

Problems for Selection Problems: Comments on Wayne Wu's *Movements of the Mind*

Antonia Peacocke

2,947 words

In *Movements of the Mind* (2023), Wayne Wu offers a sophisticated and original theory of action in general, and mental action in particular. It has been carefully crafted with deep knowledge of relevant neuroscience and psychology. The book presents a model of what it is to act that sits at Marr's (1982) computational level of analysis.

There's much to admire here; let me first mention three connected advantages of Wu's proposal.

First, Wu is exceptionally clear on the relationship between automaticity and control of action. He's already known for a view (2013) he develops more deeply in this work: the view that there is automaticity involved in all intentional action, and this automaticity is needed to ground the kind of control essential to such action. You can't act intentionally without controlling what you are doing to meet the intention in question, but this control itself requires certain processes to be *outside* of your active, person-level guidance.

The second advantage is related. Wu's theory of control in intentional action presents an alternative to an inadequate theory of control originally suggested by Harry Frankfurt (1978/2015).¹ On Frankfurt's view, controlling what you're doing is a fundamentally dispositional matter: you control what's happening by being poised to intervene in a process that is otherwise proceeding without your concurrent causal support. In his famous example, a driver "coasting downhill in virtue of gravitational forces alone" may be controlling his action just fine because he is "prepared to intervene if necessary, and ... in a position to do so" (29). Wu has given us a way to understand how ongoing control can work in intentional action. He's done that by describing how we sustain attention on the task at hand, in part by employment of our working memory (Chapters 3-4).

The third advantage is a powerful theory of working memory as a realization of intention in action, a kind of "practical memory for work" that sustains what you do as you do it (93). Wu's theory of working memory carefully distinguishes between the idea of *storage capacity* in working memory and the *executive* component of working memory; this latter coordinates and modulates a whole host of further mechanisms that subserve action (§3.3). The result is a new substantive theory of one of the most important realizers of intentional action in general.

Although these are real achievements, there are also serious challenges facing one of the book's foundational claims. This is the claim that intentional action is necessarily a solution to a Selection Problem, i.e. a problem of selecting one among a set of exclusive alternatives of what to do.

This claim provides the framework for the main arguments of the book. Wu needs this claim to make any headway in his inquiry into the "*biology* of agency" (1). It is this analysis of intentional action that gives us the possibility of analyzing such action in computational terms, insofar as the Selection Problem is amenable to treatment as a kind of computational problem to be given a computational solution. But I have deep reservations with the idea that intentional action is necessarily a solution to a Selection Problem.

I'll give two arguments to reject this claim. First, I'll question Wu's analysis of a contrast between reflexes and (intentional) actions—which analysis he uses to support this claim about action as the solution to a Selection Problem. Then I'll point out a threat of vicious regress for the way in which Wu details how Selection Problems arise within intentional action.

¹ Frankfurt (1978/2015) put the point in terms of "guidance," but the idea is essentially the same. For the sake of consistency, I'll continue to use Wu's term "control" to capture this notion.

1. Reflex vs. Action

“To uncover action’s structure,” Wu writes, “contrast action with *reflex*” (21). Reflexes, we agree, are not actions. We can see this point by considering that actions but not reflexes “express rationality, grasp of reasons, skill, or expertise”—and unlike actions, reflexes “are not targets of normative assessment, are never intentional or done freely” (21). This is all explained by a more fundamental difference between reflexes and action, according to Wu: “What is it about reflexes that rules out action? I suggest that it is the *necessitation* of response” (21).

Wu claims reflexes are incompatible with action because they aren’t selected from a space of possible alternatives; in a reflex response, there’s only one causally possible outcome. Consider some example reflexes. When a doctor hits your knee in the right spot, your lower leg kicks. When something flies at your face unexpectedly, you blink. “Pure reflexes rule out agency by eliminating alternative behavioral possibilities. The subject is purely a patient suffering a change” (21).

Although actions are clearly not reflexes, and reflexes never actions, Wu’s explanation is a substantive one, and I don’t think it can be right. There can be a reflex selected from a set of causally possible alternative behaviors; there can also be an action that is the only causally possible behavior in context. I’ll illustrate with some examples.

First, consider reflexes that aren’t necessitated behaviors in the way Wu means. Say that your body such that sometimes when something flies at your face you blink, sometimes you flinch instead, and you never do both. (Add in, if you like, that there are consistent causal mechanisms underlying each of these, that they bypass higher centers of motor control as real reflexes do, and that they happen immediately in response to something flying at your face.)

