*Penultimate Draft: Published in *Topoi*, 2021.

Oops! I did it again: The psychology of everyday action slips

Myrto Mylopoulos

Abstract: We have all had the experience of everyday mistakes like distractedly pouring orange juice into our cereal bowl rather than the milk, or inadvertently continuing on our regular route home rather than stopping at the store as we'd planned. These so-called "action slips" (Reason 1984a) are characterized as failures to execute one's intention arising in habitual or highly-learned action sequences. This paper argues that a proper understanding of slips, and thus action more generally, requires an understanding of the control structure that implements an agent's guiding intentions. Central to this structure are motor representations that are active downstream of intention and attentional processes that ensure that they reliably implement the intentions they serve.

1. Introduction

In order to understand how a system works, we must understand how it breaks down. The action control system is no different. This paper will examine a specific type of action control failure that is familiar from everyday life. Some examples will help to illustrate: You plan to stop off at the grocery store after work, but instead drive straight past it, continuing on your regular route home. You arrive at your friend's house for a visit, and instead of ringing the bell, you take out your own key and attempt to unlock the door (James 1890, p. 115). You intend to pour milk on your cereal, but wind up pouring the orange juice beside it into the bowl. Instead of throwing out the wrapper and putting the piece of gum in your mouth, you throw out the gum and shove in the wrapper. You get up to retrieve something from the next room, but upon arriving, you can't recall what it was that you intended to get. These are all instances of what are commonly called 'action slips'—generally classified as failures to act as one intends in the context of habitual or routine action.

Though familiar and quotidian, there are good reasons for why action slips in their many guises are phenomena worthy of attention from cognitive science. The practical benefit of understanding them is easy to appreciate, given the significant consequences that sometimes result from this type of error. Here, of course, what is primarily at issue is not some mild momentary embarrassment or frustration, as in the cases described in the opening paragraph, but errors that yield large-scale disasters in such contexts as airplane piloting, medical procedure, and industrial management (see Reason 1984b, p. 183-184).

Their theoretical import is also manifest, for they can doubtless provide clues as to the inner workings of the representations and processes that coordinate to produce successful action. Cognitive psychologists have been sensitive to these considerations. For some time they have been documenting and analysing both verbal "slips of the tongue" (Fromkin 1973) and nonverbal action slips, which are the focus of this paper (e.g., James 1980; Reason 1984a, 1984b).

Philosophers, too, have started to recognize the theoretical significance of action slips. In particular, some have argued (e.g., Amaya 2013, 2020; see also Peabody 2005) that they serve as important counterexamples to standard philosophical accounts of intentional action that, roughly, attempt to explain what it is for an action to be intentional in terms of causal guidance by an agent's beliefs, desires, and intentions (e.g., Davidson 1980; Mele 1992). Slips, the challenge goes, are intentional actions that are not so guided.

This paper seeks to contribute to theorizing about slips, and thus intentional action more generally, by arguing that a proper understanding of slips requires an understanding of the control structure that implements an agent's guiding intentions. In section 2, I begin by arguing that a general comparison of slips with other cases of failing to act on one's intentions reveals that we must understand the nature of the *motor representations* that serve to implement the goals or plans specified by an agent's intention, as well as the roles that (different types of) attention play in securing the reliability of such implementation. In section 3, I say more about the motor representations that partially implement our intentions: motor schemas. In section 4, I focus on attention, and identify two roles for two different types of attention in ensuring the proper guidance of behaviour by an agent's intentions and motor schemas. In section 5, building on Reason's (1984b) classic discussion of action slips, I show how we can use this understanding of motor schemas and the roles of attention to diagnosis different types of action slip. Overall, I hope the discussion to shed light on precisely what goes wrong in cases of slipping, and on what goes *right* in cases of successful action implementation.

2. The General Character of Action Slips

All action slips have in common a failure to do something that one intends to do. This feature of slips is sometimes taken to be definitional. Thus, Norman (1981) tells us that "[a] slip is a form of human error *defined to be the performance of an action that was not intended*" (p. 1, emphasis mine). Indeed this does capture a salient feature of slipping. When I drive home along my usual route, I fail to act in accordance with my intention to stop off at the grocery store. And when I wind up pouring orange juice into my cereal bowl, I fail to satisfy my intention to pour milk into it instead.

