstances that fit into our model—we specify parameters for what makes some causal variable or causal pathway (which could be construed as a pathway amongst variables) count as sufficiently similar to another. Of course these models will often be rough approximations. But we can envision how this is supposed to go. And we can allow that things often change over the course of behavior, so that variations in circumstances come into play.

Call a set of circumstances that begins in this way, specifies the relevant parameters, and types the circumstances into a (probabilistic) model of possible permutations, a comprehensive set of circumstances. This need not be an infinite number of circumstances. It depends upon correctly gauging

the circumstance-type the agent is in in the baptismal circumstance. Sometimes only a few parameters are relevant. Sometimes there will be more.[5]

Here, then, is the account:

Non-deviant causation. Given an agent J, a plan-state P, a comprehensive set of circumstances C, and a token circumstance T which is a member of C, J's behavior in service of P is non-deviantly caused in T, relative to C, if and only if (a) J's behavior in T reaches a level of content-approximation L (to the content of P, or to some part of P) that is within a sufficiently normal range for J, relative to J's behavior across C, (b) J's behavior in T is produced by way of causal pathways that, when taken, lead, with sufficient frequency, to J's reaching L across C, and (c) P is among the causal influences of J's behavior.

We want to exclude deviantly caused successes. Non-deviant causation does so by way of two requirements. The relevant level of content-approximation—the level our agent actually achieves—must be within a sufficiently normal range for a comprehensive set of circumstances. So this will be something the agent can do repeatedly, in circumstances like this. And the causal pathways operative in the actual case must be such that taking those pathways leads to something close to the behavior in question, in circumstances like this. There is no room for a deviant causal pathway to operate surreptitiously.

There is nothing magic at work here. It takes some work to get it straight. But in a way, this is simply a plausible precisification of what is implied by the term non-deviance.

3.3 An Alternative: Wu on Control and Non-deviance

Before considering objections to this account, I wish to hold it up against an alternative: Wayne Wu's account. Doing so will help clarify the substance of my own.

Wu agrees with the general view that control is essential for understanding action and agency:

[5] Compare this to Manley and Wasserman's comments about "C-cases" regarding dispositions (2008: 75).

> Agentive control yields phenomena of central philosophical interest: moral, rational, reason-based, skilled, conscious, epistemic, and free agency. To understand these specific forms of agency, however, we must understand the core phenomenon, agentive control. (Wu 2016: 101)

How does Wu attempt to account for agentive control? For Wu the operation and interaction of both intention and attention hold the keys: "If you like, talk of 'cause in the right way' or 'appropriate causal role' is spelled out by delineating more completely the psychological basis of agency, namely the intention-attention nexus" (114).

Wu claims that the intention-attention nexus is the psychological basis of agency, and that understanding this nexus will help us understand non-deviant causation. How so?

According to Wu, in normal cases of action agents face a "many-many problem" (see Wu 2011). Agents face many possible targets for action, and agents face many ways to implement action related to these targets. In order to solve the many-many problem, agents must find a one-one mapping from target to action. (They need not do so through deliberation—often agents face a non-deliberative many-many problem.)

For Wu, solving the Many-Many Problem is a necessary condition on action. Entities that never solve the many-many problem would be in a position such that only one target for action, and only one movement-type, is available to them. According to Wu:

> What is by hypothesis not available to them is the possibility that that very stimulus can be mapped to different response at that time such that the creature could react to the stimulus in this alternative way rather than the pre-set way. Their behaviour does not count as action, for at each time they are driven by what is essentially a reflex, and I take it that this never exemplifies agency. (Wu 2011: 54)

The important question now, for Wu, is how human agents solve the Many-Many Problem, and thus act intentionally. Here is where the "intention-attention nexus" is critical. Wu (2016) reviews evidence that both intention and attention are crucial for establishing one-one mappings in paradigm cases:

> The idea is that in setting intention, one sets the weights that biases which selections are made in action (psychologists speak of task sets). So, if one intends to act on X, then X is selected for action; if on Y, then Y is selected.

> This is part of intention's causal role, one that is driven by the content of intention. The point is that by registering attention in the causal theory, we can expand our conception of intention's causal role. (Wu 2016: 110)

How does the biasing role of intention, and the selective function of attention, deliver an account of non-deviant causation? Wu leverages these biasing and selective functions to offer a sufficient condition for non-deviant causation:

> S exerts agentive control in F-ing at t if S's F-ing is a result of an intention to F where this intention solves the Many-Many Problem faced by S at t by directing appropriate attention-mediated coupling. (Wu 2016: 114)

In his (2016), Wu suggests we take this condition as necessary as well. This is consistent with his (2011), where he argues for the following analysis of (bodily) intentional action.

