

The Spontaneousness of Skill and the Impulsivity of Habit

Christos Douskos

Forthcoming in *Synthese*

Abstract

The objective of this paper is to articulate a distinction between habit and bodily skill as different ways of acting without deliberation. I start by elaborating on a distinction between habit and skill as different kinds of dispositions. Then I argue that this distinction has direct implications for the varieties of automaticity exhibited in habitual and skilful bodily acts. The argument suggests that paying close attention to the metaphysics of agency can help to articulate more precisely questions regarding the varieties of automaticity exhibited in overt action.

Keywords: *habit; skill; automaticity; multi-track dispositions;*

I. Most philosophers acknowledge that intentional acts do not always ensue from deliberation. Agents do not always pause to consider their reasons for acting, nor do they need to have any conscious thoughts about how to proceed. But negative characterisations such as these do no more than mark a contrast with cases of ‘full-blooded’ agency, and it is easy to underestimate the differences between the things that are characterised in this negative manner. The various things one does in one’s morning routine, such as making a cup of tea or getting dressed, do not typically ensue from deliberation. And the same holds for the various things one does in playing a musical piece on the violin, such as moving one’s fingers in a certain way. However, I think it will be agreed that there are important differences between these two sorts of case, and that these differences account for the fact that we would ordinarily consider the former case to be a manifestation of a habit, whereas the latter an exercise of a skill. My objective is to articulate some of these differences and trace some of their philosophical implications.

Habit and skill belong to a family of dispositional notions with venerable philosophical legacy, and have come to acquire a variety of quite distinct meanings. Nowadays the idea of skill is used in connection with a very wide range of human capacities. It is sometimes used as the modern equivalent of *techne*, the kind of sophisticated capacities associated with practical expertise, as exercised in art, sport and various professions. On the other end of the spectrum, in action theory the idea of skill is closely related to that of basic acts. These include relatively

simple bodily capacities such as grasping a cup or tying one's shoelaces—the sorts of capacities studied in psychology under the heading of *motor skills*. Similarly, the etymological variants of “habit” have been used in connection with a great variety of quite different ideas. These include behaviors that merely manifest a regular pattern, like going to a concert on Friday evenings, everyday routines like driving home from work via a certain route, compulsive behaviors and addictions, and much more besides. Habit and skill are not manifested only in overt action. However, here I shall focus on bodily skills, by which I mean skills that essentially involve overt bodily movement in the sense that no such skill has been exercised unless some relevant overt movement occurs. Similarly, I shall focus on habits of action, which also essentially involve overt movement, as opposed to habits of thought or emotion.

It would seem that many authors do not think that the differences between habit and skill are significant. A striking example is Bill Pollard, who has offered the only comprehensive treatment of habitual action in contemporary action-theoretic terms. Pollard's central examples of habit formation include learning a musical instrument or a sport. And he lists among its characteristic features the increased sensitivity to the peculiarities of the circumstances as well as the improvement in accuracy and dexterity of bodily movement. But these are features of what we would ordinarily take to be the acquisition of a skill, rather than the formation of a habit. In Pollard's book (2008) on habit, skills are often appealed to but they do not get a separate treatment, not even as a species of habit. Similarly, in one of the very few precursors to Pollard's discussion, Nathan Brett (1981) examines the ‘habits’ involved in speaking, walking, typewriting or playing a musical instrument, and argues that these are not stereotyped mindless behavior because they involve attention and variation. But again, the question is whether these activities as such are properly considered habitual to start with. Generally, the tendency to overlook the differences between habit and skill can be testified in several works in analytic philosophy. This is because analytic philosophers are expected to display sensitivity to the ordinary use of a term, but they often use “habit” when discussing typical examples of skills, leaving one with the impression that there is not much to choose between the two.¹

¹ Setiya (2007: 55) talks of the ‘habit’ involved in playing a musical instrument, and uses “habit” and “skill” interchangeably in the same passage; Velleman (2007: 139) takes the ability to drive to consist in a set of habits. Papineau (2015: 399) collapses the distinction in the opposite direction. He adopts a standard view on which basic acts are exercises of skills. But then he claims that taking a familiar route when driving back home can become a

Much could be said about the sources of this use of “habit”, which has a long history and has also been adopted by several psychologists.² The fact that both habit and skill arise from a process of habituation, that they are species of the traditional category of *habitus*, surely plays a role. Moreover, a habit is formed by repeatedly doing the same sort of thing in certain circumstances, initially on the basis of conscious thought and deliberation, but the thought, attention and concomitant sense of effort diminish as the habit gets more entrenched, as does the awareness of certain aspects of the action. But these are equally salient features of skill development, at least on many prominent accounts (e.g. Dreyfus and Dreyfus 1986). These presumed common features of habit and skill are often brought together under the heading of acquired *automaticity*: the various things one does in exercising a skill and in manifesting a habit likewise ensue largely from sub-personal processes established by habituation. In this they both stand in contrast to cases of ‘full-blooded’ agency. And to the extent that one focuses primarily on this contrast, the differences between habit and skill are likely to be disregarded.

Philosophers more mindful of the tradition, however, stress that the process of habituation ensues in different kinds of aptitudes. The best-known discussion of the differences between habit and skill is to be found in Ryle (1949: ch.2). One prong of Ryle’s attack to the “intellectualist legend” is the observation that we consider the various things one does in the course of skilled activities to be manifestations of intelligence, without supposing that such activities are always guided by conscious thought or deliberation. However, Ryle does not claim that every act that does not ensue from thought or deliberation is a manifestation of intelligence. He distinguishes between “intelligent capacities”, which he takes to be exercises of a distinctive kind of knowledge, knowing-*how*, from “mere habits”. Habits are rote behaviour, each manifestation of which is a “mere replica of predecessors” (1949: 42). In this sense, they are

“learned basic action”, which is thus on a par with—to use his examples again—hitting a forehand slice in tennis or tying one’s shoelaces (or indeed, we might add, to the skills exercised in driving itself). However, we shall see that the former is a paradigm example of habit in the relevant literature. This suggests that assimilating the two cases stands in the way of providing a credible account of basic acts (see Douskos 2017a). For a critique of Pollard’s account of habit see Douskos (2017b).

² This is clear in two foundational works: William (James 1981/1890: ch.4), Bryan & Harter (1899). Nowadays psychologists are more attentive to the pre-theoretical usage of “habit” and “skill”, but this does not mean that the distinction is always clearly observed. One reason is that, as Wood and Rüniger (2016: 292) notice, “automaticity” is often used interchangeably with “habit”. This may lead one to use “habit” as a generic term that covers skills.

single-track dispositions, that is, dispositions with a relatively uniform manifestation, which are acquired by “drill or conditioning”. Ryle writes: “When we describe someone as doing something by pure or blind habit, we mean that he does it automatically and without having to mind what he is doing. He does not exercise care, vigilance, or criticism” (1949: 42). By contrast, intelligent capacities are multi-track dispositions: in exercising a skill one manifests a sensitivity to the circumstances, which may require one to do quite different things in order to achieve one’s objective on each occasion; this is one respect in which skill exercises are “indefinitely heterogeneous” (1949: 44). So exercises of skill involve innovation, whereas manifestations of habit involve sheer repetition. Moreover, in learning a skill, agents adopt a critical attitude. They identify their mistakes, and figure out ways to improve; they monitor their progress. Thus each exercise of a skill “modifies its predecessors”: one improves by practicing. By contrast, agents do not practice their habits in order to improve on them. (For a related distinction between skill and “routine”, see Annas [2011a; 2011b]).

The nature of knowing-how, or Ryle’s intelligent capacities, is the topic of an important current debate. But Ryle and his intellectualist critics (Stanley and Krakauer [2013]) agree that an account of the distinction between habit and skill is required to properly demarcate genuine cases of knowing-how from other fruit of habituation with which it is liable to be assimilated. As Ryle (1949: 42) observes, “the common assumption that all second natures are mere habits obliterates distinctions which are of cardinal importance for the inquiries in which we are engaged”. The inquiries that preoccupy Ryle have also resurfaced in the form of questions about of the relation between skill and knowledge, or the extent to which skilful bodily acts can manifest intelligence. But the “common assumption” which Ryle thinks stands in the way is still quite widespread.

In this paper I shall not inquire on the implications of the distinction I shall (partially) articulate for these and other issues. My objectives are limited. First, I shall suggest that there are different ways of acting without deliberation, and that these can be systematically distinguished on the basis of certain characteristic features. I call this the *Diversity Claim*. Second, I argue that skill exhibits at least one such feature that is not exhibited in habit. (I consider additional differentiating features in other work). Hence habit and bodily skill are two ways of acting without deliberation. I mark this distinction between kinds of automaticity by saying that habits exhibit *impulsivity*, while skills exhibit *spontaneousness*, though both terms have unwanted

connotations.³ In arguing for these claims I rely on a metaphysical distinction between propensities and capacities, a distinction that has significant similarities to Ryle's distinction between single-track and multi-track dispositions, as well as on certain action-theoretic ideas, especially regarding the ontology of action. So I hope to convey a sense of how certain metaphysical considerations are important when addressing questions about the kinds of automaticity exhibited in overt action.

