

Joshua Shepherd

Replies for Brains Symposium on 'Halfhearted action and control'

Thanks to Aaron Henry for organizing, to John Schwenkler and the editors of Brains for running such a quality site, and to Andreas, Nora, and Zachary for their trenchant comments.

Reply to Andreas Elpidorou:

In the following passage, Elpidorou makes two claims:

'[A] characterization of weak overall motivation in terms of 'degree of control an agent possesses over the action' doesn't settle the issue as to why the motivation is weak in the first place nor does it explain why relaxing control (and thus adversely affecting the performance of the action) is an appropriate response.'

In response to the first claim, I agree: my account is meant to offer a characterization of the nature of halfhearted action, while leaving open explanatory issues regarding how such action arises. Given the complexity of human motivation, my guess is that there are several routes. Indeed, Andreas usefully highlights one that I did not focus on in my paper – boredom. (Elpidorou, by the way, has a lot of interesting work on boredom.)

In response to the second claim, I have a minor quibble. I do not think of the lower degree of control associated with halfhearted action as always something the agent explicitly does, as the language of 'relaxing control' suggests. Nor do I think the lowered degree of control is always appropriate. In some cases, acting halfheartedly will be irrational. In some cases, acting halfheartedly will be morally inappropriate.

Elpidorou goes on to suggest that we understand halfhearted action in terms drawn from (behavioral) economics. So, we might think of halfhearted action as action that [a] carries a small net utility, where [b] the costs of suboptimal performance are low. This is an interesting suggestion, and one worth developing.

Upon reflection, I would resist the thought that all halfhearted action can be captured in this way. One reason is that actions that fail to carry net utility are common, and it seems these may be halfhearted. (Arezo's dance performance, discussed on pp. 271-272, could be construed in this way. Perhaps Arezo should stop practice early if she cannot practice well. But there are plenty of examples.)

We might rather think of Elpidorou's suggestion as a way to capture rational halfhearted action. Again, though, I am not sure the rationality of action can be fully captured in these terms. Take a case like Arezo's:

Arezoo is performing a difficult dance routine at the end of a long day of practice. She knows that a part of this routine requires a significant amount of attentional control—that is, effortful concentration directed to certain parts of the action. She also knows she is suffering from weak motivation. Having performed under similar circumstances, Arezoo knows she is capable of concentrating harder than she currently is, and as the relevant dance part approaches, she knows that an effort at concentration would enable her to successfully execute the part, while maintaining her current level of concentration will impair performance. Due to her weak motivation, however, Arezoo does not make an effort at concentration. Consequently, although Arezoo intentionally performs the part, her execution of the part is very poor.

Say Arezoo's net utility here is small but positive, and costs of suboptimal performance are low. I think that Arezoo's performance is irrational, since in knowledgeably acting halfheartedly Arezoo acts against the commitment embedded in her intention.

Even so, I think Elpidorou's suggestion captures an interesting range of cases. In a constructive spirit, I might suggest the following tweak. Consider a case in which costs associated with suboptimal performance are high, while risks of suboptimal performance are very low. Perhaps Arezoo is being watched by a talent scout, and her performance just then is critical to her career. But all the talent scout needs to know is whether she can perform a very simple maneuver. In such a case, halfhearted action could occur. And it could be rational, if the risk of suboptimal performance is very low. So we might add risk of suboptimal performance as an additional consideration. (We might also add available resources. I mention a different Arezoo case in footnote 9 on p. 272, involving reasons that justify the poor performance that stems from halfhearted action.)

What I think this indicates is that when we look closely at the structure of halfhearted action, we begin to see significant complexity. This complexity in turn suggests significant complexity regarding the practical rationality or irrationality of halfhearted action. Hopefully that generates some philosophical motivation towards further reflection on halfhearted action.

Reply to Nora Heinzelmann:

Heinzelmann suggests that the key functional upshot of weak motivation is not possession of less control, but allocation of less control to the action at hand. I like this suggestion. For all I know, it could be right. As Heinzelmann says, the evidence is consistent with it. And there is some initial plausibility in positing a link between weak motivation and allocation of less control.