On Wu’s view, neither such blinking nor such flinching would count as a reflex, because the other could have happened instead. But would this alternative possibility really make either of these responses look less like a reflex?² It seems not. Instead of the possibility of an alternative response, what seems to matter more to a behavior’s being a reflex is the causal pathway that produces it. It matters how the behavior in question is *actually* generated, and not just because how it’s actually generated could have produced an alternative output. It matters, for instance, whether the agent’s own person-level perspective on her situation is involved in this generation process; it matters which bodily and mental systems are involved in producing the behavior; perhaps it even matters whether it is learned or not. Insofar as it seems as though there can be branched possibilities for reflexive responses, it won’t help much to distinguish reflex from action to look at such branching possibility.

Next, consider actions that *are* necessitated behaviors in Wu’s sense. For examples, we can look to highly skilled behaviors for which we accord agents credit and responsibility of the kinds associated with intentional action. Consider an example from tennis. If you train well, you’ll learn rapid responses to your opponent’s shots. For instance, standing close to the net, when a ball approaches you in a certain way, you might practice and engrain one single response to be able to volley the ball back immediately. This response might be highly sensitive to the nature of the input; it might take into account the angle of spin, the velocity of the ball, and the wind around it. But given these features as the input to the question of what to do, there might be really only one thing that it is (causally) possible for you to do: to volley the ball back *just so*. Making such a response causally determined might actually be a successful achievement of years of training. Does this imply, as Wu would have it, that your volleying the ball back *just so* in these conditions is not an action? It doesn’t seem disqualifying at all. If it’s not, then there are actions that are necessitated behaviors, and so actions that are not, by definition, solutions to Selection Problems.

² Thanks to Nico Orlandi for discussion of this point.

In short: it's true that reflexes aren't actions, but Wu mischaracterizes the reasons why. The relevant feature of reflexes that rules out their being actions cannot be whether they are necessitated behaviors, i.e. whether they solve Selection Problems. That's because there are reflexes that solve Selection Problems, and actions that *are* necessitated behaviors that don't solve Selection Problems.

A potential response emerges from what Wu says in this connection about intentions:

"The approach to intention will, I think, jar many readers since it will clash with certain philosophical intuitions and frameworks, with how we ordinarily speak about intention, with folk psychology, and perhaps with introspection. The plea is that missing from all this has been a biological perspective that should be given at least equal weight" (9).

Perhaps the same line is available with respect to action as opposed to intention; Wu might allow a genuine science of action to inform and even to amend our conception of action. But it is not in response to the scientific aspects of Wu's view that he offers this distinction between actions and reflexes in terms of the Selection Problem; this is a foundational point that informs the science itself. For that reason, it doesn't seem right to accept this counterintuitive claim about action and reflex.

2. A Regress of Selection Problems

Even if we allow that action is necessarily a solution to a Selection Problem, and we use this to guide our computational inquiry into action, there's a further problem: a regress looms.

To introduce this regress, start with a foundational point. As Wu and I agree, what you do intentionally is intentional under a certain aspect—an aspect that captures your own conception of what you are doing. This is a descendant of the classic point from G.E.M. Anscombe (1957) and Donald Davidson (1963/2001) that action is intentional only 'under a description.' In Davidson's own famous example, an action involving a certain hand movement of yours might be intentional as *turning on the lights*, but not intentional as *alerting a prowler to your presence* (4).

Insofar as action is necessarily a solution to a Selection Problem, then what you do intentionally in your own terms needs to be one of the options presented in the space of possible behaviors. Let's say that your action is one of high-fiving Sara. Then it must be the case that *high-fiving Sara*, in those terms, is a solution to a Selection Problem given by a space of possible behaviors only one of which is *high-fiving Sara*. (Other alternatives in this space might be *fist-bumping Sara*, *cheering*, etc.) Wu notes that this kind of Selection Problem is precisely solved by your forming an intention, thereby selecting one of the possibilities for your action. This happens at a personal, not sub-personal, level. You are the one who solves this Selection Problem.

But new Selection Problems return with a vengeance once you actually start to act on your intention. According to Wu, your intention to *high-five Sara* guides your arm and body movements in part by raising yet further Selection Problems to be solved—again, by you, not by some sub-personal system. With an intention in mind, there is a *bias* on how you solve these problems, but a bias does not obviate the need for a solution.

Why should we accept that further Selection Problems must arise in this way? Once *one* Selection Problem is solved, giving you an intention, why think that there are further Selection Problems to solve to further *realize* the way that action takes place? It's not clear that all actions are going to have these dynamics once they get going.

To return to the example of your *high-fiving Sara*, we might think that the movement of your arm and hand towards Sara is determinate once it gets going; we need not take it that this movement's evolution is the result of repeated solutions to further Selection Problems that crop up along the way to the high-five. The question here is in part whether the way each action actually gets done needs to

involve you, the agent, responding to yet more inputs by answering yet further questions about what to do as these inputs come in.