In section [II] I explain that in psychological literature different kinds of automaticity are distinguished on the basis of certain features. In the sections that follow I shall limit the discussion to two features. (i) (Absence of) Deliberation and intention; (ii) (Diminishing) Attention. In section [III] I consider and reject a simple way drawing the distinction between impulsivity and spontaneousness solely on the basis of feature (i). This will enable us to focus the discussion on a specific problem-case, which I use to motivate the Diversity Claim. In section [IV] I articulate a distinction in the identity conditions of propensities (habit) and capacities (skill). This distinction implies that, in contrast to habits, skills exhibit a characteristic variability in performance. In section [V] I argue that once skill and habit are distinguished in this way, it can be shown that on a conception of attention as goal-dependent sensitivity, attention is an essential feature of the spontaneousness of skill, but the impulsivity of habit dispenses with it. We thus reach the Diversity Claim. Section [VI] contains some concluding remarks and points to some further ways in which the distinction can be developed.

Before proceeding, let me explain the action-theoretic terminology I am using. In line with the coarse view of action individuation, I take actions to be temporal particulars (events or processes) that can be variously specified (described).⁴ Specifications correspond to properties of these action-particulars, or to the various things one does in acting (doables). Actions are habitual or skillful only under some specification(s). Responding to one's e-mails in the morning may be habitual for an agent, but it does not follow that it is done in the same way every morning: one may do this using her desktop computer, her tablet, her phone, and writes different things each time. Hence the lower-level specifications are different each time but the habit

³ The term "impulsivity" is inspired by Garnder (2015), who argues that habit generates an impulse to act. But I do not mean to endorse any assumptions regarding the nature of impulses.

⁴ Important contributions to the literature on action-individuation include Anscombe (1963), Davidson (1981: essays 1 to 8).

manifested is the same. Some of these lower-level specifications, such as typing a word or phrase, are specifications under which the agent exercises a bodily skill—her typing skills. So an action may well be habitual under some specification(s) while skillful under some other(s).

Whether or not a thing one does exhibits automaticity depends on the occasion. This point is taken for granted in the literature on basic acts: on some occasion one might type her password either “at one go” (i.e. by the operation of automaticity), while on some other occasion she attends to each key she presses in sequence. Thus “is basic” predicates things done on occasion, or specifications of actions (Hornsby 1980: 69 and Enc 2003: 54). Similarly, on some occasion the agent may carefully deliberate about whether or not to do something she is in the habit of doing, and usually does with without any thought or deliberation (i.e. by the operation of automaticity). So I shall use “act” as equivalent to “thing done on occasion” or “action under a specification”, for this allows me to speak simply of a habitual or skillful *act* as exhibiting automaticity.

Action specifications can be related in several ways, two of which will be central in what follows. On the one hand, two specifications may specify the same action at different levels of detail. One waved goodbye by making a hand movement with specific spatiotemporal properties. Here the associated specifications specify the same event or process. On the other hand, an action might be composed of what are intuitively discrete sub-actions in sequential order: think of making a cup of tea or a cake, for instance. We may call these *composite actions*. Here the ‘by’ relation may also pertain to ‘same level’ specifications of discrete actions.

A composite action is skilled or habitual under some specification just in case each of its component actions is, under *at least one* specification, a manifestation of a habit or the exercise of a skill respectively. This is to say that we may speak of a composite *acts* as being skilful or habitual—a *habitual routine*, in the latter case. A non-composite skilful bodily act consists in a bounded stretch of *bodily activity*. Roughly speaking, notion of *activity* here implies a continuity that forbids any non-arbitrary segmentation of the act into further component acts (see Hornsby 2013); think of the unfolding movement involved in reaching for a cup, for instance, or the way one ties her shoelaces once she has acquired the relevant skills. The discreteness of actions, on which the composite non-composite distinction relies, is not a clear-cut issue. And as the examples above illustrate, in bodily skills the discreteness of actions is sometimes affected by the operation of automaticity.

With these preliminaries in place, let us start by explaining the Diversity Claim.

II. The Diversity Claim

Acquired automaticity is such a multifaceted aspect of human nature that is hard to characterise it precisely. In psychological literature what is meant by “automaticity” varies according to a host of theoretical assumptions, as well as the phenomenon under investigation. Indeed, some psychologists dispute that “automaticity” expresses a unitary concept. It has thus become common to approach the idea of automaticity by singling out its characteristic features. Among these are the *absence of deliberation and intention*, *diminishing attention* and the concomitant *sense of effort*, *diminishing awareness*, as well as *improved speed*, *accuracy* and *effectiveness* in acting. In recent psychological literature these and other features do not figure in definitions of automaticity as individually necessary and jointly sufficient conditions. It is generally acknowledged that there is no point in trying to define an all-purpose concept of automaticity. Rather, they are features the different combinations of which serve to distinguish between and define different *kinds* or *varieties* of automaticity. (For a defence of this ‘feature-based’ approach to automaticity, a discussion of the different features and how they dissociate, see Moors & De Houwer [2006; 2007]; for an influential earlier discussion, see Bargh [1994]. Snow [2006] and Fridland [forthcoming a] discuss different philosophical implications of decompositional approaches to automaticity).

The features that I take to be the most critical for the distinction between the automaticity of habit and skill are the absence of *deliberation and intention*, *diminishing attention*, and *improvement in performance*. However, in this paper I shall limit the discussion to the first two. (I shall elaborate on these features in due course; the notion of intention is closely related to that of a *goal*; I shall comment on this relation below). Now suppose we pick out the class of acts that exhibit the first feature: the acts that do not ensue from deliberation. One may still distinguish between the kinds of automaticity exhibited by acts in this class on the basis of the other features. Thus I shall argue that since both habitual and skilful acts exhibit the absence of deliberation, or at least since they often do so, we should distinguish between the *kinds* of automaticity habit and skill exhibit precisely because they differ with respect to the feature of attention. We may thus say that habit and skill exhibit different ways of acting without deliberation: habit exhibits *impulsivity* whereas skill exhibits *spontaneousness*. This is a case of the Diversity Claim.

My discussion is circumscribed in several respects. Several other features of automaticity are relevant to the distinction between spontaneousness and impulsivity. Moreover, the features I do discuss can be understood in quite different ways, and I shall rely on specific accounts of attention and deliberation. So in selecting certain features and their construal, the following discussion will neglect certain aspects in which habit and skill might be similar, while at the same time fail to identify additional differences between them, as well as more subtle differences between their different varieties. However, it is not my objective to provide a comprehensive account of the phenomena under investigation. In order to support the Diversity Claim, it is sufficient to show that the automaticity of habit differs from the automaticity of skill with respect to *at least one* feature under some well-established understanding of it.

III. Absence of Deliberation: the “Why?” and “How?” question

In this section I shall consider a way of drawing the distinction between impulsivity and spontaneousness solely on the basis of feature (i), the absence of deliberation. We shall see that this is not satisfactory because it does not apply to certain sort of case. This will enable me to limit the scope of the discussion to this problem-case, which I use to motivate the Diversity Claim.

Following Owens (2011), I take deliberation to be a conscious intentional activity, the aim of which is to resolve a question. Deliberation is conscious because it involves phenomenal attention. Further, deliberation is an activity in the sense that the deliberator is the author or agent, not just a subject to which deliberation occurs. And it is an activity one intentionally engages in, since it is conducted with a specific aim: to resolve a certain question. There is no such thing as deliberating unintentionally. Practical deliberation terminates with the formation of intentions, that is, by deciding or making up one’s mind about what to do and/or how to do it.

I assume that the absence of deliberation and intention is an invariable feature of the automaticity of skill and habit. I am not aware of anyone who denies this. This assumption should be distinguished from the stronger claim that skill and habit invariably involve automaticity. In psychological literature it is common to build automaticity in the definition of habit and/or skill.⁵ Some philosophers agree.⁶ However, the stronger claim might well be

⁵ Some examples: Ouellette and Wood 1998: 55 Verplanken and Wood 2006: 91; Graybiel 2008: 361; Wood and Neal 2009: 580; Verplanken and Aarts 2011: 104; Gardner *et al.* 2012: 2; Wood, Labrecque, Lin & Runger 2014.

disputed, and it is orthogonal to my purposes here. The claim that the *automaticity* of habit or skill invariably involves certain features does not by itself entail that these are essential features of habit or skill as such.