Though this condition is reasonably viewed as necessary for an action to count as a slip, it is clearly not sufficient. For instance, I might intend to take the garbage out after dinner, but find myself lounging on the couch with a book instead. Here what I do is at odds with a previously formed intention, but it need not be an action slip. Perhaps I just changed my mind—I felt the allure of the couch and I decided to take the garbage out tomorrow morning. I am irresolute (see Holton 1999), but I have not slipped.

Notice that in the case just described, I act against my distal or prior intention to take out the garbage, but I do not act against another intention, i.e., the one that prompts me to head to the couch. This is a proximal intention. Proximal intentions can be understood as representations that (i) specify an agent's goal or plan, (ii) unlike distal intentions, are present-directed and tokened in the situation of action, and (iii) function to help structure and guide an action's course to completion (see Fridland 2019 for further discussion of the functional role of proximal intention). I will say more about how such guidance works later on, but for now the point is merely that it is *this* kind of intention that is at issue when we talk of action slips as involving a failure to do what one intends to do: When I slip, I fail to do what I *intend at present to do*. Indeed, this is in keeping with another standard way of describing action slips, i.e., as "errors that violate their own *governing* intentions" (Baars 1992, emphasis mine; see also Amaya 2013, p. 559).[1]

But even failing to act on a proximal intention is not sufficient to constitute a slip. Consider a case where I am standing at a red light with the intention to cross the street when it turns green. I find myself lost in thought to the point that I do not notice the light has changed and thereby fail to act on my intention to cross. Here I have failed to do what I intend at present to do. But I have not slipped. As noted by others (e.g., Amaya 2013, p. 566), slips characteristically involve an *incorrect implementation* of an intention, not merely a *lack* of implementation.

[1] From here onwards, I use the term 'intention' to refer to proximal intentions as described here, unless otherwise indicated.

Not all incorrect implementations of intention are slips, however. Consider a novice gymnast, who intends to perform a cartwheel for the very first time, but ends up stopping short and executing a handstand instead. Though the novice has incorrectly implemented their intention to do a cartwheel, they have not slipped. This is because slips are not cases of failing to implement an intention for a novel action sequence, but rather failures to implement intentions for *routine or habitual action*. Indeed, this is among the more frustrating and puzzling features of action slips. They are failures to do what one knows all too well how to do.

How should we understand what happens when such intentions are implemented incorrectly? Recently, Amaya (2020) has argued that in cases of slipping, one's acquired *habit*s mediate between intention and action resulting in the discordant behaviour. On Amaya's view, habits are "dispositions to execute intentions in ways which have worked in the past and that have been internalized through a process of rehearsal" (p.6-7). Indeed, that slips are cases of acting out of habit can be further supported by the fact that they occur in recognizable patterns, rather than in a random fashion. As Reason (1984b; also Amaya 2020, p. 4) stresses, "… perhaps the most important distinguishing feature of absent-minded slips is that no matter how inappropriate or embarrassing they are, these errors are *nearly always recognisable as belonging to our own personal repertoire of actions.* We know them as our own. Obviously they are not what we intended at that moment, but we can usually see in them a curious kind of logic..." (p.181).

Furthermore, on Amaya's view habits are "ways of acting for reasons" (p. 7). This is what puts slips, he maintains, in the category of *intentional* agency: "By acting habitually, one winds up making an error. Yet, to the extent that one acts habitually in the pursuit of the intention, one also winds up *intentionally* doing various things" (p. 9, emphasis in original). Amaya cashes this out by maintaining that if wee describe the actions that constitute slips in a "piecemeal" fashion, we get actions that are intentional under a description, e.g., reaching for the carton, opening it, pouring its contents into the cereal bowl.

