


This article was downloaded by: [University College London]

On: 28 July 2014, At: 03:42

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Cognitive Neuroscience

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/pcns20>

The neural antecedents to voluntary action: A conceptual analysis

Parashkev Nachev^a & Peter Hacker^b

^a Institute of Neurology, University College London, London, UK

^b Department of Philosophy, University of Kent, Canterbury, UK

Published online: 30 Jun 2014.

To cite this article: Parashkev Nachev & Peter Hacker (2014): The neural antecedents to voluntary action: A conceptual analysis, *Cognitive Neuroscience*, DOI: [10.1080/17588928.2014.934215](https://doi.org/10.1080/17588928.2014.934215)

To link to this article: <http://dx.doi.org/10.1080/17588928.2014.934215>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Versions of published Taylor & Francis and Routledge Open articles and Taylor & Francis and Routledge Open Select articles posted to institutional or subject repositories or any other third-party website are without warranty from Taylor & Francis of any kind, either expressed or implied, including, but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement. Any opinions and views expressed in this article are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor & Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

It is essential that you check the license status of any given Open and Open Select article to confirm conditions of access and use.

Discussion Paper

The neural antecedents to voluntary action: A conceptual analysis

Parashkev Nachev¹ and Peter Hacker²

¹Institute of Neurology, University College London, London, UK

²Department of Philosophy, University of Kent, Canterbury, UK

The inferential standards for *testing* hypotheses are settled; those for *constructing* them rarely even discussed. If the fit to the data of a hypothesis matters, then so must its fundamental coherence. That is indeed prior to any other question. Here we make use of conceptual analysis in testing the coherence of hypotheses in cognitive neuroscience and apply it to the study of the antecedents to voluntary action. We show that many influential experiments in the literature are premised—often covertly—on erroneous conceptions that render their hypotheses incoherent. The inferences drawn from the data are therefore invalidated proximally to any objection empirical replication could counter. We further demonstrate the empirical consequences of these errors in generating artifactual observable effects that have no general significance and impede further progress. We conclude with a basic framework for constructing robust hypotheses in this difficult and important field.

Keywords: Voluntary action; Conceptual analysis; Antecedents to action; Readiness potential.

The relation between voluntary action and the brain is an empirical matter, to be determined by scientific experiment. It is properly the business of cognitive neuroscience, which combines the study of thought and behavior with the study of brain anatomy and physiology. If philosophy is to intervene here, we must first explain why it should.

In common with most sciences, cognitive neuroscience does not have a fixed set of tools. To the basic equipment of psychology and physiology have recently been added sophisticated devices borrowed from computer science, information theory, signal analysis, statistical physics, and many other disciplines. This is as it should be—the methodology of a science must evolve with its subject, in response to its demands. Equally, the demands of a subject must not be tailored to the available tools—a special danger in

cognitive neuroscience, where we can study the brain at the macroscopic (e.g., fMRI) and to a lesser extent microscopic (e.g., single cell recording) levels but have no means of seeing the intermediate, “mesoscopic” level at which a great deal of the critical organization may be concealed (e.g., in medial frontal cortex; Nachev, Kennard, & Husain, 2008).

Now when we speak of scientific tools we usually have technical tools in mind—devices that improve the quality of the data experiments generate. But science is not *just* about data. Although exploratory analysis of data is sometimes suggestive, the *first* move in good science is the specification of a set of hypotheses the data are *subsequently* used to help distinguish between. Without hypotheses, data get us nowhere just as without a vehicle fuel gets us nowhere. A scientist has to think about the tools

Correspondence should be addressed to: Parashkev Nachev, Institute of Neurology, UCL, 33 Queen Square, London WC1N 3BG, UK. E-mail: p.nachev@ucl.ac.uk

© 2014 The Author(s). Published by Taylor & Francis.

This is an Open Access article. Non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly attributed, cited, and is not altered, transformed, or built upon in any way, is permitted. The moral rights of the named author(s) have been asserted.

used to *construct* hypotheses as well as the tools used to test them. And whereas the latter are usually technical, the former *can only be conceptual*. Crucially, for a scientific investigation to have power both must be adequate to the task. Moreover, since hypothesis generation precedes hypothesis testing, the conceptual tools are often more important than the technical ones. For a hypothesis has to make sense before the question of testing it experimentally can coherently arise. If a hypothesis involved an unrecognized division by zero in its mathematics, any data collected in an attempt to confirm or infirm it would be beside the point because it would not make any sense.