The claim that the automaticity habit and skill involves the absence of deliberation, however, requires further elaboration. The aim of practical deliberation is to resolve a question, but not always the *same* question: practical deliberation can specify one's objectives, as well as one's means. Correspondingly, for any act of A-ing we may ask *why* the agent A-ed, as well as *how* she A-ed. According to a standard semantics of questions, the meaning of a question is determined by the set of its possible answers. The domain of possible answers to a "Why?" question figures *reasons*. The domain of a "How?" question figures *ways*.⁷

This distinction corresponds to an important distinction in the typical explanatory role of habit and skill ascriptions. Just as reason explanations, a habit ascription explains *why* one A-ed, or undertook to A, on some occasion. Mary has just finished her dinner and she is washing her teeth. We may explain why Mary is washing her teeth now either by citing her habit of brushing her teeth after dinner, or by pointing to her reason for doing so. So habit ascriptions have the same explanatory role as ascriptions of reasons: they both explain why an agent came to do something on some occasion.

By contrast, an ascription of skill in A-ing does not typically explain *why* one undertook to A on some occasion: being a skilled sniper is no reason to start shooting people here and now. Skill ascriptions typically address questions of method or technique. Lacking the relevant skills, one might have to think carefully about *how* to move one's fingers in order to tie one's shoelaces.

These authors build automaticity in their definitions of habit, and claim that the automaticity of habit involves the absence of deliberation and intention. Regarding the automaticity of skill, see Logan 1985 for a review and discussion of earlier literature.

⁶ Pollard 2008: 55-56, Appendix, builds automaticity in the definition of habitual action, and the absence of deliberation in his definition of automaticity. Brett 1981: 357, 364 takes habit to be a form of automaticity, and claims that it involves the absence of deliberation. Regarding skill, the claim that the lower-level aspects of bodily actions are to be explained by the operation of automaticity is implied by the very idea of basic acts (e.g Enc 2003: ch.2; Papineau 2013; see also Wu 2015: 4); but we should keep in mind that not all bodily skills are plausibly basic. For discussion regarding the automaticity of bodily skill, see Fridland forthcoming a; for considerations against the view that sophisticated skills invariably involve automaticity, see Montero 2016.

⁷ This semantics originates in Hamblin (1958) and has been developed in several ways by subsequent semanticists.

But once the skill is acquired, the *same* “How?” question is resolved without deliberation. It is the agent’s skills that account for the exact way she moves her fingers, a way of which she might have only a demonstrative conception. So while in some cases the “How?” question can be resolved by figuring out the means (ways) in deliberation, in other cases the same question will be resolved by the agent’s skill. Generally, if S is skilled in A-ing, she will know how to go about—a suitable *way* to A—in a range of varying circumstances: ascriptions of skill typically address a “How?” question, the answer to which figures ways of acting. Ways are properties of actions that may be specified either by descriptive or demonstrative expressions (see also Stanley [2011: 58]).

The full answer to the “Why?” question reached in deliberation will specify a relation: “(to) A for reason *r*”. The answer to the “How?” question will be “(to) A by way *w*”. The answers to these two questions specify different relations. Once these questions are distinguished, we can see that automaticity or deliberation can address only one of them. One can decide to tie her shoelaces (deliberation addresses the “Why?” question) and let her skills take care of the way she does it (automaticity addresses the “How?” question). But is it also possible that automaticity addresses both questions. One habitually picks up the phone when it rings. Here her habit may explain *why* one picks up the phone while her bodily skill determines the exact bodily trajectories by way of which one does it.

This suggests a simple way in which spontaneousness and impulsivity can be distinguished merely on the basis of feature (i), the absence of deliberation. The impulsivity of habit is responsible for eliciting an act of A-ing in response to certain circumstances, explaining *why* one A-ed on some occasion. By contrast, the spontaneousness of skill explains *how* one A’s, so is responsible for the guidance of the on-going act. They both exhibit the absence of deliberation, but with respect to different questions. And by resolving different questions they determine different relations.

However, the suggestion above does not suffice to distinguish impulsivity from spontaneousness. For one thing, impulsivity may also be operative with respect to the “How?” question. Just like going to work early in the morning is a habit, one may also be in the habit of going to work by car. One does not have to deliberate about how to go to work each day. So one may equally have a habit of A-ing by B-ing, and B is the answer to *how* she A’s. Moreover, the habit of going home via a certain route comprises a sequence of acts, taking this or that turn etc.

Let us call *habitual routines* those cases where the automaticity of habit takes care of at least one means or step. Habitual routines are usually composite acts: they consist in multiple actions in sequential order, actions that are habitual under some specification(s).⁸ Here impulsivity guides an on-going act.

Conversely, in certain cases the spontaneousness of skill can explain *why* an agent came to do something. Think of how the various consecutive responses at a game of tennis are determined, for instance. Here the agent's tennis skill may explain *why* she has opted to respond by some component act, say a forehand slice to the far right corner, as opposed to doing something else. (The point is not restricted to composite acts). When we take into account the complex structure of human action, it becomes clear that distinguishing impulsivity from spontaneousness on the basis of the distinction between the "Why?" and "How?" question does not provide a comprehensive solution.⁹

This sets the stage for formulating a problem that can be used to motivate the Diversity Claim. We exercise a bodily skill by doing something or other. One returns a tennis volley by a forehand slice, and hits a forehand slice by moving in a certain way. Similarly, in habitual routines one habitually takes certain means: one goes back home from work by taking some route. These are structurally analogous: one A-ed by B-ing. And in both cases the answer to the "How to A?" question is determined by the operation of automaticity. Just as with skills, in habitual routines automaticity accounts for *how* one acts, the guidance of the on-going act. However, according to the Diversity Claim—and common sense—there is still a difference between the two cases.

In what follows I seek to articulate this difference. The first step is to take a closer look at the metaphysics of habit and skill.

⁸ Psychologists have long recognized that the automaticity of habit guides "routine action sequences". See for instance Graybiel 2008: 361; Gardner 2014: 280; Wood & Neal 2007: 851; Cooper & Shallice 2006; Bargh & Gollwitzer 1994: 78. Wood & Runger (2016: 292) observe that many authors take the sequencing of multiple acts to be a distinguishing feature of the automaticity of habit (e.g. Dezfouli, Lingawi and Balleine 2014).

⁹ This does not conflict with the point that habit and skill explanations typically address the "Why?" and "How?" question respectively. When an agent is A-ing by B-ing, the "How to A?" question corresponds to a "Why to B?" question (Anscombe 1963). The question is different but the relation specified in the answer is the same. However, in what follows I shall articulate the problem I use to motivate the Diversity Claim in terms of the "How?" question. This is partly a matter of convenience.

IV. Identity conditions for propensities and capacities

A habit is specified as ‘A-ing when in circumstances *c*’ (brushing one’s teeth after dinner), ‘B-ing when A-ing’ (taking a certain route when going back home), or equivalently, ‘A-ing by B-ing’ (getting back home by taking a certain route). Thus the idea of eliciting *circumstances* is to be broadly construed. It does not include only features of one’s environment. In habitual routines the performance of an act may constitute itself the circumstances where one takes the next step. Or, the decision to do something, say a cup of tea, may constitute itself the circumstances in which one habitually takes a certain means.¹⁰

The reason that the canonical specification of a habit has to include both the doable (specification) which constitutes its manifestation and the eliciting circumstances (broadly construed) has to do with the habit’s explanatory role: given an agent’s habit of A-ing in *c*, the fact that *c* obtain on some occasion explains why the agent A-ed on that occasion. We may explain why John had a beer with his dinner tonight by citing his habit of having a beer with dinner. We may explain why an agent took a certain route yesterday evening by citing his habit of coming back from work by taking that route. This is why in psychological literature the claim that the manifestation of a habit is elicited by circumstances one has repeatedly encountered in the past is embedded in almost all characterisations of habit.¹¹ The fact that appreciation of *both* the manifestation and the eliciting circumstances (broadly construed) is necessary to provide for a habit’s explanatory force implies that both of these enter into a habit’s identity conditions. Having a beer with one’s dinner is not the same habit as having a beer with one’s breakfast, because only the former can explain why one had a beer with dinner last night. Going to work by

¹⁰ Both of these are countenanced in psychological literature. Orbell and Verplanken (2010: 374) argue that the circumstances that may elicit a habitual act include “preceding actions in a sequence”. Aarts and Dijksterhuis (2000a, 2000b) have argued that when one is in the habit of A-ing by B-ing, forming intention to A automatically activates the goal to B. I shall come back to the difference between these below.