I think we can go further than this, and point out that the actions that constitute slips are to *some* degree sensitive to the content of the agent's governing intention, even though they ultimately fail to implement it. It is due to the content of one's intention to pour something into the cereal bowl that one pours orange juice into it. And it is due to the content of one's intention to put something in one's mouth and throw something else out that one throws out the gum and puts the wrapper in one's mouth. In each of these cases, if the agent's intention had been different, e.g., to eat a spoonful of cereal, to offer the piece of gum to a friend, so too would have been the actions.

In my view, these two features of slips, viz, their sensitivity to the content of intention and their recognizable patterns, are best explained, not in terms of a *disposition or set of dispositions* to execute intentions in ways that have proved successful in the past, as Amaya suggests, but rather in terms of the types of *representations* that constitute (at least partially) the categorical bases of these dispositions. In other words, we must look to the representations that reliably implement our intentions. I will return to this point, but first, we can be more specific still about what exactly a slip involves.

Consider cases of the phenomenon known as "the yips" in sporting performances. In typical cases of yipping, an agent experiences serious detriments to their ability to perform certain actions in their skill domain. A common type of yipping occurs in the context of expert golfers attempting routine putts. In such cases, the golfer forms the intention to sink the putt, but instead their wrist spasms out of control, preventing them from doing so. But while this is a case of incorrectly implementing a "governing" intention to perform a familiar action, it is not a case of slipping. Why not?

The main difference here concerns the role of attention. According to an influential theory of what goes wrong in the yips, known as "explicit monitoring theory", the yips are the result of disruptive attention. The thought here is that performance anxiety can cause one to attend to motor routines that would normally be carried out automatically, and this results in detriments to performance (see Baumeister 1984; Beilock and Carr 2001). Here attention to the details of one's action implementation makes it more likely that one will fail to act correctly.

In cases of slipping, by contrast, the reverse relationship holds between attention and successful implementation. Here a failure to properly allocate attentional resources makes it more, rather than less, likely that one will slip. Indeed, slips are often construed as paradigm cases of "absentmindedness", where this is defined as being "inattentive to ongoing activity" (Manly et al. 1999, p. 661). Empirical evidence based on diary studies (e.g., Reason 1984a, b) and clinical studies with neuropsychologically impaired patients (e.g., Robertson et al. 1997) supports the view that there is an important role for attention to play in preventing slips. For this reason, some have (rightly) maintained that slips "yield important clues as to the role of attention in the guidance of routinized behavior" (Reason 1984a p. 517).

In light of the foregoing, we can see that a full understanding of action slips, and therefore intentional action, requires that we understand: (i) the representations that implement our intentions to perform habitual or routine actions, and (ii) the role of attentional processes in such implementation. I take each of these in turn.

3. Intention Implementation: The Role of Motor Schemas

In the background here is a commonly-adopted model according to which actions are controlled by way of hierarchically arranged goal representations and processes (see Pacherie 2008; Mylopoulos and Pacherie 2017, 2019; Fridland 2019). The levels of the hierarchy are causally structured in a means-end way, such that the representations embedded within a given level specify the means of implementing the goal representations, if any, on the level(s) above them. As such, the goal representations in the hierarchy go from more general to more specific as one moves "downstream" (i.e., lower down the hierarchy) from any given level, except the bottom one.

At the top of the hierarchy are an agent's distal intentions, specifying actions to be done in the future. Typically, though not always, an agent will form proximal intentions that serve to implement their distal intentions. As mentioned, these intentions specify an action to be performed in the present situation. They also interface with downstream goal representations of the motor system that, in turn, serve to implement them. It is in these representations that we are primarily interested.

In order to understand the nature of these motor representations, we can turn an influential theoretical framework in the motor control literature: schema theory (see Schmidt 1975, 2003; Schmidt & Lee 2013; Norman 1981). According to schema theory, the motor representations that are activated by an agent's intention are *motor schemas*, i.e., representations that store information or knowledge in long-term memory about the general form of an action type. They do so by representing its invariant features, i.e., those that remain the same—or with negligible differences—across several token performances of it. These include such features as an action's spatiotemporal organization and sequence, as well as its being directed (or not) towards a target object. For instance, the motor schema corresponding to the action type of reaching-and-grasping-an-object might specify the following elements, in the following order: extend elbow, move arm forward towards object, open grip and straighten fingers, close grip to match object shape and size.