CONCEPTUAL ANALYSIS IN NEUROSCIENCE

Scientists do not require instruction on recognizing divisions by zero. In most domains of science no special conceptual problems arise because the use of the concepts deployed is straightforward. It is not so in cognitive neuroscience, for reasons that are not easy to convey. Cognitive neuroscience unavoidably traverses the boundaries between the neural cum physiological and the psychological, as well as the boundaries between the neural/physiological and the behavioral. The concepts in these three domains are categorically dissimilar. Their logical articulations are unlike, and the logical connections (of implication, exclusion, compatibility) between the different domains are exceedingly difficult to get into clear view. The relationships are not at all like those between, say, ionic and molecular descriptions of chemical phenomena. The different domains are not reducible one to another, and what explanations are appropriate to one domain may be inappropriate to another. How the description of phenomena in one domain bears on the description of phenomena in another is highly problematic.

The task is all the harder for seeming, at first glance, so easy. Any mature speaker of our language is master of the rich psychological vocabulary of our language, is able to use the verbs, nouns, adverbs, and adjectives that characterize human beings (and, up to a limited point, other animals). But to be master of the use of words does not imply possession of an ability to give an overview of the use. One may use such terms as “voluntarily,” “on purpose,” “deliberately,” “intentionally” perfectly correctly and without any hesitation, but when questioned about their relationships, one may have the greatest difficulty responding. Can one act intentionally without acting

voluntarily?—Yes, when one does something under duress. Can one act voluntarily without acting intentionally?—Yes, when one does something knowingly but neither because one wants to do it nor for any further reason, as when one gestures while one speaks; or when one knowingly does something that is an unwanted or not wanted consequence or by-product of one’s intentional action, as when one wakes one’s wife when putting the cat out at night. Is everything one does intentionally also done deliberately? No, for much of our intentional behavior requires no deliberation or decision, but is done as a matter of course, as when one opens the door when going out or gets on the bus when the bus one awaits arrives. Can one do something with the further intention of doing something else (e.g., go to London to visit Jack), without doing what one does with the purpose of doing that other thing? Yes, for example, when one goes to the theater with the intention of catching the last bus home. For one’s purpose in going to the theater is not to catch the last bus home.

Most of us would have great difficulty in coming up with these answers off the cuff. That shows that one can master the use of an expression without having an overview of that use. It also demonstrates the fact that mastery of use does not imply mastery of comparative use. One may be able to find one’s way round an old town with unerring exactness, but be quite unable to draw a map. But to construct fruitful and coherent hypotheses in cognitive neuroscience, a conceptual map is needed lest the questions be incoherent, the presuppositions misguided, and the hypotheses tacitly unintelligible.

Some will find these arguments unpersuasive—perhaps inevitably so, since blindness to our defects here—what neurologists call anosognosia—is a constituent of the problem. But all scientists agree that we need a standard for determining whether or not a set of data supports a hypothesis, indeed we have a well-developed conceptual apparatus for it—statistical inference—and a commonly agreed criterion: <0.05 probability of error. If this is so, then it is perverse not to demand a standard for determining whether or not a hypothesis *makes sense in the first place*. The application of such a standard requires apparatus for determining not empirical truth but conceptual sense—coherence. Our tool here is conceptual analysis—a process of examining the constituent concepts of a hypothesis and their putative relations, unraveling the nexus of connections so that any incoherences between them can be laid bare. It is kin with the familiar process of looking for confounds—alternative explanatory factors

that void a hypothesis, so framed, of the power to be tested experimentally, thereby rendering the data irrelevant to the question they are collected to answer. What we are proposing is therefore something good scientists do already—we merely set it within a comprehensive framework that helps the novice learn how to do it well. Though conceptual in method, such an analysis has consequences for empirical research. It is no less pertinent to the formation and validation of empirical hypotheses than the mathematical theorems invoked in formulating those hypotheses.

THE CONCEPTUAL LANDSCAPE OF VOLUNTARY ACTION

To make an assertion about the neural processes that make voluntary action possible one must be clear about what voluntary action *is*. Without such clarity one's experiment may be confined to a conception of action that has no application outside the experimental setting and is therefore of no broad consequence, or worse, one may be shipwrecked on the rocks of a misconception of what *counts* as an action (for example, confusing a reflex reaction with an action, or conflating bare doing with acting). Unless what we discover ramifies into the legal, moral, wider social reality, it will be of no interest to anyone. The first step here, therefore, has to be a survey of voluntary action, spanning the full landscape so that the boundaries of each aspect can be brought into view (see Hacker, 2007, for a full discussion).