¹¹ For example, Wood & Neal (2007: 845) claim that “habits are repeated responses that come to be cued by recurring features of the context”. Orban & Verplanken (2010: 374) claim that “habits are triggered by features of the context”, and are thus a form of “cue contingent automaticity”. In their review article on the psychology of habit, Wood & R nger (2016: 292) take “activation by recurring context cues” to be a distinctive feature of the kind of automaticity habit exhibits. In the eight definitions of habit listed in Gardner (2015: 279), all but one make reference to stable features of the context as automatically eliciting habitual acts.

train is a different habit from going to work by bus, because only the former habit can explain why the agent is boarding the train now. If the habit of A-ing in *c* were identified merely as a habit of A-ing, it would fail to explain why the agent A-ed on some occasion, as opposed to some other where *c* did not obtain.¹²

Human habits differ from typical physical dispositions in that an agent has a habit only if she has repeatedly manifested that habit in the past: she has repeatedly A-ed in *these* circumstances or in *this* way. So as long as the same habit is manifested, what one does *and* why/how is the same across different occasions. Mary is not in the habit of having a beer with her breakfast no matter how many times she had a beer with her dinner. John is not in the habit of coming back home from work by this route no matter how many times he came back via some other route. Pollard (2008: 49-55; 207) makes essentially the same point when he claims that “S has the habit of A-ing in *c*” entails “S has A-ed a sufficient number of times in *c*”. Repetition or invariance with respect to what answers the relevant question is built into the idea of habit. As long as one is manifesting a given habit, the answer to the relevant question is settled; and the relation specified in the answer is invariant across the different occasions where the same habit is manifested.

This does not imply that *other* specifications of the actions in which a habit is manifested will also be the same across occasions. When John manifests his habit of A-ing by B-ing, say driving back from work by taking route *r*, he invariably A’s by B-ing. But he B’s by doing a multitude of different lower-level acts on each occasion: the different ways he handles the wheel, the different lanes he takes etc. However, these answer a *different* (“How to B?”) question and are not part of the identity of the habit of A-ing by B-ing. The present claim is that actions that manifest some habit are invariant only with respect to the specifications that are part of that habit’s identity conditions (these may be more than two in the case of complex habitual routines). Ryle’s (1949: 42) claim that actions that manifest a habit are “replicas of their predecessors” is clearly exaggerated.

¹² We should not be misled by the fact that in everyday habit ascriptions circumstances might be cited somehow obliquely: “She is in the habit of putting sugar in her tea” (i.e. when she is having tea). Moreover, in everyday explanations circumstances are part of the conversational background, and hence may not be explicitly mentioned. In other cases they might be too generic to afford a non-trivial specification and thus likely to be omitted in a habit ascriptions for pragmatic reasons.

I turn now to the identity conditions of skills. To start with, it is clear that answers to the “Why?” question cannot figure in the identity conditions of a skill: the same skill can be exercised for very different reasons. One may play the lyre to entertain her audience, or just for practice. One can use her medical skills to cure or harm. Any skill can, at least in principle, put in the service of different projects. The answer to the “Why?” question will be different across occasions. So reasons for exercising a skill cannot be part of the skill’s identity conditions, in contrast to (non-routine) habits where the eliciting circumstances are part of the habit’s identity conditions. Nevertheless, one might think that the identity conditions of skills have to include answers to the “How?” question, in which case there would still be an analogy between the identity conditions of skills and habitual routines. I shall argue that this is not the case either.

The first thing to notice is that, as a matter of fact, when A-ing is the exercise of a skill, and one has repeatedly A-ed, the answer to the “How did she A?” question is typically not the same across different occasions. This is evident in the case of sophisticated skills such as tennis playing, where one does a very different sequence of things on each occasion as required by the circumstances. One exercises a skill in a composite act of A-ing by a different combination of the component acts. But even in simpler bodily skills, such as the non-composite acts of grasping a cup or catching a flying ball, the initial configuration of the relevant objects and bodily parts is hardly ever the same. So the lower-level bodily activity that resolves the “How to A?” question, the ways of A-ing which are the bodily trajectories in this case, will be different across the different occasions one A’s. Since on various occasions where the agent’s A’s the answer to the “How?” question is different, the specification figuring in this answer cannot be part of a skill’s identity conditions.

One might object that some simple skills of some agents may exhibit invariance with respect to what answers the “How?” question. Consider, for instance, tying one’s shoelaces, or putting one’s signature. These involve a more or less invariant body-environment configuration. Indeed, developing a signature seems to consist precisely in managing to minimize variation: one had better try to resolve the “How?” question in the very same way.

There are several ways to deal with this objection, based on points made in current literature. First, one might insist that there still exist minute variations in the ways one puts one’s signature or ties one’s shoelaces, though these might not be readily observable with the naked eye.¹³

¹³ This answer might appeal to research in psychology, which indicates that even if we hold everything else fixed the

Second, one might retreat to a more restricted conception of skill, withholding the title of skill for more sophisticated capacities that exhibit the required variability.¹⁴ None of this gets to the heart of the matter, however. The identity claim above is a modal claim. In the case of habitual routine, the specification figuring in the answer to the “How?” question is part of the habit’s identity conditions: one would not have exercised her habit of A-ing by B-ing if she A-ed in some other way. By contrast, suppose that on some occasion someone signed or tied her shoelaces in a slightly different way, that is, by a slightly different bodily movement. The point is that she would still have exercised the *same* skill even though the answer to the “How?” question is different. She did not exercise it as optimally as she usually does, presumably. But surely, this possibility must be allowed for any skill. So the contrast between habit and skill holds even on the (disputable) assumption that some skills of some people do not *actually* exhibit variability with respect to the “How?” question. Even though improvement in some skills involves obtaining consistency across different occasions in the relevant aspects of the performance, one exercises the same skill as one improves on it (that is, as one gradually comes to resolve the “How?” question in a better way, say by a more accurate movement), whether she exercises it more or less optimally, or indeed when she fails. And this requires that in exercising a skill one is always confronted with an open “How?” question, even though one might manage to resolve this question in the same way on multiple occasions.

bodily trajectories will exhibit substantial variability across occasions. The inherent “trajectory redundancy” is a central feature of the motor system (Todorov & Jordan 2002; Haith & Krakauer 2013a; see Fridland forthcoming b for discussion). It is true that on this approach to motor control minimizing variability in the relevant respects is a central aspect of skill development. But this does not imply that variability can be completely eliminated in human agents; and in any case, this claim is restricted to what counts as ‘relevant respects’, whereas the variability claim at issue here includes all the spatiotemporal properties of overt movement.

¹⁴ Fridland (2014) distinguishes between skills and mere (bodily) abilities. Annas (2011a) distinguishes between skill (understood as practical expertise) and “routine” (understood as akin to habit). But Annas claims that simple bodily capacities, such as tying one’s shoelaces (to use her example [2011: 18fn.3]) sides with ‘routine’. While what counts as skill for these authors clearly exhibits variability, I do not think that these distinctions withstand scrutiny. Nevertheless, I do not want to rule out the possibility that in some cases the detailed bodily movements or stretches of bodily activity by way of which one acts are so inflexible and reflex-like that they should be explained by the operation of impulsivity (there is some evidence pointing to this direction; see Haith & Krakauer 2013b: 15-16 on “motor habit”; see also the discussion “habitual lags” in Toner, Montero & Moran 2015). The claims defended here imply no constraints on which kind of automaticity can be operative at which level in the structure of action.

Thus on any occasion where one has exercised a bodily skill in A-ing, it will be true that one A-ed by B-ing, for it is metaphysically necessary that there will be something specific by way of which one has A-ed, say a bodily movement with certain spatiotemporal properties. But there is no specific thing which one *invariably* has to do when exercising her skills in A-ing. So the best we can do is to say that, necessarily, in exercising a skill in A-ing one A's by *w*-ing, where *w* is a variable ranging over ways of acting the value of which is determined on occasion. By contrast, when one is in the habit of A-ing by B-ing, both A and B provide for the identity conditions of that habit. This is to say that on any occasion in which one manifests *that* habit, one A's in the same way: by B-ing.

In other words, skill exercise is indefinitely heterogeneous or variable, at least potentially. By contrast, habits involve sheer repetition and inflexible routine, though of course the specifications under which the action is not habitual will vary. This is the kernel of truth in Ryle's distinction between skill and habit in terms of single/multi-track dispositions (Ryle meant additional things by it). Habit is a propensity to act in the same way in response to certain specific circumstances (broadly construed). Skill is a capacity to do whatever it takes to achieve one's objective in a range of different circumstances, though in certain cases the circumstances and ensuing act(s) might turn out to be the same.

The discussion above points to a crucial difference in the functional demands placed on the two kinds of automaticity. Deliberation seeks to answer a "Why/How to A?" question. The answer to this question is a relation: (to) A by B-ing, or (to) A because *c* (where *c* is, or provides, a reason). When one acts *without* deliberation, this relation is specified by the operation of automaticity. So automaticity and practical deliberation specify the answer to a practical question: they likewise explain why and/or how an agent acted as she did, and the answer to that question denotes a relation. The difference between impulsivity and spontaneousness is this. Impulsivity invariably elicits the same kind of act in response to similar circumstances. If one is in the habit of A-ing by B-ing, and her objective is to A, impulsivity invariably elicits B as a means (or, in the non-routine case, impulsivity invariably elicits A in response to *c*). Impulsivity pertains to invariable relations. By contrast, the function of spontaneousness is more complex: in exercising a skill in A-ing, spontaneousness has first to determine, for some way of A-ing *w*, the value of *w* which is a/the suitable way of A-ing on the occasion, say B, and then elicit B. This is a consequence of the variability of skill.