In addition to invariant features, action types and their corresponding schemas also possess *surface features*, which are those that vary across token performances (e.g., direction, force, speed, object type) in a way that allows the action to be tailored to the agent's present context. Surface features are represented in the motor schema by open parameters, the values of which are assigned through a process of *parameterization*. The initial setting and adjustment of these parameters as the movement unfolds is subserved by way of stored information that over time yields mappings of the relations among initial conditions of the body and the agent's environment, parameter values, sensory feedback and action outcomes. Parameterization thus plays the role of "scaling" the representation of the action to the agent's present context in a way that reliably brings about the intended outcome (for further discussion, see Mylopoulos & Pacherie 2017, Pacherie & Mylopoulos 2020; Mylopoulos 2020). A reaching-to-grasp movement might be parameterized differently depending on whether one is reaching for a fragile glass or sturdy mug, for example.

Importantly, once the actions to which an agent's complex motor schemas correspond become well-rehearsed, highly learned sequences—as in the case of our habits and routine actions—the schemas may be executed largely automatically, in a way that does not typically require the formation of new intentions. For instance, an agent's intention to tie their shoelaces might interface with a complex motor schema that specifies the steps involved in the habitual routine that the agent engages in so as to satisfy the intention (e.g., hold ends of shoelaces, form overhand knot, and so on), as well as the order within which to perform them, without the need to form a new intention corresponding to each step.

How does this interfacing between an agent's intention and motor schemas work? While many proposals are available (see, e.g., Butterfill & Sinigaglia 2014; Mylopoulos & Pacherie 2017, Shepherd 2019, Burnston 2017, Ferretti & Caiani 2019), here I adopt one according to which intentions interface with motor schemas by way of the specialized action concepts they deploy in their content, i.e., *executable action concepts (EACs)*. These action concepts differ from those that one might acquire by way of observing instances of others performing certain action types, in that they are acquired primarily on the basis of performing those action types oneself. If one possesses an executable action concept of Φ-ing, then one knows how or is able to Φ. On this view, the deployment of an EAC in the content of an intention serves as a pointer to a specific location in LTM, where the appropriate motor schemas for implementing the corresponding action type are stored.

Once these motor schemas are selected and activated by way of intention-embedded EACs, they must be correctly parameterized and actively maintained. Both of these processes, and therefore the implementation of intention, rely in important ways on attention. In the next section I explain how.

4. Intention Implementation: The Role of Attention

In this section I identify and discuss two roles for attention in intentional action control. The first is that of selecting targets in one's environment to act upon, i.e., providing target information for the parameterization of motor schemas.[2] This is carried out by *selective attention* (also referred to as "endogenous attention"). The second role, which has not been as widely discussed, is that of task

[2] While overlooked for a long time, the importance of attention for intentional action control is being increasingly recognized by philosophers (e.g., Fridland 2014, Wu 2016, Buehler 2019). Some even go so far as to argue that attention is a necessary feature of intentional action (e.g., Wu 2016). Others defend the weaker claim that, while not strictly necessary for intentional action, attention-guided action is a paradigm instance of it, so that a full understanding of the former requires an understanding of the latter (e.g., Buehler 2019; see also Buehler 2021 on the role of the executive system in agentive guidance of visual attention). I will only endorse the weaker claim here.

maintenance, i.e., keeping motor schemas active. This is carried out by *sustained focal attention*. I start with the first role.

In addition to interfacing with motor schemas, a primary function of intention is to guide one's attention in ways that are appropriate to the present action. The main function of the type of attention in which we are interested here—selective attention—is to prioritize and select relevant perceptual information (e.g., object locations, target objects, object features, and events) from the environment for further processing by other cognitive systems, including those subserving action control.