Necessarily, for all agents, S, bodily action types B, motivational states, M, towards B-ing and action contexts C:

> S's motivational state M (paradigm: an intention to B) structurally causes the selection required in implementing a solution to the Non-deliberative Many-Many Problem appropriate to action B for S in C if and only if S intentionally B's in C in the sense that she B's because she is in motivational state M. (Wu 2011: 68)

So, for Wu, agents never exercise agentive control unless a relevant intention structures attention-mediated coupling. As Wu would have it:

> It turns out then that deviance is quickly addressed in a way that we expected. For causal deviance revealed that the causal theory did not adequately explicate agentive control. Such control reflects a complex psychological capacity of the agent, so that defects in control must be defects in the implementation of that capacity. By highlighting the intention-attention nexus, we illuminate agentive control. (Wu 2016: 115)

Let us take a step back before considering this proposal. Proponents of causalism about intentional action often forward causalism as a way of naturalizing action theory—of showing how actions fit within the causal order described by science. Now, if the goal is to naturalize action theory, one

might think that the best way to fill out a causalist theory of action is with reference to cognitive psychology and neuroscience. And one might think that worries about deviant causation are ultimately worries born of ignorance regarding the science. Once we specify the relevant causes, our work is done. Alvin Goldman once endorsed an approach like this (1970: 62). According to Goldman, deviant causation is not philosophically interesting, or philosophically challenging. Discovering the nature of non-deviant causation is simply a matter of elucidating in some detail the causal processes "characteristic of intentional action" (62). Wayne Wu's work on control and non-deviance comes from a perspective similar to Goldman's. Accordingly, he directs attention to empirical details. Wu thinks that if we see how (human) agents actually exercise control, we will understand what goes wrong in cases of deviant causation.

In my view, there is something right in this approach. I have argued that in order to explain non-deviant causation in a particular case, we have to go some way towards modelling the causal system the case concerns. We have to specify causally relevant parameters, and we must have some understanding of how these parameters vary across similar circumstances. Only then does the causal pattern under investigation stand out as non-deviant. One way to think of what Wu is doing is this. Wu is suggesting that in human agents, the causal parameters will necessarily involve certain links between attention and intention.

That is partially an empirical claim, although the claims involving necessity seem to go beyond what science could be expected to show. And ultimately, regarding the full range of action-types and agent-types, I think Wu's conditions cannot be expected to hold. This is just a comment about my prior expectations. There are a wide range of agent-types and action-types. It is doubtful that a particular link between attention and intention, as implemented in human actions—and arguably only in some human actions—would set the standard for non-deviant causation in every case.

Further, in requiring that the relevant attention-mediated coupling be appropriate, both of Wu's conditions appear to assume some account of non-deviant causation. Notice that this troubles Wu's (2011) analysis of intentional action as well: that analysis appears to depend on the assumption that any implementation of a solution to the Many-Many Problem will be non-deviant. But that fails to address the possibility that deviant causation could afflict this part of the process as well. Does Wu's account of agentive control have the resources to specify "appropriate" coupling in a satisfactory way?

It looks like the answer is no. Consider the following two possibilities. On the first possibility, attention-mediated coupling happens in two ways: deviant and non-deviant. If so, then Wu's account fails to account for the non-deviant cases, and so fails to account for non-deviant causation. On the second possibility, attention-mediated coupling is by default appropriate (or non-deviant). So long as a path through the behavioral space is charted, such that the agent's behavior approximates the content of her intention, then the coupling has been appropriate. If that is right, then the word "appropriate" here does no additional work. But that cannot be right. Intention and attention do not work flawlessly—deviant causation cases are precisely cases in which the causal role of intention for behavior is deeply flawed. We have been given no reason to think that the same kinds of flaws will fail to afflict attention.[6] Wu's account fails to explain non-deviant causation.

The problem here is not Wu's interesting work in connecting intention, attention, and typical instances of human action control.[7] The problem is rather that Wu's account, while specifying mechanisms he takes to be important for the implementation of control, cannot guarantee the non-deviant working of these mechanisms. Wu is looking in the wrong direction.