In the next section I argue that spontaneousness essentially involves attention because attention is required to determine the suitable value(s) of *w*: the way(s) of acting suitable for reaching one's objective on the occasion. By contrast, the claim that habitual acts involve invariant relations supports and explains the widespread view that the automaticity of habit dispenses with attention.

V. Attention and goal-dependent sensitivity

Let us start by two points that serve to clarify the conception of attention I rely on.

The first point concerns the objects of attention. Since actions are intentional under some specification(s), the first feature of automaticity, the absence of deliberation and intention, is not a feature of actions as such. Now the notion of automaticity would be incoherent if its defining features did not apply to the same sort of entity. So attention cannot be directed to action-particulars however specified. This seems right: an agent can hardly direct her attention to all the things she does in acting. As it is commonly observed, skill acquisition allows us to direct attention away from lower-level aspects of the action in order to concentrate on more 'strategic' ones. This suggests that the primary objects of attention are the things one is doing (properties, specifications, acts), or what psychologists call "tasks". (Wu [2011; 2013] argues at length that "is automatic" predicates action specifications or properties; Mole [2011: ch.4] and Koralus [2014] provide several arguments for the view that the primary objects of attention are tasks).¹⁵ Attention to elements in one's environment is theoretically derivative from attention to what one is doing. As Mole argues, one directs her attention to a physical object to the extent that this object is *suitably involved* in what one is doing. For example, the tennis ball is suitably involved in the act of hitting a forehand slice (see Mole [2011: 72-73] for how "suitably involved" may be defined). Attention to what one is doing directs the patterns of focal attention as required for the exercise of a bodily skill.

The second point concerns a distinction between phenomenological and explanatory or functional conceptions of attention. The difference between doing something attentively or less

¹⁵ Much research on bodily skill addresses the question of which properties are attended to and which are not, and how skill learning involves changes in the patterns of attention. There is some controversy regarding which exactly aspects of bodily action are attended and which are not (Montero 2010; 2016), but the idea that one cannot direct her attention to *all* the properties of a bodily action is hardly controversial.

attentively is not only phenomenological. It has also certain explanatory consequences. We may understand the claim that an agent does something attentively as the claim that he has certain dispositions to register correctly the relevant features of the situation, which in turn explains why he is disposed to adopt the appropriate course of action. This is why we often explain success or failure by saying that the agent was or wasn't sufficiently attentive. This is a central theme of Ryle's treatment of "heed concepts" (1949: 135-49), to which he appeals in distinguishing habit from skill. And a related conception of attention is usually operative in cognitive psychology, and hence in accounts of automaticity.¹⁶

The conception of attention as underlying the capacity to opt for a suitable course of action can be sharpened by appeal to the semantics of questions. On the Erotetic Theory of Attention proposed by Koralus (2014), a task of A-ing determines a question, that is, a set of possible answers. One engages in A-ing attentively to the extent that one is sensitive to what counts as an adequate answer to this question (notice that "question" and "answers" need not be discursively articulated).¹⁷ On such a conception attention is essentially involved in any exercise of bodily skill, precisely because exercising a skill involves resolving a question: the question of how to A in the circumstances one is confronted with.

We saw that on any occasion where one exercises a skill one is faced with an open "How?" question, a question which determines a set of possible ways of acting on the occasion. Within this set, there is a subset of way(s) that successfully resolve this question, in the sense that they

¹⁶ As Smithies (2011: 250) observes in discussing these two conceptions, "[i]n cognitive science, attention is usually defined in terms of its functional role, rather than its phenomenology" (2011: 250).

¹⁷ It is only certain core aspects of this view that are apposite here. Though he does not consider infinitival questions, Koralus allows that any task may consist in sub-tasks, and hence that attention to the main task might consist in sensitivity to "How to A?" questions. But his claim that questions encode *completion* conditions makes hard to see how he can account for the continuous (processual) mode of guidance exemplified in skilled bodily activity (see below). Wu's (2011; 2015) account of attention is also closely related to the conception sketched above. Wu (2011: 52) claims that attention is metaphysically necessary to bodily guidance, and hence bodily skill, for reasons akin to the ones adduced here. It is a consequence of Wu's view that habit dispenses with attention, a consequence he seems happy to endorse (2015: 5-6).

are the most suitable, or just sufficient, ways of acting on the occasion.¹⁸ This subclass is jointly circumscribed by at least two factors. First, the class of ways that successfully resolve the “How?” question is partly circumscribed by one’s objective: a “How to A?” question is resolved by a different set of ways than a “How to Z?” question on the same occasion. For example, if one wants to return a tennis serve to point *a* in the court, rather than to point *b*, the suitable class of bodily movements will differ. In this sense, this sensitivity is *goal-dependent*.¹⁹ Second, the suitable class of ways is partly circumscribed by occasion-specific properties. Supposing one means to return the serve to *a*, the appropriate class of bodily trajectories is circumscribed by one’s initial bodily posture and position and the spatiotemporal properties of the incoming ball. On the present conception, attention is defined as a sensitivity to what counts as a suitable answer to the “How to A?” question, that is, as what enables the agent to discern what are the most suitable ways of A-ing in the circumstances she is faced with, which will of course require her to direct her focal attention accordingly. Thus in exercising a skill one essentially manifests a *goal-dependent sensitivity to the peculiar circumstances of the occasion*.²⁰ And since skill exercise involves resolving a “How to A?” question, attention is an essential feature of skill.

Resolving the “How?” question is not a one-off affair. Exercises of bodily skills are activities flexibly and continuously adjusted on the basis of feedback. One would typically persist in the face of perturbations, unexpected obstacles, or changes in the configuration of the elements in one’s environment one interacts with, and try to do whatever it takes to achieve one’s objective. This ‘closed-loop’ continuous mode of guidance of bodily activity accounts for a distinctive

¹⁸ This does not imply that there is always a way, or more than one way, to *successfully* resolve the “How?” question. But even so, the “How?” question must still be open for the agent. For no skill exercise is guaranteed to succeed, and there are many ways to fail.

¹⁹ Bargh and Chartrand (2000) notice that the literature on skill development is concerned with goal-dependent automaticity, a claim repeated by Moors & DeHouwer (2007: 17). See also Bargh (1997: 28-29). A similar claim is implied in Logan (1985: 368). Stanley and Krakauer (2013: 5) write “skill can be considered the practice-related improvement in a goal-directed action”. Fridland (2014b: 2741) concurs, but adds that we should add explicit reference to *attention* to this definition. The claim that goal-dependence is a feature of spontaneousness is widely shared.

²⁰ A point of clarification: features of the environment can elicit a habitual act, and can guide the exercise of a skill. These are likewise *circumstances*. I say that the *peculiar circumstances of the occasion* guide the exercise of a skill to mark the fact that, in contrast to habit, the agent need not (and typically has not) encountered these circumstances before. As explain below, this is an important difference between spontaneousness and impulsivity.

feature of bodily skill: a failed attempt to catch a ball or to return a tennis serve cannot fail to be impressed on the agent's experience. Generally, agents can hardly fail to become immediately aware of failures in skilful bodily acts, or at least, when one is not in a position to tell, one would be left wondering whether she has succeeded or not. We shall now see that that there is a striking contrast with habit on this score, and that this points to a crucial difference between habit and skill with respect to the variety of automaticity involved.

A hallmark feature of the automaticity of habit is conspicuous in a familiar kind of mistake: habitual action-slips. A stock example here is the absent-minded driver. When coming back home from work the driver takes the familiar route though he knows full well that it is closed for maintenance, instead taking an alternative route as she had earlier decided (i.e. intends) to do. What is striking here is that the agent lacks any sense that something might be going wrong as she takes the familiar route. She will be irritated or amused when it later occurs to her that she has been led astray by the force of habit, though in some cases one might never become aware of one's slip.

In discussions of habitual action-slips it is generally assumed that what we casually call "absent-mindedness" to make sense of them is a failure of attention. This is how William James (1981: 119) accounts for them, this is how they are accounted for in the influential discussions of Reason (1990: 68), Norman (1981), Norman and Shallice (1986: 12), as well as by the overwhelming majority of psychologists. This is also widely accepted in philosophical literature (Elian & Roessler 2003: 4; Roessler 2003: 388-389; Wu 2011: 63; forthcoming; Romdenth-Romluc 2013: 5). The general presumption is that what sets such slips apart from other kinds of mistakes in action is that all it takes to prevent slips is to pay sufficient attention, or a timely awareness that one is about to act contrary to one's previously formed intention.²¹ Thus habitual action slips are an important source for the widespread view that the automaticity of habit dispenses with attention.²² However, since "attention" can be understood in various ways, what I

²¹ It is important to distinguish action slips, which are due to sheer absent-mindedness, from other superficially similar mistakes that might be due to failures of judgment or weakness of will. Amaya (2013) provides a discussion of the relevant distinctions. As Amaya (2013: 12) explains, the distinguishing mark of action slips is that "[slips] have a quick and easy cure: awareness of an imminent slip is enough to prevent it".