As others have proposed (e.g., Fridland 2014, Wu 2008, 2011, 2016), intention can be seen as a form of influence on attentional selection—not by specifying in its content how attention should be directed, but by automatically biasing certain attentional responses. In particular, intention-guided attention mediates the "coupling" between a target of action (e.g., orange juice) and motor response (e.g., reach-and-grab) by providing detailed information about the spatial location and sensory features of the target for the parameterization of the motor schema. As Wu (2016) explains:

"What attention is providing to the agent is the relevant target but this involves providing content to systems that process that target in a way that sets parameters for and programs an appropriate response. For example, in visually guided reach, one parameter will be spatial location of the attended target so that an appropriate reach can be programmed to that location…" (p. 113)

Important to note, for our purposes, is that the 'coupling' operation of intention-guided attention is non-trivial, since in many cases there will be multiple object representations, only one of which corresponds to the target object. The more objects there are in a visual scene confronting the agent, the more difficult it is to "screen out" non-target stimuli.

This basic finding is explained well by the *biased competition model* of attention, a central assumption of which is that targets and non-targets compete for limited processing resources in visual search (Desimone & Duncan 1995). On this model, attentional selection of competing targets is influenced by a set of biases. Some of these are bottom-up, stimulus-based biases, e.g., bias towards a suddenly appearing object, bias towards non-targets that are visually similar to targets, bias towards a brightly coloured object, and so on. But attentional selection is also served by top-down biases. Indeed, as on the present picture, intention can be seen as a central source of such top-down bias, providing or activating what is known as an *attentional template* for the target object, i.e., a "short-term description [in working memory] of the information currently needed [that] must be used to control competitive bias in the visual system, such that inputs matching that description are favored in the visual cortex" (ibid, p. 200).

The second role of attentional processes in the implementation of intention is that of *sustaining focus* on a present action plan and helping to ensure that it unfolds correctly. *Sustained focal attention* is differentiated from selective attention in terms of (i) its being internally-directed, operating as it does over, e.g., representations in working and long-term memory (see Chun 2011), (ii) its distinct neural localization (see Posner and Peterson 2012), and (iii) its distinct function, which includes task maintenance and inhibition of dominant responses.

Of particular interest for us is its role in maintaining the activation of target motor schemas, i.e., those selected by an EAC, in working memory. Here, too, we may adopt the biased competition model, where the representations competing for limited attentional resources are target and non-target motor schemas. Consider what happens when one performs the Stroop task as an example of such competition. Here one must prioritize the schema that corresponds to naming the colour of the word one is reading rather than that which corresponds to reading the word itself, but both are clearly activated by the task. And one may do so by deploying sustained focal attention towards this particular response, thus "keeping it in mind" (see Miller & Cohen 2001).

.I turn now to explaining how we can make sense of the psychology of action slips by appeal to motor schemas and these two forms of attention.

5. The Psychology of Action Slips

In this section, I explain how a proper understanding of the control structure that implements intention, and in particular motor schemas and selective and sustained focal attention, helps to explain the psychological mechanisms underlying each category of action slip identified by Reason (1984b) in his classic discussion, as laid out in Table 1.

Let's start with the coupling role of selective attention. Certain types of slip quite clearly reveal the importance of this role. Following Reason (1984a), I refer to these as 'wrong object' slips. Consider the example of pouring orange juice over your cereal instead of the milk that is beside it. A credible interpretation of what is happening here is that selective attention fails to appropriately prioritize the perceptual representation of the milk over the competing non-target representation, i.e., the orange juice, and thereby fails to select the correct object for coupling with the appropriate response. This may be the result of the target object and sharing many features with the non-target object, e.g., the shape and size of the carton and even the general spatial location. It may also be the result of attentional capture (i.e., attention driven by bottom-up bias) by the non-target stimulus due to its salience (e.g., brightly coloured orange juice carton; see Theeuwes 1992) or association with past reward (e.g., satisfying a desire for a sweet beverage; see Anderson 2013). In any of these cases, the motor schemas open parameters corresponding to, e.g., the direction of the reach information will be incorrectly filled in, as information about the wrong object is being used.