In order to say what it would be for any particular mechanism or causal process to run non-deviantly, we have to move up a level of abstraction. This requires rejecting the thought that there is anything magic or special

[6] I would lodge a similar complaint against Gwen Bradford's (2015) recent, very interesting, discussion of what she calls competent causation. This may not be the same thing as non-deviant causation, but it is close. Bradford deploys the notion in an account of achievement—achievements necessarily involve competent causation. (According to Bradford, an achievement is composed of a process and a product. The process must be difficult, and the product must be competently produced by the process.) What is competent causation? Bradford offers the following:

> Given that A causes some outcome E, A competently causes E via his activity, phi, when A, while phi-ing, has the requisite amount of [Justified True Beliefs] about his phi-ing causing E. The requisite amount is a percentage from the total possible beliefs about the activity, and [Justified True Beliefs] that are about the overall structure of the activity are worth more than discrete beliefs. (Bradford 2015: 79)

This is enough to see that Bradford's account of competent causation is not an account of non-deviant causation. They can deviantly cause behavior as well as intentions. Just tell Davidson's climber case in a way that what causes the trembling hand is not the climber's intention to let go of the rope, but his belief that letting go of the rope is how he will effect the murder. Bradford's competent causation says nothing about this kind of case, but presumably deviantly caused effects are not competently caused effects. Thus Bradford's competent causation requires an ancillary account of non-deviant causation.

[7] This work is of course important. In this connection, see Buehler (2019) for criticisms of Wu and of earlier work of mine on control. As Buehler notes, however, my concern with control here is different from Wu's (and Buehler's). I do not here seek to explain how human action control works. I am concerned with the nature of control and non-deviance.

about the mechanisms that do implement non-deviant causation. And this is just as well. Action is a many-splendored phenomenon, and the set of all action-types is incredibly diverse. So the possible causal pathways from plan-state to successful behavior are likely to be diverse as well.

3.4 Questions, Objections, and Replies

In discussing this view with others, a number of questions and objections arise. I will consider several.

3.4.1 Cases of Easy Success

Davidson's (1980) climber intends to kill another person by letting go of a rope. The intention's acquisition unnerves him, leading him to deviantly let go of the rope. Suppose this person is constructed such that literally anything he does will lead to letting go of the rope. Say his grip is exceedingly loose. So if he breathes in too heavily, or has a slight motor tremor, or looks up towards the top of the cliff, or down towards the bottom, the rope is gone, the other climber plummeting. Doesn't my account implausibly imply that, in such circumstances, whatever the killer does is non-deviantly caused?

Here is an initial reply. The case is a bit under-described. Either the killer conforms to his plan or he doesn't. If the plan was to let go of the rope in a certain way, then his behavior may fail to conform. His success will occur in spite of the low degree of control exercised.

But we can fill out a better version of an easy success case. Say an angel is balanced very precariously on the head of a pin. Say also that this angel has, in most circumstances, little control over their behavior. (Maybe the devil has afflicted this angel with a condition that makes their intentions generate absolutely random movements.) In this case, however, the angel is well situated. Given how precarious their position is on the pin, any movement at all will cause a fall. This angel now intends to fall. And fall they will, although the proximate cause of the falling will be some movement generated by the intention. Say, in addition, that their plan is very simple. The angel intends to fall by causing itself to fall (perhaps by way of some basic action). They do not have a detailed conception of how they will cause themselves to fall. But they are rightly confident that they can do it. (The angel is perhaps analogous to a human who intends to do something simple, via some

opaque causal route that runs through the motor cortex.) And every time they form the intention to fall, they successfully execute their simple plan. They cause themselves to fall. Doesn't my account implausibly imply that, in such circumstances, whatever the angel does is an instance of non-deviant causation?

Yes, that is what my account implies. Although it is not implausible. It is rather a unique situation. Sometimes success is easy to come by. Sometimes plans are easy to execute. What the account requires is that, in addition to the agent's behavior conforming to the agent's plan, the agent's causal pathway is reliably productive of conformity. That is true in every token circumstance for our angel. They bring about success, perhaps, in a million different ways. That's okay. If there are a million ways to get it, choose one.

On to the next one.

3.4.2 Cases of Rare Success

Say that an agent engages in planned behavior that succeeds, across a very big set of circumstances, 4 percent of the time. Can my account explain the difference between hits and misses? In the 4 percent of cases that involve success, the agent is by definition behaving in an abnormal way. How can these be cases of non-deviant causation?

My reply begins by observing that the case is under-described. Is there a relevant difference in the causal pathways taken between hits and misses? If so, we can split the sets to get more information. Remember that the set of circumstances at issue is comprehensive. It will not miss a difference between hits and misses, if the difference is relevant.