We must draw some preliminary distinctions between doing something and acting or taking action, between act and activity, between acting and refraining from acting. All agents, inanimate as well as animate, do things. All acting is doing, but not all doing is acting. Inanimate and non-sentient agents do things, but do not act. A rolling stone is doing something, namely rolling downhill, but it is not acting, and a plant may be growing, but not acting. On the other hand, such agents with causal powers may do things *to* other agents, i.e., act *on* them and so bring about or prevent change in the patient upon which they act. Nevertheless, insentient beings cannot *take action*. They may have an action, as many machines and organs of an animal do, but they cannot act or perform deeds—for only beings that can take action can also act. To take action is to act voluntarily in response to a circumstance (e.g., a threat) or in pursuit of a goal, given apprehension of an opportunity. So a creature that can take action is one that can do or refrain from doing something

voluntarily, that can have or seize an opportunity, that can opt for or choose one course of action over another.

Sentient creatures may *do* things that fall short of acting (e.g., fall asleep). Among things they do are *reactions* that cannot be initiated at will. Some of these may be controllable and inhibitible (e.g., stifling a sneeze, cough, or yawn). While inanimate doings are mere motion or change, when self-moving intelligent beings do such things they *act* (walk or run, climb or jump) or *take action*. This may, but need not, involve acting *on* another thing. Acting on another thing is to bring about or prevent change in or to the other thing by one's action. An animate being may move as an inanimate being moves (e.g., slip, fall), and its limbs may move without the animal moving them (e.g., reflex actions, being subjected to the force of another thing). But characteristic animal movements are forms of behavior of which *it makes sense* to say that it is voluntary, done on purpose in pursuit of a goal. Voluntary behavior is behavior that is controlled by the agent, that exemplifies a two-way power to do or refrain from doing. Voluntary behavior stands in contrast to both involuntary behavior (when what is done on some occasion is not under one's control) and non-voluntary behavior (e.g., when what one does is done under duress).

Human beings, like other sentient animals with wants, have the power to move or refrain from moving, to act or refrain from acting, at will. "To move" here does not mean *causing a movement*, but *making* one. When a human being's movement is an act, it is of a kind that falls within the ambit of the variety of teleological explanations appropriate to human action. A woman may have moved her hand in order to ..., or because she wanted to ..., or because she thought that ..., for such-and-such a reason, or out of such-and-such a motive, and so on. The movement is to be understood as liable to the range of explanations of the exercise of two-way powers by a *rational being*—a being sensitive to reasons for acting and responding. Such behavior is typically under the control of the agent. She is answerable for it. It is something she can do or refrain from doing at will, and for which she may have a reason for doing that explains and justifies or purports to justify what she does.

We distinguish act from activity, the active from the passive, and act from omission. An activity is a sequence of acts. It may be a repetition of acts (e.g., walking, running, hammering) or an ordered sequence of acts (e.g., starting up the car and driving off) or a sequence of acts given their unity by their purpose within a framework of rules

(playing tennis or chess) or without any such framework. There is not one divide between action and passion, but a number of different ones. We contrast what we do (when we are active) with what happens to us (when we are acted on). Among things that happen to us, we distinguish between what happens and what is done to us. Among things done to us, we distinguish between things we demand or allow to be done to us (so we are not merely passive), and things done to us against our will. We further contrast thought and action, the thinker with the doer, and action with inaction—as when we look on and do nothing, take no action, refrain from acting.

Inaction too may take various forms. It may result from lack of ability or lack of opportunity. Opportunity is relative to ability—what is an opportunity for the skillful may be none for the incompetent. A necessary condition for not acting to constitute omission is that there be both ability and opportunity to act. Depending on agential knowledge, the context, and the requirements of the circumstances, one's not doing something may be omitting, abstaining, or refraining.

Characterizing action, then, in a manner that allows the neural processes that make it possible to be illuminated is far from straightforward. Though we have access to the immediate, circumstantial features of the environment, the principal difficulty is in knowing how to make sense of features of the agent that are not so easily described for they are not simply states but *abilities*. Here it is tempting to lean on conceptual devices that acquire their sense in a different domain and offer no solid support in this one. The only safe means of proceeding is to pause at each step, testing the ground by considering all the ramifications of each move, before moving forward. And if a particular direction leads to a dead end, the only option may be to retrace our steps to the very beginning. Here we examine the antecedents to action—the aspects of action preceding it that have a bearing on its neural relations.