Moving on, there are at least three kinds of action slip that can be understood, at least in part, as failures of sustained focal attention. The first two pertain to the correct ordering in the steps of an action sequence specified by a motor schema. The first of these is *omission*. This is when steps in the action sequence one is presently pursuing are skipped or left out, with the result that subsequent steps are performed too early in the sequence, or the sequence is terminated prematurely. This kind of slip is evident in the cases of, e.g., leaving before one receives one's change at the grocery store, heading inside without retrieving one's keys from the door after unlocking it, and neglecting to wash one's hair while in the shower as intended.[3]

The second is *repetition*. This refers to cases wherein steps in the action sequence specified by a motor schema are inadvertently repeated. This kind of slip is evident, for example, when one puts sugar in one's coffee twice, instead of just once as planned. Or when one takes one's vitamin pill twice in the morning. In both of these cases, sustained focal attention is required to ensure that one does not "lose one's place" in the action sequence, and that the relevant motor schema remains actively maintained in working memory.

The third kind of slip that can be understood as the product of a lapse in sustained focal attention is that of habit intrusion (also known as a "double capture slip"; see Reason 1984 b, p. 538). This is among the most common types of slip, with some diary studies estimating that it

---

[3] A reviewer helpfully points out that the first example may also be explained, in part, by a failure in selective attention, i.e., to attend to the cashier's outstretched hand offering the change. Doubtless there are other instances of slipping that may also be explained by failures in both of these kinds of attention.

accounts for up to 40% of them (Reason 1984a). This type of slip takes place when at some point in the implementation of an intention, an incongruous motor schema is activated, which possesses overlapping features with the present schema, either in terms of the elements in the action sequence it specifies or in the setting within which it is activated. Both the recency and frequency of the competing, incongruent schema can contribute to its activation.

Take the familiar case of going to the kitchen sink to wash a dish but, lost in thought, finding yourself washing your hands instead. Here there is overlap in the context (i.e., at the sink) and type of task (i.e., turning on the tap to wash something) that helps to explain why a different, competing, non-target schema gets activated at some juncture, as well as the frequency (and perhaps also recency) of this schema.

There are at least two ways to diagnose the relevant breakdown in sustained focal attention here. One is a failure to inhibit the competing schema or response. The other, as I am suggesting, is sustained attention to the present task, which is perhaps the result of the "parallel mental activity" of being engaged in a separate stream of thought (see Reason 1984b, p. 538). Though these roles are not mutually exclusive, and are both important for keeping a task on track, there is some evidence that it is lapses in the latter that are more explanatory and predictive of habit intrusion slips.

One such piece of evidence comes from Manly et al.'s (1999) study attempting to identify behavioural correlates of scores in the Cognitive Failures Questionnaire developed by Broadbent et al. (1982) to assess the occurrence of everyday action slips. Items include questions such as: "Do you start doing one thing at home and get distracted into doing something else 'unintentionally'?" (as reported in Manly et al. 1999, p. 662). Manly et al. developed a task they call the "Sustained Attention to Response Test" (SART) that is one of the rare experimental paradigms to successfully elicit action slips. The task is an instance of what is known as a classic "vigilance" paradigm, which involves the automatization of a simple task response (e.g., pressing a button in response to a stimulus) as well as the need for a certain type of response to a rare target. The SART task involved the visual presentation of target stimuli, which were digits between one and nine. Participants were required to respond as quickly as possible with a key press, unless the target was the digit '3', in which case response was to be withheld.

Manly et al. (1999) found that the sensitivity of performance on the task to scores on the CFQ decreased when demands for sustained attention were weakened (e.g., the inter-target intervals were decreased, the number of trials per block was decreased) even though the requirement to withhold response to an infrequent target was still present. In other words, when demands for sustained attention were weakened, performance on the task failed to distinguish between high (reported more everyday action slips) and low (reported fewer everyday action slips) CFQ scorers. This suggests that strong performance on the SART is primarily determined by the capacity to sustain attention to the task, and that this, in turn, is what explains the finding that stronger performance on the task correlates with lower scores on the CFQ, and vice versa.