Maybe it's worth saying that I don't think I absolutely have to disagree with Heinzelmann's suggestion. I am not trying to give an analysis of halfhearted action. So I think what Heinzelmann is suggesting can be made consistent with my claim, which is that an action is

halfhearted in virtue of motivational weakness, and that one functional upshot of motivational weakness is less control. When I say lowered possession of control is one key functional upshot of weak motivation, I don't mean it's the only one: I agree with her 'that motivation is but one factor affecting performance directly or indirectly.'

Still, I think it is worth examining Heinzelmann's claim about allocation more closely.

What does it mean to speak of the allocation of control? Sensibly, I think, one might interpret this as the allocation of control-relevant resources: perhaps attention, or effort, or inhibitory control, or planning resources, depending on the case.

If that is how we should understand allocation, how should we understand the allocation process? Is this an explicit process – does the agent decide to allocate less control, or decide not to intervene as less control is allocated? Or is this something that happens outside of the agent's awareness? Is the allocation automatic? Sub-personal? This is something worth exploring empirically. My guess is that in some cases, allocation of control resources occurs explicitly, while in others, allocation takes place implicitly, perhaps even automatically (if that term has any meaning). On some accounts of cognitive control, after all, resource allocation co-varies with levels of motivation, and is guided by computations (regarding, e.g., the opportunity costs associated with the task at hand) that seem largely sub-personal.

This issue of understanding an allocation process may have downstream implications regarding the account of control's possession one favors – whether mine or another. I first characterize the possession of control abstractly, in terms of abilities that enable one to (flexibly, repeatedly) bring behavior to match the content of motivational states across ranges of circumstances. One can select a restricted range of circumstances for various theoretical purposes, in order to get a closer look at the abilities that constitute the control. The science of action does this in various ways – adding pressure to some situations, or levels of skill, or varying the learning process, or examining the role of anticipatory imagery. How, then, might my account attempt to accommodate Heinzelmann's suggestion regarding control's allocation?

One option is to zoom out, and to include the allocation process within the measure of control's possession. On this option, an agent might possess less control when weakly motivated, in part because that agent will be disposed to allocate less control resources to the task at hand. I think this is consistent with my account – what is added is additional empirical detail regarding a process of allocation that would be important for understanding how weak motivation undermines control.

A second option is to zoom in, and to exclude the allocation of control resources from the measure of control's possession. The nice thing about this option is that it allows you to say what does seem plausible – that in at least some cases it looks like weak motivation does not impair control's possession, even if it impairs control's exercise. In these cases, then, Heinzelmann's interesting point about weakness of will seems to hold.

Taking this latter option raises interesting questions:

How tightly does the allocation of control resources correlate with levels of motivation? What are the influences on allocation? Might there be a role, e.g., for an agent's level of commitment, or for antecedent processes of planning?

Might weak motivation impair control independently of the resources allocated by the agent? If so, the points I make in the paper about attention might hold, and Heinzelmann will have uncovered an additional interesting wrinkle in the structure of action control.

As mentioned above, is allocation explicit or not? If it is, this zoom-in option is more plausible than the zoom-out option. But in cases where allocation is not explicit, I find the zoom-out option more tempting.

Is the allocation process discrete, taking place before the agent begins to act? Or is it dynamic, interacting with the range of processes – of inhibition, attention, effort, planning, etc. – that unfold together as the agent exercises control?

If this latter possibility is the more plausible one (and I think it is, although I'm not too confident about that), we might come to think of an allocation process as one way to interpret the functioning of the mechanisms underpinning action control. And if we come to think that, we are not too far away from the first option, on which control's exercise and possession within some range of circumstances include the allocation processes influenced by motivation.

Reply to Zachary Irving

Irving's comments triangulate the views in my paper with the neuroscience of mind-wandering and the case of creative action. The result is fruitful: I think creative actions, in addition to being an intrinsically interesting and under-explored area of action theory, pose good questions for my perspective on control and on halfhearted action.

Irving's discussion of creative action includes Katie, who is brainstorming ideas for a new novel. A key part of her activity in so doing is, in a way, to give up control over the direction her thinking takes. As Irving puts it, Katie 'can skillfully initiate and maintain a stream of unguided thoughts. She's *good* at losing control, letting her mind wander, and letting the ideas come to her.'