THE ANTECEDENTS OF ACTION

When one raises one's arm the contraction of the muscles is naturally preceded by activity in the motor neurons innervating the muscle. Since all neural signaling takes a non-zero length of time, such activity must occur in advance of the movement. It is therefore tempting to pursue the chain of events further up the neural tree in the belief that at its apex must be the one neural event, or at least a limited set of neural events,

that is the ultimate cause of the movement, or at least explanatorily more potent than any event contemporaneous with the movement itself. It seems any explanation of voluntary action must explain its antecedents first and foremost.

The neural antecedents of action are important for a second reason. It is constitutive of the concept of voluntariness that what is done voluntarily lies within the range of actions that the agent can do or refrain from doing at will, and actualizes a possibility for action the agent is aware of as being available. Where a movement is invariably associated with a somatic event it might instead be determined by the cause of the event. For example, the knee jerk that follows a tap on the patellar tendon may either be a voluntary movement, coincident with the tap, or an involuntary movement caused by activation of the stretch reflex. We can distinguish between these possibilities to the extent to which we can dissociate the somatic event from what the subject voluntarily does. The movement must be involuntary if the subject is unable to refrain from doing it when the event occurs, and we have no grounds for doubting its voluntariness if they can do it without the event. If we *cannot* demonstrate a dissociation—the event always matches—then the question remains, strictly speaking, unanswered, but it makes better sense to treat the action as voluntary as long as the subject retains the two-way control constitutive of voluntary action. Now a difficulty arises where such an indissociable somatic event is significantly antecedent to the action. Though our two-way control remains, there now seems to be doubt about the identity of the controller, for unless we invoke backward causation, something else seems to be determining the choice for us, in advance, making the action *covertly automatic* (e.g., Walter, 2001; Wegner, 2002). The presence of such antecedent “brain somatic” markers may therefore not only explain actions but also make us uneasy about their fundamental nature.

We need to examine each of these aspects in turn.

Antecedence and causation

The notion of causation circumscribes a diversity of concepts mapped out in detail elsewhere (Hacker, 2007). For our purposes we need only examine one prototype of causation customarily deployed in explanations of biological phenomena. A biologist expects a causal explanation to specify a set of biological components and to describe how one component acts on another to bring about a change

in it. For example, a muscle fiber contracts when activated by the motor neuron that innervates it. Sometimes, as in this example, the relation between the components is obvious, but often it is opaque. In such circumstances, we must look for criteria that help us distinguish as best we can between rival hypothetical possibilities. Amongst the criteria commonly used in biology, two are pre-eminent—temporal coincidence or priority, and dependence. If the action of component *A* on component *B* occurs *after* the change in *B*, then that action cannot be its cause. If the presence or absence of the action of component *A* (naturally or experimentally induced) makes no difference to the occurrence of a change in component *B*, then that action is unlikely to be its cause, and certainly less likely than that of a component to which a change in *B* is shown to be more strongly dependent.¹

Now in creating a biological picture of causation here, we cannot use these two criteria in the same way. Temporal priority is merely a binary factor—the action of a component *at any time* preceding the change in *B* may be its cause, giving us no grounds for favoring or neglecting a component that acts earlier or later. By contrast, dependence can be indexed by continuous fractional measures (e.g., sensitivity and specificity), giving us a direct comparative index of the component's causal power, as far as our observations allow us to grasp it. If the change in *B* is shown to be more strongly dependent on the action of component *A*₁ than on the action of component *A*₂, the causal power of *A*₁ is greater whatever the temporal relation between *A*₁ and *A*₂ (as long as both precede or coincide with the change in *B*). Where the dependence is the same, the temporal factor cannot be used to rank two components, for either earlier or later action could be interpreted as being more important.

Where, then, does equating degree of antecedence with causal power come from? It is a feature of many biological mechanisms, especially at the molecular and cellular level, to follow a *serial* pattern of organization—one component acts on another at some given probability of success (e.g., an enzyme interacting with its substrate or a receptor with its ligand), a series of such interactions forming a chain of causation where each step is dependent on the preceding (Figure 1). Within such a serial chain, the probability of each link is *conditional* on the preceding, and since each link

takes up a finite length of time, temporal priority and dependence of the final outcome become *circumstantially* correlated. The best criterion is *still* dependence, it is just that in such circumstances temporal priority is a good enough proxy.

Now the brain is certainly not a simple serial system, at least not at the neural level. In the context of voluntary action, it is well established that *multiple parallel* pathways are simultaneously co-activated antecedent to action (e.g., Cisek & Kalaska, 2005). Temporal priority here *cannot* be assumed to coincide with causal power: We must always closely examine the nature of the association between the activity of a candidate component—whatever its timing—and the voluntary action. We must do this not only for neural antecedents but *also* for any behavioral or cognitive antecedents others have attempted to bring into causal play, either as causal elements in themselves, or as reliable *temporal markers* of accompanying neural processes that are supposedly causal.