Type of slip	Analysis	Examples
Wrong object	Correct motor schema is activated, but object parameter is incorrectly specified due to failure of selective attention.	Pouring orange juice onto one's cereal instead of milk. Throwing out one's gum and keeping the wrapper.
Omission	Correct motor schema is activated, but is incorrectly executed due to failure of	Forgetting to wash one's hair while in the shower.

	sustained focal attention to the task in working memory. Steps are omitted.	Leaving store after purchase without taking change.
Repetition	Correct motor schema is activated, but is incorrectly executed due to failure of sustained focal attention to the task in working memory. Steps are repeated.	Putting sugar into one's coffee twice. Taking one's morning vitamin twice.
Habit intrusion	Incorrect motor schema is activated during a task due to failure of sustained focal attention to the task in working memory.	Going to the kitchen sink to wash a dish but finding yourself washing your hands instead. Removing one's shoes and socks instead of just one's shoes as intended.

Table 1. A non-exhaustive taxonomy of types of action slip based on that provided in Reason (1984b, p. 531), their analysis in terms of interactions between attention and motor schemas, and real-life examples of each, drawn from diary studies (e.g., Reason 1984a, 1984b, Jonsdottir et al. 2007).

5. Conclusion

In this paper, I have proposed a specific understanding of the control structure underlying an agent's intentional action—in particular intention implementation by motor schemas and the roles of attention in securing such implementation—and argued that many everyday action slips can be helpfully explained as varieties of breakdowns in this structure. That they can be so explained by appeal to it lends further theoretical support to the positing of the structure.

Acknowledgements

This work was supported by a Social Sciences and Humanities Research Council of Canada Insight Development Grant on "Consciousness, Intentional Control, and Automaticity" (430-2017-00811).

Sources

Aguilar, J. (2012). Basic causal deviance, action repertoires, and reliability. *Philosophical Issues, 22*(1), 1–19.

Amaya, S. (2013). Slips. *Noûs, 47,* 559–576.

Amaya, S. (2020). Out of habit. *Synthese*, https://doi.org/10.1007/s11229-020-02780-3

Anderson, B. (2013). A value-driven mechanism of Attentional selection. *Journal of Vision*, 13 (3), 1–16.

Anscombe, G. E. M. (1963). *Intention, 2nd Edition.* Oxford: Blackwell Press.

Baars, B. (1992). *Experimental Slips and Human Error: Exploring the Architecture of Volition*. New York, NY: Plenum Press.

Baumeister, R. F. (1984). Choking under pressure: Self-consciousness and paradoxical effects of incentives on skillful performance. *Journal of Personality and Social Psychology, 46*, 610–620.

Beilock, S. L., & Carr, T. H. (2001). On the fragility of skilled performance: What governs choking under pressure? *Journal of Experimental Psychology: General, 130*, 701–725.

Bratman, M. (1984). Two Faces of Intention. *The Philosophical Review, 93*(3), 375 - 405.

Bratman, M. (1987). *Intention, Plans, and Practical reason*. Cambridge, MA: Cambridge University Press.

Broadbent, D. E., Cooper, P. F., FitzGerald, P., & Parkes, K. P. (1982). The Cognitive Failures Questionnaire (CFQ) and its correlates. *British Journal of Clinical Psychology*, 21, 1–16.

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

Buehler, D. (2021). Skilled guidance. *Review of Philosophy and Psychology*. https://doi.org/10.1007/s13164-021-00526-9

Burnston, D.C. (2017). Interface problems in the explanation of action. *Philosophical Explorations*, *20*(2), 242-258.

Butterfill, S.A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, *88*(1), 119-145.

Carrasco, M. 2011. Visual attention: The past 25 years. *Vision Research*, 51, 1484–1525.

Clark, A. J., Parakh, R., & Roy, E. A. (2008). Inducing slips of action: Creating a window into attention failures. Brain and Cognition 67(S1), S17).

Davidson, D. (1980). *Essays on actions and events*. Oxford: Clarendon Press.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222.

Ferretti, G., & Caiani, S.Z. (2019). Solving the interface problem without translation: The same format thesis. *Pacific Philosophical Quarterly*, *100*(1), 301-333.