²² As Ouellette and Wood (1998: 55) observe, there is general agreement that the automaticity of habit involves diminishing attention (see also Graybiel 2008: 361; Wood, Labrecque, Lin and Runger 2014: 378-79). Other definitions often figure the feature of (diminishing) awareness, rather than attention (van t'Riet, Sijtsema, Dagevos

need to show is how this widespread view can be construed on the conception of attention as goal-dependent sensitivity sketched above. This will allow me to articulate more precisely the contrast between impulsivity and spontaneousness.

Several psychologists explicitly claim that habit involves a distinctive form of automaticity (Neal, Wood, Labrecque & Lally 2011: 6; Wood, Labrecque, Lin & Runger 2014: 376; Dezfouli, Lingawi & Balleine 2014: 1; Orbell & Verplanken 2010: 374; Wood & Runger 2016: 3; see also Bargh 1997, though Bargh prefers to avoid the term “habit”). However, there are two contrasting views regarding one of its defining features. One view asserts that habits exhibit a kind of *goal-dependent* automaticity, while the other view denies this. I will have to show that habit dispenses with attention on both views.

In their comprehensive review article on the psychology of habit, Wood & Runger (2016: 292) claim that “[t]wo defining features of habit automaticity are (a) activation by recurring context cues and; (b) insensitivity to short-term changes in goals (a.k.a., not goal dependent)”. I have explained the import of feature (a) in section [IV] above. But why is (b) so central? Why not mention the most frequently cited features, such as diminishing attention or awareness?

Researchers insist on insensitivity to current goals because this feature has a clear behavioral manifestation and hence it can be operationalized to measure of habit-strength (Orbell & Verplanken 2003; 2010; Neal, Wood, Wu & Kurlander 2011; Dezfouli, Lingawi and Balleine 2014). As Orbell & Verplanken (2010: 378) explain, since a habit is elicited by contextual cues rather than intentions, its strength can be measured by “the likelihood of an individual making an unintended response to a cue previously associated with that response. The stronger the cue contingent automaticity associated with the response, the more likely an accidental, nonintentional response would be made”. In other words, if intentions are a kind of goal (Moors and De Houwer 2006: 303), the automaticity of habit is measured by the tendency to do things that run counter to one’s goals: action slips. But habitual action slips are due to failures of attention. So on the construal of attention as goal-directed sensitivity outlined above there is not much to choose between the claim that impulsivity exhibits diminishing sensitivity to current goals and the claim that it exhibits diminishing attention. Diminishing sensitivity to changes in

& de Bruijn. 2011: 586; Nilsen, Roback, Brostrom, & Ellstrom 2010: 1). There are important questions regarding the relation between awareness and attention in habitual acts, which I set aside here.

current goals is a behavioral manifestation of diminishing attention phenomenologically construed.

Let me illustrate. The absent-minded driver intends to go back home (A) via the alternative route, which requires her to turn right at junction j (C). So she has an intention to C (or to [A-by C-ing]). At the same time, she is in the habit going home (A) by turning left at j (B), which is the familiar route. On the occasion, impulsivity elicits the response to (B), as she habitually does. However, B-ing is not a suitable way of A-ing on that occasion—the usual route is closed—and the agent knows this. But this is just to say that the agent is not sufficiently sensitive to what counts as a suitable answer to the “How to A?” question on that occasion: she is not paying sufficient attention.

The alternative view on the automaticity of habit appeals to a strand of research that suggests that goals can be activated by a process that satisfies certain features of automaticity, such as absence of awareness, attention and deliberation. This claim is central to the ‘auto-motive’ model of goal-dependent action developed Bargh and colleagues (Bargh, 1990; 1997; Bargh & Gollwitzer, 1994). On the basis of this model, it is reasonable to assume that habits automatically elicit *goals*, and hence that habits rely on a kind of goal-dependent automaticity (Aarts & Dijksterhuis 2000a; 2000b; see also Snow 2006).

This view would account for action slips by supposing that the driver’s intention to A by C-ing is not effectively shielded from the competing automatically acquired proximal goal to B; hence the intention to C (or to [A-ing by C-ing]) is ‘defeated’ by the goal B. So the ensuing act of B-ing is also goal-dependent.²³ But on this view habitual action slips are also due to failures of attention. The absent-minded driver knows that on that occasion, where the familiar route is closed, this automatically acquired proximal goal of B-ing is not a suitable way for her to A (going home, the non-proximal goal). This is why she has formed an intention to A by C-ing. It is just that she fails to take this into account at the critical moment. But again, this is just to say

²³ Here I rely on the discussion of how habit-intrusions activate goals that interfere with intentions in Danner, Aarts, Papies & de Vries (2011). See Neal, Wood, Labrecque & Lally (2012) for a criticism of this second view by proponents of the first view.

that she is not sufficiently sensitive to what counts as a suitable way of A-ing on that occasion: she is not sufficiently attentive.²⁴

Thus impulsivity exhibits insensitivity with respect to the “How?” question, that is, to whether what one is currently doing is conducive to one’s objective on the occasion. This explains how habitual action slips are possible and why agents might not become aware of them until much later. By contrast, we saw that spontaneousness essentially involves this kind of sensitivity. As long as one exercises a skill in A-ing, say in hitting a forehand slice or catching a flying ball, one necessarily exhibits sensitivity to what counts as a suitable way of A-ing. If one does not keep track of the ball and does not seek to adjust one’s posture and hand movement accordingly, one would hardly count as even *trying* to hit a forehand slice. And if one is to count as exercising her skill in catching a flying ball, one has at least to visually locate the ball and try adjusting her hand movements accordingly on the basis of sensory feedback. One might fail to hit or catch the ball, but then one would immediately become aware of this, precisely because skill exercise involves moment-to-moment monitoring of what one is doing: one is continuously alert to the “How?” question.²⁵ This kind of sensitivity is essential to bodily guidance, and bodily guidance is essential to exercising a bodily skill. Spontaneousness essentially exhibits sensitivity to whether what one is currently doing is conducive to one’s goals.

This contrast is closely related to the metaphysical distinction between habit and skill outlined in section [IV]. In exercising a skill an agent resolves an open “How?” question, and attention is defined as sensitivity with respect to what counts as a suitable answer to that

²⁴ So whether impulsivity automatically activates an impulse to B as on the first view, or a goal-state to B as on the second view, what matters is that the agent fails to detect its inconsistency with her intention to A (or to [A by C-ing]). Intentions are formed on the basis of deliberation, and their formation is subject to the normative constraints of consistency with other intentions and knowledge of the agent (Bratman 1987, *inter alia*). But as a goal or an impulse is automatically activated, the agent does not have the opportunity to bring the norms of practical reason to bear. This explains why automatically activated goals or impulses often defy such consistency constraints, and thus conflict with the agent’s intentions. It is this common feature of automatically activated goals and impulses that matters here. So the term *impulsivity* is meant to be neutral with respect to these two views on the automaticity of habit, though on the second view it might be a misnomer, given the different theoretical associations of “impulse” and “goal” in psychology.

²⁵ Amaya (2016) argues that basic acts are immune from action slips. I argue that this is a distinctive property of bodily skills, basic or otherwise.

question. This is why attention is an essential feature of spontaneousness. By contrast, as long as one is engaged in the habitual routine of A-ing by B-ing, the answer to that question is settled: the impulsivity of that routine invariably elicits B in response to A. So as long as impulsivity is operative, there is no open question to be resolved and no role for attention: impulsivity dispenses with attention.

Before closing this section I would like to consider two related objections. I have argued that impulsivity dispenses with attention. But surely, one will object, this is too strong. We are often attentive enough to prevent slips. The driver might well realise before it's too late that B-ing is not conducive to A, and adjust her course of action accordingly. Doesn't this show that we do pay *some* attention when acting impulsively? Moreover, on different evenings the driver might be more or less prone to make a slip as she drives back home. Doesn't that show that one can be attentive to various degrees when impulsively manifesting a habit on various occasions? How does that square with the claim that impulsivity simply *dispenses* with attention? Second, if it is allowed that some degree of attention may be operative in impulsivity, and that different manifestations of the same habit may exhibit varying degrees of attention, then it might seem that the difference between the two forms of automaticity turns out to be merely a difference in degree, rather than a difference in kind.²⁶ For it is clear that different skills also require varying degrees of attention, and that the same skill can be exercised more or less attentively on different occasions.