With this in mind, let us consider a few antecedent phenomena implicated in causal accounts of voluntary action. We focus on only a few because the critical points are obviously extendable to others—to be exhaustive here is merely to be exhausting. We begin with cognitive antecedents.

URGES

It is held that urges are antecedents to action that may illuminate the neural processes underlying action in general (Haggard & Libet, 2001; Haggard, 2008). At first sight, urges appear to satisfy the requirements for antecedents that may be licitly implicated in causal accounts of voluntary action. First, urges are definitionally antecedent to action. One can have an urge to *V* without *V*-ing, as when we “suppress” the urge, but if we do not stifle the urge, then it antecedes the *V*-ing. If an urge to *V* follows *V*-ing, then it can only be an urge to *V* again. Where an urge accompanies *V*-ing—which is only possible where *V*-ing takes an appreciable length of time—it does so only to the extent to which the *V*-ing is incomplete, for example, when one's bladder is not fully empty, or an itch remains. This case thus effectively reduces to the circumstances where the urge precedes the action. Secondly, an urge is an occurrent phenomenon with genuine duration—we can report when it began, how long it lasted, and when it ended. Thirdly, an urge is always transitive—it is an urge to *V*, where *V* is a specific action, or at

¹The picture is more complex than this, but our simplifying here for the sake of brevity does not change the conclusions. For further details, especially the notion of causal conditions, see Mackie, 1974.

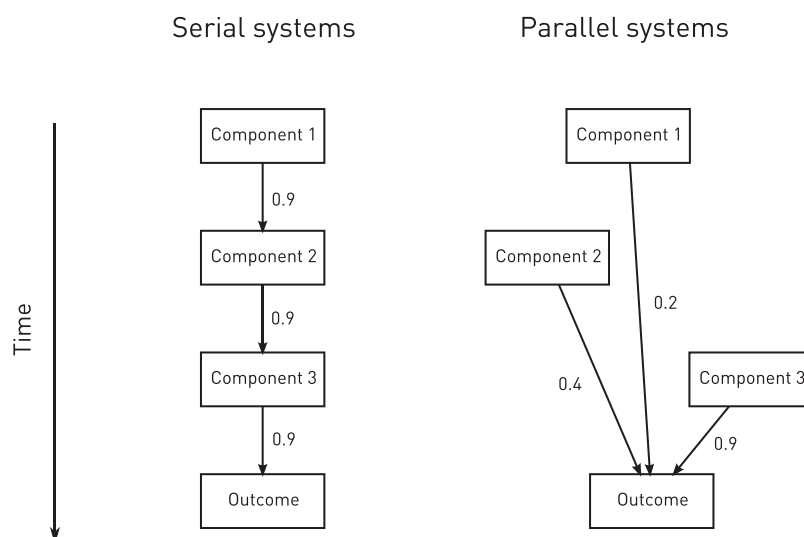


Figure 1. Temporal priority and causality.

Notes: Many biological mechanisms, especially at the molecular and cellular level, follow a serial pattern—one component acts on another with some given probability of success (e.g., an enzyme interacting with its substrate or a receptor with its ligand), a series of such interactions forming a chain of causation where each step is comparably dependent on the preceding (left diagram). Within such systems causality and temporal priority coincide—it therefore makes sense to speak of the first event as being causally the most important. This is true *only* of serial systems. Where a system is parallel, such as the brain, time no longer coincides with causal power (right diagram). Here, what matters is the relative dependence of the individual components (arrow line fractions in diagram), not their temporal order, and temporal order cannot be used as a proxy measure of causal power.

In the context of causal accounts of voluntary action, if an agent can change his action with perfect control 200 ms before it is done, then that a weak bias may be demonstrated 1000 ms earlier does not significantly change its voluntariness.

least—for example, in the case of fidgeting—a circumscribed set of possible actions.