Fridland, E. (2014). They've lost control: reflections on skill. *Synthese*, 91(12), 2729–2750.

Fridland, E. (2019). Intentions at the interface. *Review of Philosophy and Psychology*, https://doi.org/10.1007/s13164-019-00452-x.

Fromkin, V. (1973). (Ed.) *Speech errors as linguistic evidence.* The Hague: Mouton.

Holton, R. (1999). Intention and weakness of will. *Journal of Philosophy*, 96, 241–262.

Hursthouse, R. (1991). Arational actions. *The Journal of Philosophy, 88*(2), 57–68.

James, W. (1890). *The Principles of Psychology* (Vol. I). New York: Holt.

Jónsdóttir, M. K., Adólfsdóttir, S., Cortez, R. D., Gunnarsdóttir, M., & Gústafsdóttir, Á. H. (2007). A diary study of action slips in healthy individuals. *The Clinical Neuropsychologist* 21(6), 875-883.

Manly T., Robertson I. H., Galloway M., Hawkins K. (1999). The absent mind: further investigations of sustained attention to response. *Neuropsychologia*, 37(6), 661-670.

Mele, A. R. (1992). *Springs of Action: Understanding Intentional Behavior*. Oxford University Press, Oxford

Mele, A., & Moser, P. (1994). Intentional action. *Noûs, 28*(1), 39-68.

Mylopoulos, M. & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, 8(2), 317–336.

Mylopoulos M. & Pacherie E. (2018). Intentions: the dynamic hierarchical model revisited. *WIREs Cognitive Science*, e148.

Mylopoulos, M. (2020). The intelligence of motor control. In *The Routledge Handbook of Philosophy of Skill and Expertise* (Eds.) C. Pavese and E. Fridland. Routledge Press.

Norman, D. A. 1981. Categorization of action slips. *Psychological Review*, 88(1), 1-15.

Pacherie, E. & Mylopoulos, M. (2020). Beyond automaticity: The psychological complexity of skill. *Topoi*, online first.

Pavese, C. (2019). The psychological reality of practical representation. *Philosophical Psychology,* 32(5), 784-821.

Peabody, K. 2005. Trying slips: Can Davidson and Hornsby account for mistakes and slips? *Philosophia*, 33, 173–216.

Posner, M. I. and Peterson, S. E. (2012). The attention system of the human brain: 20 Years After. *Annual Review of Neuroscience*, 35, 73-89.

Reason, J. (1984a). Little slips and big disasters. *Interdisciplinary Science Reviews*, 9(2), 179–189.

Reason, J. (1984b). Lapses of attention in everyday life. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 515–549). Orlando: Academic Press.

Robertson I. H., Manly T., Andrade J., & Baddeley B. T., Yiend J. (1997). 'Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia,* 35(6), 747-758.

Schack, T., Essig, K., Frank, C., Koester, D. (2014). Mental Representation and Motor Imagery Training. *Frontiers in Human Neuroscience*, 8(328), 1-10.

Schmidt, R. A. (1975). A schema theory of discrete motor skill learning. *Psychological Review* 82(4), 225.

Schmidt, R. A. (2003). Motor schema theory after 27 Years: Reflections and implications for a new theory. *Research Quarterly for Exercise and Sport* 74(4), 366–375.

Schmidt, R., & Lee, T. (2014). *Motor Learning and Performance, 5th Edition. with web study guide: From Principles to Application*, Champaign, IL: Human Kinetics.

Shepherd, J. (2019). Skilled action and the double life of intention. *Philosophy and phenomenological research*, *98*(2), 286-305.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics* 51, 599–606. https://doi.org/10.3758/BF03211656

Thompson K.G., Biscoe K.L., Sato T.R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience,*. 25, 9479–9487.

Verwey, W. B. (1996). Buffer loading and chunking in sequential keypressing. *Journal of Experimental Psychology Human Perception and Performance*, 22, 544–562.

Wu, W. (2008). Visual attention, conceptual content, and doing it right. *Mind,* 117 (468), 1003–33.

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.