The first point to be made in reply is that none of my claims implies that in acting habitually we cannot pay varying degrees of attention. But this is not because impulsivity involves (degrees of) attention. It is because a manifestation of a habit can be *more or less automatic/impulsive*. It is widely acknowledged that automaticity comes in degrees, if only because it is gradually acquired and lost. So nothing I have said precludes the claim that, even when a habit is firmly entrenched, impulsivity can be operative in varying degrees on different occasions. Indeed, there is independent support for this claim.²⁷ Once this is acknowledged, it is not hard to accommodate

²⁶ I would like to thank an anonymous reviewer for raising this second worry.

²⁷ Current research suggests that the tendency to make action slips is related to factors other than 'habit-strength' (i.e. the degree to which automaticity is established), which is to say the extent to which automaticity determines the course of action depends on additional factors. These include factors that vary irregularly from one occasion to another, such as acute stress, fatigue, or preoccupation with other matters. So even for well-established habits, the

the above observations. The claim that impulsivity dispenses with attention implies that a purely impulsive act would exhibit complete absence of attention. But this may be quite rare, if possible at all. Agents often retain some degree of awareness or attention, and to the extent that they do they are less likely to make a slip. This is all we need to explain why on some occasions agents may be more likely to make a slip than others, as well as why and how agents can avoid slips.

There is more to be said regarding the second objection. The functional difference between impulsivity and spontaneousness articulated in [IV] is incompatible with the idea that the difference between them is merely one of degree. Impulsivity invariably elicits an act of the same kind in response to similar circumstances (broadly construed). By contrast, spontaneousness determines, for some way of A-ing w , that w is a/the suitable way of A-ing, and elicits w , a way in which the agent need not (and often has not) A-ed before. Attention is sensitivity to what counts as a suitable answer to that question, and hence it is an essential feature of spontaneousness. By contrast, impulsivity dispenses with attention, because as long as the impulsivity of a given habit is operative the answer to the relevant question is settled: there is no work for attention to do in impulsivity. To the extent that one does attend to the relevant question one considers alternative courses of action, which is to say that one considers *refraining* from manifesting the given habit: impulsivity is disrupted.²⁸ This is why attention is antagonistic to the operation of impulsivity: the more sensitive one is to the relevant question, the less impulsively one acts. By contrast, increasing attention is anything but antagonistic to the operation of spontaneousness. To put it more succinctly: attention is external to the operation of impulsivity, but it is an inherent aspect of the operation of spontaneousness—i.e, a feature of this variety of automaticity. So whereas agents can be attentive in various degrees in both skilful and habitual acts, this does not mean that the difference between impulsivity and spontaneousness is a matter of degree.

degree in which impulsivity determines a course of action varies from one occasion to another; attention co-varies accordingly. See Wood, Labrecque, Lin & Runger 2014: 377-380 for discussion of the relevant research.

²⁸ In the example above, where the driver is about to make a slip, sufficient sensitivity with respect to the “How?” question would lead one to *refrain* from manifesting a habit. But impulsivity will cease to operate to the extent that the driver becomes sensitive to the “How to A?” question, even if it turns out that she is on the right track (a false alarm, as it were). Notice, though, that one would still be manifesting that habit, for the habit still explains the act. This suggests that automaticity is not a necessary characteristic of habit.

Moreover, the difference cannot be a matter of degree for reasons which are independent of the claims defended in this paper. First, according to the first view on the automaticity of habit mentioned above, this is not a form of goal-dependent automaticity. But the automaticity of skill is goal-dependent on all counts. The point is not that goal-dependence as such cannot be a matter of degree. Proponents of the first view allow that as a habit is formed, the influence of goals gradually diminishes (Neal, Wood, Labrecque & Lally 2012). The point is rather that the development of impulsivity diminishes goal-dependence while the development of spontaneousness does not. It is this difference which is not a matter of degree. Second, we saw that all views on the automaticity of habit agree on one central feature: an act can be impulsively elicited only if the agent has repeatedly acted in the same way in response to the same circumstances (broadly construed). But the agent need not, and typically has not, encountered before the circumstances involved in the guidance of a bodily skill in all their specificity. Impulsivity requires repeated past coupling of circumstances to acts, whereas spontaneousness does not. So one would have to acknowledge that spontaneousness and impulsivity are different kinds of automaticity even if one rejected the arguments above.

The role of attention in habit and bodily skill is a vast topic. There are several ways to cash out the widespread view that habit dispenses with attention, and a wealth of theories about the role of attention in bodily skill. But habitual and skilful acts come in a bewildering variety, and most accounts are geared towards some variety or other (e.g. the role of perceptual attention in object-directed action). So I have relied on a conception of attention that allows me to draw a contrast between habit and skill that holds across the board, though as a result the distinction is expressed in rather abstract terms. Of course, there are alternative conceptions of attention on the table. But as I mentioned in section [II], my objective here is only to identify certain distinguishing features between two ways of acting without deliberation. In order to do this it is sufficient to show that habit and skill differ with respect to *some* way to understand the features of automaticity. I need not have any substantive disagreement with someone who thinks that attention is primarily a phenomenological notion. Indeed, it is plausible that habit and skill, as well as their different varieties, can also be distinguished on phenomenological grounds. But here I have tried to steer clear of these difficult questions.

VI. At the end of section [3] I formulated a problem that can be used to motivate the Diversity Claim. In exercising a skill, as well as in manifesting a habitual routine, one does something by doing something else: one A-s by B-ing, and the answer to the “How?” question is resolved by the operation of automaticity. What is the difference?

A habit of A-ing by B-ing is a propensity to B *only* in response to circumstances in which one is A-ing, and requires that one has repeatedly B-ed when A-ing in the past. By contrast, a skill in A-ing is a capacity to respond in whatever way is required to A in a range of different circumstances; but the agent need not (and often has not) have encountered these circumstances before in all their specificity, nor does she need to have A-ed in this way in the past. This difference imposes different functional demands on the varieties of automaticity involved. Impulsivity pertains to invariable relations between acts and circumstances (broadly construed): it invariably elicits an act of B-ing in response to A. By contrast, spontaneousness determines, for some way of A-ing *w*, the/a value of *w* which is a/the suitable way to A in the circumstances obtaining on that occasion, say B, and then elicits B as a response. This functional difference has direct implications in turn for the features of automaticity involved.

First, on a construal of attention as goal-dependent sensitivity to the circumstances, spontaneousness essentially involves attention because attention is responsible for selecting a/the suitable way of A-ing. But since impulsivity pertains to invariable relations, there is no corresponding role for attention thus construed. Impulsivity dispenses with attention. Since different varieties of automaticity are distinguished on the basis of their constituent features, this difference with respect to the feature of attention suffices to support the Diversity Claim.

Second, in other work I argue that impulsivity and spontaneousness can be distinguished on the basis of another feature of automaticity: improvement in performance. The idea of skill development implies that of improvement in performance: when developing a skill in A-ing one gradually becomes better at A-ing. This can be naturally accommodated by the claim that skills exhibit variability with respect to what answers the relevant question. The various ways in which the “How to A?” question can be resolved are not equal with respect to what counts as A-ing well. Depending on the way one A’s, say the exact trajectories of a bodily movement, one’s A-ing may be *faster*, *swifter*, more *accurate*, more *elegant*, or generally instantiate to an increasing degree whatever graded property or properties constitute improvement in A-ing. Since the capacity to resolve the “How?” question in an optimal way largely depends on the development

of automaticity by practice and training, improvement in performance is a feature of spontaneousness. By contrast, the claim that one has formed a habit of A-ing by B-ing does not as such imply that one improved in A-ing. When the automaticity of habit becomes operative in one's morning routine, it does not make one faster, more efficient or less likely to make mistakes. Improvement in performance is not a feature of impulsivity. This can be explained by the claim that impulsivity pertains to invariable relations: the answer to the relevant question is always the same—one A's in the same way.

There is much more to be said regarding the differences between habit and bodily skill and the varieties of automaticity they exhibit. But I hope I have said enough to motivate the idea that theoretically important differences drop out of view when habit and skill are bunched together on the grounds that their manifestations do not involve conscious thought or deliberation but largely rely on automatic processes established by habituation.

References

- Aarts, H., & Dijksterhuis, A. (2000a). Habits as knowledge structures: Automaticity in goal-directed behavior. *Journal of Personality and Social Psychology*, 78, 53–63.
- Aarts, H., & Dijksterhuis, A. (2000b). On the automatic activation of goal-directed behavior: The case of travel habit. *Journal of Environmental Psychology*, 20, 75–82.
- Amaya, S. (2013). Slips. *Noûs* 47 (3), 559-576.
- Amaya, S. (2016). Slip-Proof Actions. In Roman Altshuler & Michael J. Sigrist (Eds.), *Time and the Philosophy of Action* (pp. 21-36). Routledge.
- Annas, J. (2011a). *Intelligent Virtue*. Oxford University Press.
- Annas, J. (2011b). Practical Expertise. In John Bengson & Marc A. Moffett (Eds.), *Knowing How: Essays on Knowledge, Mind, and Action*. (pp. 101-12). Oxford University Press.
- Anscombe G.E.M. (1963). *Intention*. Harvard University Press.
- Bargh, J. (1997). The Automaticity of Everyday Life. In Wyer, R. (Ed.) (1997). *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1-61). Lawrence Erlbaum.
- Bargh, J. (1990). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, 93-130). New York: Guilford.