So suppose that there is no difference between the causal pathways. The success is thus abnormal, but the causal pathway that leads to it is not. Sometimes the best an agent can do is put herself in a position to succeed, even when success is rare. In such a case the successes and the failures would be examples of non-deviant causation, even though the success qua success would be due to a certain amount of luck. This is a feature of the world, not a problem for my account.

3.4.3 Cases of Strange Success

Here is the basic idea. We imagine an agent's behavioral profile across a well-selected set of circumstances. The profile is consistent, but the causal

routes to success are consistently strange. Say, for example, that every time Kawhi takes a shot, it bounces multiple times on the rim. The bounces are unexpected. But again and again—at least much of the time—the strange bounces lead to a made shot. There are probably better examples. The worry is that my account implies that these unexpected bounces are non-deviant, and this is purportedly implausible.

My reply is, first, that the violation of expectation could be a fault on our end, not on the end of the behavioral profile. But let's assume that the causal routes taken really are strange. What do we mean by this? Perhaps we mean: in nearby sets of circumstances, these causal routes are not reliably productive of success. But why think that is a problem for success in these circumstances? Sometimes a situation calls for causal pathways that are normally not great. When I was a kid they used to say that the way to a man's heart is through his stomach. Maybe the way to a man's heart is not generally through his brain. But in some circumstances, maybe that's just the way. Strange success might look strange, but when the set is well selected, this too is non-deviance.

3.4.4 Cases of Partial Deviance

In such cases the agent displays good control with most of the plan. Perhaps an obstacle is encountered, deviant causation occurs, and the agent moves on without a hitch. In such a case success depends upon a deviant joint in the causal architecture.

This is not an objection against the account. I mention it only to note that it is not an infrequent feature of our world. Intuitions will differ regarding how much deviant causation is consistent with intentional action. I wager that some is—that is, we should not rule out all deviant causation as necessary for some behavior's qualifying as intentional action. What an account of non-deviant causation allows us to do is to say what parts of a plan's execution are controlled, and what parts are due to deviance.

3.5 A Further Criticism

Erasmus Mayr (2011) offers counterexamples to what he calls sensitivity accounts of non-deviant causation. These are accounts that lean on the insight that the causal processes at issue in intentional action should be in some sense sensitive to the content of a relevant intention. My account has

affinity with this basic idea. So it should be informative to see how my account fares against Mayr's counterexamples. After discussing Mayr's counterexamples, I consider one of my own.

Mayr notes that sensitivity accounts come in two varieties. First, we might understand sensitivity as a kind of counterfactual dependence between content and behavior. Second, we might understand sensitivity as involving an explanatory connection between content and behavior (as in, e.g., Schlosser 2007). Mayr focuses on the latter kind of sensitivity account, which makes sense given Christopher Peacocke's (1979) subtle version of it. In the present context, however, this is unfortunate. For my account is closer in spirit to a counterfactual dependence version of a sensitivity view. I maintain that non-deviant causation is not a matter of the explanatory connection between content and behavior, but rather of a level of reliability displayed by the agent in executing some plan.

Even so, in discussion Mayr touches on a view close to my own. Mayr's initial target is Peacocke's view, which requires a one-to-one function between explanandum (such as behavior) and explananans (such as intention). Given this target, his use of the following case (due to Scott Sehon) makes good sense:

> A baseball pitcher tries to move his arm at a speed of 70 mph, and succeeds. Though the fast movement of his arm is an intentional action, we need not assume that the movement would have had a correspondingly greater or smaller velocity had the agent's intention been to move his arm at 69 or 71 mph. Instead, we may well believe that even if the pitcher had intended to move his arm at 69 or 71 mph, his actual speed would have remained the same. So we need not assume that there is a one-to-one function correlating the behavior with the content of the intention in order to classify the pitcher's behavior as an (intentional) action.
>
> (Mayr 2011: 119–20)

Mayr argues that cases of course-grained abilities such as this one undermine any one-to-one correlation between intentions and behavior. For intentions may be much more precise than behavior allows. This strikes me as a problem for a view that depends on one-to-one correlations. But it is no problem for the account I offer. In some cases noise makes relevant differences to success and failure; regarding the exercise of control, and regarding the measurements that underly non-deviant causation, a range representation is thus often more appropriate than a specific number. As a result, course-grained abilities are easily accommodated.