On closer examination serious difficulties arise. First, one normally acts without feeling any urge to act. Indeed, urges are both relatively rare and specific to circumstances that obviously differ from common instances of voluntary action (Jackson, Parkinson, Kim, Schürmann, & Eickhoff, 2011). Moreover, absence of a felt urge to act is clearly not a criterion for the involuntariness or non-voluntariness of what was done. Urges are most commonly felt needs to do something (e.g., sneeze or yawn) that we cannot initiate at will, but can sometimes suppress, either to the point of disappearance or for a while, or felt needs to do something that is under partial control but that we can sometimes initiate at will—such as scratching an itch or voiding bladder or bowels. Where we rarely refer to urges as felt needs that are less somatic and more “intentional” in character—e.g., to kiss someone one loves, to hug a child, to hit someone who is being intolerably offensive—we use the term *analogically* from the less voluntary use as a way of *illustrating* the strength of our desire. Language is full of such analogical uses—one does not need a real weapon to kill a conversation, a physical key to unlock a

problem, a source of water to rain on someone’s parade, etc.—that mislead if taken literally. We know the use here is analogical because we can describe such circumstances perfectly well without recourse to the notion of an urge.

Second, urges are characteristic of circumstances where actions are normally withheld, at least for a time. Indeed, it is striking that urges are commonly spoken of in relation to “automatic” actions we do not positively choose to execute but only to delay, such as a sneeze or a yawn. Where actions are sometimes automatic and at other times voluntary—coughing, blinking, urinating—urges tend to arise in the former case and not the latter. If an action is executed *at the same time* as the occurrence of a felt need (as in the case of an infant voiding its bladder or sneezing) we would not speak of an urge—because the action is already being carried out. Inhibition of the movement—either voluntarily or non-voluntarily through physical restraint—therefore must be constitutive of an urge, for an urge implies both a felt need for an action *and* the prevention of its satisfaction. Far from being causal antecedents to action generally, if urges have any physiological connection with the neural substrate that makes actions possible, it is with circumstances of suppressing or

being prevented from doing something, rather than with ordinary voluntary behavior (Nachev, 2011; Rothwell & Edwards, 2011).

Empirically, the most commonly cited observation, due first to Itzhak Fried and his colleagues (Fried et al., 1991), is the report of an urge to move without any overt movement when focal areas of the brain are electrically stimulated macroscopically, i.e., with large electrodes that will typically influence the activity of millions of neurons simultaneously. The areas from which such responses can be evoked are generally close to “motor” areas where stimulation more commonly (or at higher intensity) produces overt movements (Fried et al., 1991; Lim et al., 1994). But these are highly artificial—and definitionally abnormal—circumstances, characterized by many confounding factors that offer as good or better an explanation for the results as a general link between urges and actions. It is in the fundamental nature of macroscopic electrical stimulation of the brain that it cannot be strictly localized to the area being stimulated, for the current will have remote effects both physically, by potential spread at the site of stimulation, and neurally, via axonal connections. It is also in the nature of such stimulation that its effects on the tissues may be stimulant, depressant, or any combination thereof, and may relate to the intensity of stimulation in complex (e.g., non-monotonic) ways. Making sense of such results is therefore difficult.

In any event, it cannot be assumed, as some have implied (Haggard & Libet, 2001), that just as electrical stimulation of neural tissue may be physiologically subthreshold, insufficient to evoke a neural response, so stimulation that only evokes an urge is psychologically subthreshold, insufficient for an overt movement. Indeed, several aspects of the phenomenon go against the received interpretation. First, urges were rare in Fried’s study, being observed at only 15 of the 129 eloquent sites. If an urge is subthreshold to an overt response, then it ought to be more easily elicited than an overt response—in clinical studies the standard practice is to start with low intensity stimulation at each site and increase until either a predefined limit is reached or a response observed. It also ought to be commoner than an overt response, for the volume of brain tissue above a threshold approached incrementally from below is likely to be smaller than the volume above a substantially lower threshold—this is a consequence of the commonly observed current density spread function (Tehovnik, 1996). Second, where an overt movement was observed following higher intensity stimulation, it was not always congruent with the urge, sometimes even involving a different limb. If the urges elicited here are closely related to the action

executed, then they must be specific to it. Third, overt sensory responses of various kinds were more than twice as common as urges, inevitably downgrading our confidence in the specificity of the patient’s report here. Given the natural ecology of urges we have described and its dominance by circumstances of movement *suppression*, urges in this context are arguably better explained by co-activation of the neural substrates of *incompatible* movements than they are by subthreshold stimulation of what would, at stronger current, evoke an overt movement. We do not need to supply a robust alternative account of the effects of stimulation here—the critical point we wish to make is that the data clearly admit of too many equally or more plausible alternative explanations to draw strong empirical support for any general role of urges in voluntary action.