- Bargh, J. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. Wyer & T. Srull (Eds.) *Handbook of Social Cognition* (pp. 1-40). Hillsdale, NJ: Erlbaum.
- Bargh, J., & Chartrand, T. (2000). The mind in the middle: A practical guide to priming and automaticity research. In Reis, Harry T. & Judd, Charles M. (Eds.). *Handbook of research methods in social and personality psychology* (pp. 253-285). New York, NY, US: Cambridge University Press.
- Bargh, J., & Gollwitzer, P. M. (1994). Environmental control of goal-directed action: Automatic and strategic contingencies between situations and behavior. In W. Spaulding (Ed.), *Nebraska symposium on motivation* (Vol. 41, pp. 71–124). Lincoln: University of Nebraska Press.
- Brett, N. (1981). Human Habits. *Canadian Journal of Philosophy* 11, 357-376.
- Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6, 345-375.
- Cooper, R., & Shallice, T. (2006). Hierarchical schemas and goals in the control of sequential behavior. *Psychological Review*, 113, 887–916.
- Danner, U., Aarts, H., Papies, E., & de Vries, NK. (2011). Paving the path for habit change: Cognitive shielding of intentions against habit intrusion. *British Journal of Health Psychology*, 16, 189-200.
- Davidson, D. (1980). *Essays on Actions and Events*. Oxford University Press.
- Dezfouli A., Lingawi NW., & Balleine BW. (2014). Habits as action sequences: hierarchical action control and changes in outcome value. *Philosophical Transactions of the Royal Society* 369 (1655).
- Douskos, C. (2017a). Habit and Intention. *Philosophia* 45 (3), 1129-1148.
- Douskos, C. (2017b). Pollard on Habits of Action. *International Journal of Philosophical Studies* 25 (4), 504-524.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine: The power of intuition and expertise in the era of the computer*. New York: The Free Press.

- Elian, N., & Roessler, J. (2003). Introduction. In Elian, N. & J., Roessler, (Eds.) *Agency and Self-Awareness Issues in Philosophy and Psychology* (pp. 1-47). Oxford: Clarendon Press.
- Enc, B. (2003). *How We Act*. Oxford University Press.
- Fridland, E. (2014). Skill Learning and Conceptual Thought: Making our way through the wilderness. In Bana Bashour Hans Muller (Ed.), *Contemporary Philosophical Naturalism and Its Implications* (pp. 77-100). Routledge,
- Fridland, E. (2014). They've lost control: reflections on skill. *Synthese* 191 (12), 2729-2750.
- Fridland, E. (forthcoming a). Automatically minded. *Synthese*, 1-27.
- Fridland, E. (forthcoming b). Skill and motor control: intelligence all the way down. *Philosophical Studies*, 1-22.
- Graybiel, A. M. (2008). Habits, rituals, and the evaluative brain. *Annual Review of Neuroscience* 31, 358–387.
- Haith, A., & Krakauer, J. (2013a). Theoretical models of motor control and motor learning. In A. Gollhofer, W. Taube, & J. B. Nielsen (Eds.), *Routledge handbook of motor control and motor learning* (pp. 1–28). Routledge.
- Haith, A. M., & Krakauer, J. W. (2013b). Model-Based and Model-Free Mechanisms of Human Motor Learning. *Progress in motor control. Advances in Experimental Medicine and Biology*, 782 (pp. 1–21). New York: Springer.
- Hornsby, J. (2013). Basic Activity. *Aristotelian Society Supplementary* 87: 1-18.
- James, William (1981). *The Principles of Psychology*. Harvard University Press.
- Logan, G. D. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology*, 39, 367-386.
- Montero, B. (2010). Does bodily awareness interfere with highly skilled movement? *Inquiry* 53 (2), 105–122.
- Montero, B. (2016). *Thought in Action: Expertise and the Conscious Mind*. Oxford University Press UK.
- Moors, A., & De Houwer, J. (2007). What is automaticity? An analysis of its component features and their interrelations. In Bargh, J., (Ed.) *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes* (pp. 11-50). New York, Psychology Press.

- Moors, A., & De Houwer, J. (2006). Automaticity: A conceptual and theoretical analysis. *Psychological Bulletin*, 132, 297-326.
- Neal, D. T., Wood, W., Wu, M., & Kurlander, D. (2011). The pull of the past: When do habits persist despite conflict with motives? *Personality and Social Psychology Bulletin* 37, 1428–1437.
- Neal, D., & Wood, W. (2009) Automaticity In Situ and in the Lab: The nature of Habit in Daily Life. In Morsella, E., Bargh, J.A., & Gollwitzer, P.M., *The Oxford Handbook of Human Action*, pp. 428-442. Oxford University Press.
- Neal, D, Wood, W; Labrecque, JS; Lally, P; (2012). How do habits guide behavior? Perceived and actual triggers of habits in daily life. *Journal of Experimental Social Psychology*, 48 (2) pp. 492-498.
- Nilsen, P., Roback, K., Broström, A., & Ellström, P. (2012). Creatures of habit: Accounting for the role of habit in implementation research on clinical behaviour change. *Implementation Science*, 7, 53.
- Norman, D. (1981). Categorization of action slips. *Psychological Review*, 8(1), 1-15.
- Orbell, S., & Verplanken, B. (2010). The automatic component of habit in health behavior: Habit as cue-contingent automaticity. *Health Psychology* 29, 374–383.
- Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin* 124, 54–74.
- Owens, D. (2008). Deliberation and the first person. In Anthony E. Hatzimoysis (ed.), *Self-Knowledge*. Oxford University Press: 261-277.
- Papineau, D. (2015). Choking and the Yips. *Phenomenology and the Cognitive Sciences* 14 (2), 295-308.
- Pollard, B. (2008). *Habits in Action*. Vdm Verlag Dr. Mueller.
- Reason, J. (1990). *Human Error*. New York: Cambridge University Press.
- Roessler, J. (2003). Intentional Action and Self-Awareness. In Johannes Roessler & Naomi Eilan (Eds.), *Agency and Self-Awareness: Issues in Philosophy and Psychology* (pp. 383-405). Oxford: Clarendon Press.

- Romdenh-Romluc, K. (2013). Habit and Attention. In R. Thybo Jensen and D. Moran (Eds.) *The Phenomenology of Embodied Subjectivity* (pp. 5-23). Dordrecht: Springer.
- Ryle, G. (1949/1984). *The Concept of Mind*. University of Chicago Press.
- Setiya, K. (2007). *Reasons without Rationalism*. Princeton University Press.
- Snow, N. (2006). Habitual virtuous actions and automaticity. *Ethical Theory and Moral Practice* 9(5), 545-561.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Stanley, J., & Krakauer, J. (2013). Motor skill depends on knowledge of facts. *Frontiers of Human Neuroscience* 29. doi:[10.3389/fnhum.2013.00503](https://doi.org/10.3389/fnhum.2013.00503).
- Todorov, E., & Jordan, M. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5, 1226–1235.
- Toner, J., Montero, B., Moran, A. (2015). The perils of automaticity. *Review of General Psychology*, 19(4), 431-442.
- van t'Riet, J., Sijtsma, S. J., Dagevos, H., & de Bruijn, G.-J. (2011). The importance of habits in eating behaviour. An overview and recommendations for future research. *Appetite* 57, 585–596.
- Velleman, D. (2007). *Practical Reflection*. CSLI Publications.
- Verplanken, B., & Orbell, S. (2003). Reflections on past behavior: A self-report index of habit strength. *Journal of Applied Social Psychology* 33, 1313–1330.
- Verplanken, B., & Wood, W. (2006). Interventions to break and create consumer habits. *Journal of Public Policy & Marketing*, 25, 90–103.
- Wood, W., & Ruenger, D. (2016). Psychology of habits. *Annual Review of Psychology*, 37, 289-314.
- Wood, W., Labrecque J., Lin, P., and Runger, D. (2014). Habits in Dual-Process Models. In J. Sherman, B. Gawronski, Y. Trope (Eds.), *Dual Process Theories of the Social Mind* (pp. 371-385). Guilford Press.
- Wu, W. (2011). Confronting Many-Many Problems: Attention and Agentive Control. *Nous* 45, 50-76.
- Wu, W. (2015). Experts and Deviants: The Story of Agentive Control. *Philosophy and Phenomenological Research*. DOI: [10.1111/phpr.12170](https://doi.org/10.1111/phpr.12170)