Interestingly, in commenting on this case, Mayr considers and rejects a crude version of the kind of account I offer:

> A defender of the sensitivity strategy might react to such cases by claiming that no precise one-to-one fit between the intention's content and the ensuing behaviour is required for agency, but only a 'sufficiently tight' correlation—a correlation of certain ranges of the intention's content to ranges of behaviour. However, such a move seems to give up the central idea of strongly differential explanation as the key to agency, because this kind of explanation requires one-to-one-correlations. Besides, the move will not help the causalist to escape the problem of antecedential deviance, as limiting cases of 'coarse-grained' abilities to act—'on–off' abilities—make clear.
>
> (Mayr 2011: 120)

Mayr rejects an account that requires "only a sufficiently tight correlation" between intention and behavior for two reasons. First, moving to this kind of account leaves behind strongly differential explanation—that is, explanation of B by A that requires a one-to-one function mapping A to B. Since my account of non-deviant causation does not appeal to strongly differential explanations, this is no problem for me. Second, this kind of account cannot account for limiting cases of course-grained abilities to act.

Why not? Mayr offers this case as an example. We have two agents, A and B. A can wiggle his ears, but "wiggling is an 'on-off' affair with him…he can only 'simply wiggle', but cannot wiggle in specific ways" (120). Agent B cannot wiggle, but in this case intends to anyway and deviantly does so. Mayr comments:

> With regard to the sensitivity of the ensuing bodily motion to the content of the intention, both cases are identical. Both agents' ears would not have wiggled had they failed to have the intention, and no relevant further dependences between bodily motions and intentions hold in either case. So, an adherent of the sensitivity approach must either rule both as cases where the differential explanation criterion is satisfied, and therefore both as cases of actions, or neither. But both solutions are wrong. (2011: 120)

Perhaps this case works against the proponent of strongly differential explanation. But once we have moved away from this requirement it is less convincing. The proponent of my account will say in response that Mayr has failed to consider a large enough set of circumstances. Presumably Agent B's deviantly successful attempt to wiggle his ears is not replicable to

the same degree as Agent A's (otherwise why would Mayr say Agent B "cannot wiggle"?). So, once we consider a wider range of cases, it should become clear that Agent A has some degree of control over ear wiggling, while Agent B has a very low degree of control.

3.6 A Still Further Criticism

Scott Sehon has been a leading critic of causal accounts for some time. Sehon sometimes leans on the problem of deviant causation. In one place, he criticizes an account that is in some ways similar to mine—Jesus Aguilar's. Aguilar stresses the importance of the reliability of whatever mechanisms mediate between mental state and behavior (2012: 10). Sehon is not impressed. His chief problem with this proposal is that, in his view, Aguilar is unable "to show that 'reliable' is being used as a causal term and not as a way of sneaking in . . . teleological concepts" (2016: 104).

How would this usage of reliability sneak in teleological concepts? Sehon's discussion here is odd. He quotes Aguilar's elucidation of reliability, on which it is "a measurable capacity of a process or system to satisfy a given goal," and on which levels of reliability can be assessed in terms of probability of satisfying a goal (Aguilar 2012: 7). So far, so good. He then offers an explication of Aguilar on which it is the agent—not the disjunct "process or system"—that must be reliable. He then complains that the reliability in question "does not specifically concern the causal chain between the reason state and A's B-ing; the reliability condition applies to what the agent would have done in different circumstances and even with different mental states" (2016: 107). And he uses this claim about a difference in mental states to treat Aguilar's notion of reliability as requiring an appeal to sensitivity to reasons, rather than as concerning causation; Aguilar's reliability "amounts to saying that the behavior was goal-directed if it was caused by an appropriate mental state, and, in a variety of other circumstances not necessarily even involving that mental state, the agent's behavior would have been sensitive to her reasons" (107).

This interpretation of Aguilar strikes me as wrong. But I am concerned here to see how this criticism might apply to my own account. I can resist the interpretation of reliability as involving rationalizable behavior—behavior that involves sensitivity to reasons across circumstances that do not include the motivating mental state. For the notions of repeatability and normalcy at

work in my account hold fixed the relevant plan-state. So whether Sehon's criticism of Aguilar is misguided or not, it does not apply to my account.

3.7 Conclusion

According to opponents of causalism about action, the account is bankrupt. Causalists need some specification of control, but without a solution to problems of causal deviance they write empty checks. But the opponents are wrong. Quoth Megan Thee Stallion: money good. Non-deviant causation occurs all the time, and it requires no magic to explain what happens when it does. It is the normal operation of causal pathways at a time and place, where normal is explicated by reference to the background of a well-delineated causal model that captures the time and place.