Many of these concerns extend to a recent stimulation study of the parietal cortex in which a reported desire to move (in the absence of any overt movement) is described (Desmurget et al., 2009). Moreover, the picture here is emphatically *against* a causal link between intention and action, for higher intensity stimulation at sites evocative of such reports at low intensity produced not overt movement, but the mistaken report that a movement *had already occurred*. The latter finding is readily intelligible as the proprioceptive or kinesthetic analog of a somatic sensory illusion easily elicited from neighboring parietal sensory regions, but what about the former? All we can say is that it is very difficult to know exactly what the patients were describing, for their vocabulary is here being unnaturally applied to circumstances entirely outside the experience from which it has acquired its sense. Reasoning parsimoniously, if one felt a weak proprioceptive sensation of movement, one might well report it as a desire to complete the movement. Whatever the correct interpretation, a causal link between intention and movement is not amongst the possible inferences here.

Moreover, across the spectrum of pathological conditions, there is a clear dissociation between abnormalities of urges and action that makes a causal link between them implausible (Jackson et al., 2011). There are disorders of movement, such as chorea, where an obvious excess of voluntary movement is not associated with any reported urges to move—the patients simply move. Conversely, there are disorders associated with an excess of urges to move, such as Gilles de la Tourette syndrome (Shapiro, Shapiro, Gerald, & Feinberg, 1988), where the movements themselves may be much less frequent than the urges. Critically, in

pathological conditions where urges are a major part of the clinical picture, the movements are always perceived to be unwanted by the patient: Inhibition of movement is invariably present. To the extent to which urges are a distinctive marker of anything here they are of the desire to inhibit movement.

INTENTIONS

At first sight, intentions appear to remedy the defects of urges as causal antecedents while retaining their virtues. If a subject avows an intention—though she need not for the action to be intentional—then the intention avowed will antecede or coincide with the action. Intentions are identified by the intended action, and so are specific to the action. But unlike urges, intentions are not ecologically limited to circumstances of response suppression—there is none of the element of compulsion and its resistance here—and so generalization across the spectrum of voluntary action seems possible.

Unfortunately, a closer look reveals immediate difficulties. To intend to do something is not merely to be able to say what you are going to do in advance of doing it, even if declarations of intentions are usually good grounds for predictions. “I am going to vomit,” “I am going to shiver,” “I am going to have a seizure,” may be perfectly accurate predictions of future movements (amongst other things) but they cannot be declarations of intentions. It is not the accuracy of the prediction that distinguishes intentions but the retained ability to do *otherwise*, coupled with either wanting to do the intended act or with reasons for doing it. Indeed, unlike a prediction, an expression of intention is *not* falsified by failure to carry out the intended action. That someone intends to V is the ground for a prediction that she will V *conditional* on other things in the world being as she expects them to be. Saying that I intend to drink wine from a glass does not commit me to the action if in the interim I see someone put poison in it, if I notice that the edge is cracked, if I spot a better vintage nearby, and so on. To attempt to connect intentions with actions within a causal biological model therefore necessitates bringing in *all* the conditionals on which fulfilling the intention implicitly depends. In any given circumstance, one may be able to cite a set of conditionals, but it is hard to see how one could ever specify them comprehensively, for the horizon of possible conditionals here is uncircumscribable in principle. One could try to limit the possibilities *within an experimental task*, but all that does is to make the

task inadequate to the problem it is supposed to illuminate, for the lack of constraint is precisely why we speak of voluntary action as *free*. Furthermore, the subject of an experiment is always free to stop playing.

In any event, most voluntary actions and most intentional actions are not preceded by any kind of antecedent intention—one merely acts—and one’s act may be voluntary without being intentional, or it may be both voluntary and intentional or, indeed, non-voluntary and intentional, as when one acts under duress. To act intentionally does not require the formation of an antecedent intention, any more than it requires antecedent deliberation and decision.

Even if we insist on pursuing intentions, we cannot easily pin them down, so as to correlate them with any kind of neural process. Intentions are not felt. They are not kinds of sensation (neither local, like pains, nor global, like weariness). They have no somatic location, they do not have degrees of intensity, and they have no phenomenological properties (they neither throb nor sting, and are neither sharp nor dull). There is no such thing as feeling an intention. (As Schopenhauer observed, the Will is not an Idea—intentions are formed not felt. Hence they have no “onset.” They do not set in, like rain, nor do they come over one, like nausea.) Unlike occurrent feelings, emotions, and sensations, intentions have no “genuine duration”—that is, they cannot be interrupted and later resumed, and they persist through periods of sleep. Intention-formation in advance of action may be preceded by deliberation and decision. One’s decision to V (whether or not it was based on deliberation) occurs at a time, but not at a split second of time (just as lunch may be at one o’clock, but not at one o’clock, 10 seconds and 25 milliseconds). But intending need not be preceded or accompanied by deliberation and decision—the manifold intentional acts we perform in the course of the day are not uniformly preceded by deliberation and decision (the speech-acts of a normal conversation are not). One may intend to V and one may V intentionally without making any choice (the possibility of a choice is not a choice)—as when one picks up a knife in order to butter a piece of toast, or opens a door with the intention of going to work. One may form an intention in advance of acting, but nevertheless not act on the intention for any number of reasons (one may change one’s mind, forget, there may not be an opportunity, someone else may give one what one intended to obtain, etc.).