Wu emphasizes in various ways that these Selection Problems are solved at the person level. This comes out especially clearly when he talks about how you are ‘active’ in performing an intentional action in part by *continually* engaging in practical reasoning about how to accomplish your action (Chapter 4). This isn’t just a matter of not forgetting what you’re doing, under that description that marks it as intentional (say, *high-fiving Sara*). When you are acting on an intention to *high-five Sara*, your intention is accomplished in part by vigilance in your task—“a propensity to attend to task-relevant targets”—and your ongoing cognitive attention to your task, which is a “subject-level activity of guiding response” (72). Vigilance and attention are both person-level phenomena.

Wu also claims that you need to “fine-tune” your intention as you act (132ff.). Fine-tuning is a matter of forming some even more specific intention as your situation evolves; it involves “additional reasoning such as identifying appropriate targets or determining means to ends, leading to more precise intentions” as well as “maintaining and updating one’s conception of action” (125-6, 132ff.). This further reasoning solves further Selection Problems, and it happens at the person level.

Taken to an extreme, this line would lead us to a vicious regress. Forming an intention to *high-five Sara* is to solve a Selection Problem. Then, Wu says, to act on this intention is to raise yet more Selection Problems about *how* to high-five Sara; these are problems whose solutions constitute yet further intentions. But insofar as they are yet further intentions, by parity of reasoning, executing these fine-tuned intentions should also raise further Selection Problems, and so on forever. On Wu’s view, all of this sets up things for *you* to do. It can make doing anything at all seem miraculous.

The natural response here is to move from a Selection Problem confronting the agent—which might happen when an intention is first formed—to Selection Problems confronting some *sub*-personal systems that actually determine how that intention gets accomplished. This seems like a friendly move, given Wu’s careful thoughts about the essential interplay between automaticity and control in (intentional) action. Part of the power of his discussion of automaticity is that it is able to capture the way in which these realizing processes can be automatic rather than themselves controlled by another intention. Even independently of this point, I wouldn’t think that the way intention biases the connection between inputs and outputs once you’ve started to act must itself take the same shape as your forming your intention. To put it in his terms, I wouldn’t think that any further Selection Problems that arise *while* acting on an intention need be problems for you as the agent, instead of being problems that face some sub-personal computational systems involved in helping you accomplish what you are trying to do.

Conclusion

Wu’s book breaks exciting new ground for inquiry into the nature of intention and intentional action. While his approach to control is distinctive and thoroughly original, there are also some serious challenges available to his fundamental claims about action. I’ve made two such challenges to his claim that action necessarily involves solving a Selection Problem, both in forming an intention and in actually executing that intention.

I’ve argued that action does not need to be a solution to a Selection Problem. Some actions are the only causally possible responses to an input. The contrast between action and reflex doesn’t help establish this necessary link between action and Selection Problems, as Wu misdescribes the crucial contrast between action and reflexes. Some reflexes themselves solve Selection Problems.

A vicious regress threatens Wu's claim that more Selection Problems arise in controlling your action as it happens. If you yourself must solve ongoing Selection Problems to fine-tune your intention over and over again in acting on that intention, it's not clear how to get anything done.

In closing, let me gesture towards a further concern emerging from reflection on issues with Selection Problems. Wu's carefully scientifically grounded methodology promised a theory of control in action that avoids the modal structure of Frankfurt's view. But the model produced by this scientific inquiry is ultimately a matter of iterated solutions to Selection Problems. If control is most fundamentally a matter of solving Selection Problems, we end up with another kind of modally distant control on our hands—one run not quite in terms of what you're poised to do, but rather in terms of what you could have done instead. What you do seems to count as *your* doing mostly because *you* could have done something else. But this is a shame; Wu's exciting new kind of inquiry into intentional action and its realization should be powerful enough to give us a model of control as ongoing causal contact with what you are doing as you are doing it. In further development, it may well offer us just that.

References

- Anscombe, G.E.M. (1957). *Intention*. Oxford: Basil Blackwell.
- Davidson, Donald (1963/2001). "Actions, reasons, and causes." In *Essays on Actions and Events* (Oxford: Clarendon Press), 3-20.
- Frankfurt, Harry (1978/2015). "The Problem of Action." In Jonathan Dancy and Constantine Sandis, eds., *Philosophy of Action: An Anthology* (Malden, MA: Wiley Blackwell), 26-32.
- Marr, David (1982). *Vision*. San Francisco: W.H. Freeman and Company.
- Wu, Wayne (2023). *Movements of the Mind: A Theory of Attention, Intention and Action*. New York: OUP.
- Wu, Wayne (2013). "Mental Action and the Threat of Automaticity." In Andy Clark, Julian Kiverstein and Tillman Vierkant, eds., *Decomposing the Will* (New York: OUP), 244-61.