Suppose that Katie has no motivational deficit when she brainstorms. What should I say about her control in such a case? Irving writes,

'Shepherd should predict that Katie is wholehearted in virtue of a strong motivation to generate new ideas. This much seems correct. But Shepherd also predicts that Katie's strong motivation will increase her *control* over writing by *focusing* her attention. Phenomenologically, this prediction seems suspect. When Katie wholeheartedly throws

herself into a piece, her strong motivation may well lead her to *lose* control, to let her attention wander between disconnected ideas.'

I have two thoughts about this. First, my account would make the general prediction that normal or significantly strong motivation would not impair attention. But that's not the same thing as predicting enhanced focus of attention. The account is silent on the relationship between strong motivation and attention. One's degree of attentional focus may increase in lock step, and because of, increase in motivational strength. Or the relationship could be more complicated, involving a non-linear relationship, or involving the impact of motivation on attention only up to a point, or whatever. My account of halfhearted action is not intended to answer such questions regarding action that is not halfhearted.

Second, as Irving goes on to discuss, Katie's action has an interesting structure. She loses control over one aspect of it, but not over all of it. Indeed, I would argue that she loses control over one aspect of it precisely because this better enables the exercise of control over the larger behavior. And this may well require attentional focus to certain aspects of her action, even as Katie shields the other aspect from too much attention or effort or meddling. That is, it seems plausible to me that Katie will be attentionally attuned to the deliverances of her brainstorming. A big part of what makes Katie good at brainstorming, I would think, are her capacities to recognize and categorize the interesting thoughts that emerge, and to ignore the boring ones. As the process unfolds, then, she may need to focus on some thoughts and not others, to inhibit certain urges, or to pin down an interesting direction her thought is taking for further analysis, before letting her thoughts slide again. It is not as though brainstorming takes place in the complete absence of attention or control. It's just that the attention needs to be directed in a certain way.

Irving recognizes this, and goes on to say interesting and plausible things about the structure of creative action. I like his distinction between higher-order and first-order control. (Interestingly, this distinction maps a similar space as Ellen Fridland's distinction between strategic control and motor control (in 'They've lost control,' 2014 in *Synthese*), and as my distinction between executive and implementational aspects of control (in 'Conscious control over action,' 2015 in *Mind and Language*). But the distinction is clearly not the same one. So I look forward to reading Irving's paper on this.) The thing I wish to stress here is that these plausible things seem compatible with my account. I think I can maintain that failures of attention would impair the control Katie exercises in creative action.

In understanding why I want to say this, it's important to be clear regarding what it means to say that Katie has exercised some degree of control. Irving notes that my account characterizes control in terms of repeatable flexibility. Then he argues that this account gets things backwards when it comes to creative action. When Katie writes with normal (or strong) motivation, Irving predicts that 'her ideas would become more variable, less flexibly repeatable, and less controlled as a result.' And when Katie writes with impaired motivation, Irving claims that her ideas would be more repeatable in the sense that 'she would be more likely to rehash stereotypes and ideas she's written about before.' As a result, Irving claims that 'flexible repeatability is a marker of halfhearted creative action.'

This may well be, but this is not the way my account of control deploys the notion of flexible repeatability. As I use it, flexible repeatability is indexed to goal satisfaction. What the agent does with flexible repeatability is bring her behavior to match the content of some relevant motivational state, such as an intention. In my view, then, Irving's cases of creative action do not problematize my account of control or of halfhearted action. Rather, they raise the following interesting questions:

How should we think about the intentions behind creative action? Katie wants to avoid producing hackneyed ideas. She intends to brainstorm in a way that facilitates quality writing. How best to capture the content of such an intention, and what kind of background knowledge, know-how, etc. is relevant to our characterization of Katie's capacity to form such an intention? How should we think about the guiding role such a sophisticated intention might play?

In what sense are creative actions under an agent's control? Does the notion of control capture everything at issue in creative action?

I would think not: there often seems to be an element of practical rationality at issue in how agents conduct these kinds of actions. And I want to say that there is an aesthetic dimension to many creative actions that my account of control is not designed to capture. I think there are implications for our understanding of the nature of skill – but that is a subject for future work.