These objections are commonly countered by evidence of consistency in the reported timing, for example, when subjects are asked to match the alleged

time of the “occurrence” or “onset” of their freely chosen intention to press a button to the hands of a moving clock in front of them (Libet, Gleason, Wright, & Pearl, 1983). But a meaningless question can have a consistent answer. Moreover, identification of the putative time of the “occurrence of intending” in such experiments is consistent *only* where circumstantial features of the task are kept the same (Banks & Isham, 2010). Where these features are changed the result changes. For example, one can illusorily delay the perceived time of the button press in this task by pairing the action with a sound played shortly thereafter. Since the time of the actual movement remains unchanged, if intending is causing the movement, the time of the “occurrence of intending” should remain unchanged too—in fact, it is also delayed, in synchrony with the illusorily delayed time of the button press. Interference with the connection between the reported intention and the movement cannot be the cause of this effect because the sound occurs *after* the button press (Banks & Isham, 2009). The only possible conclusion is that the reported time of the “occurrence of intending” does not accurately reflect any contemporaneous process. Furthermore, if one uses a digital rather than an analog clock the time of the “occurrence of intending” is substantially delayed, occurring much closer to the time of the movement (Banks & Isham, 2010). The difference between digital and analog clocks is likely explained by the well-known “flash lag” effect—the illusion that a gradual change in an object occurs earlier than a sudden one (Hazelhoff & Wiersma, 1924). Whatever the explanation, these simple variations show that the consistency argument, for whatever it is worth, cannot stand.

In short, intention cannot play a major role in causal explanations of action because its relation to an action cannot be causal without the action’s ceasing to be intentional. Intention does not bring about an actuality any more than a plan does. Its temporal parameters are ill-defined, not because we lack the right behavioral paradigm or empirical tools to nail them down, but because they are constitutively so.

Let us now consider a few candidate neural antecedents.

LATE ANTECEDENT ACTIVITY

It has been known for half a century that voluntary movements may be preceded by several hundred milliseconds of a characteristic modulation in neural activity (Kornhuber & Deecke, 1965). Originally observed in the electroencephalographic signal in the anterior midline, the *Bereitschaftspotential*, as it is

commonly called, is robust enough to be used routinely to distinguish voluntary from automatic movements in neurological practice where the patient’s authority on the question is for whatever reason in doubt. A summary of the aggregate activity of millions of neurons is naturally polymorphous, and so there are circumstantially differing variants of the signal that likely reflect different views of the same underlying activity (e.g., Haggard & Eimer, 1999). It is not necessary for our purposes to examine each variant, for the key features on which the critical conceptual analysis depends are shared—monotonicity, usually in the form of a slow change most pronounced at or just before the point of moving, and only relative specificity for the nature of the movement that follows.

Let us first consider the question of specificity. Imagine that *in the context of a behavioral task* a *Bereitschaftspotential* reliably predicts a button press 400 ms before it is made. Such circumstances are widely held to imply a direct relation between the neural activity and the movement. But if *the same signal* were to be followed by a button release, a finger movement at 90 degrees to the button, or any other hand action materially different from the one *specified by the task*, we could hardly speak of any kind of specificity here. Indeed, no-one has been able to demonstrate specificity for the *actual* movement made so far in advance, only for a very wide class of movements crudely defined by features such as the laterality of limb effector (Eimer, 1998). The specificity here is therefore far too poor for us to have any confidence in the nature of any underlying relation.

To this it is commonly objected that we *know* what the subject is doing (because she is following our instructions)—the crucial point is that we can predict the *timing* of the action, the *when* already knowing the *what*. But that conclusion rests on a fundamental mistake about the criteria for individuating an action, for the following reasons.

We naturally distinguish between what might be called act-categories and act-individuals (von Wright, 1963)—firing a rifle is an act-category, firing a rifle to kill a specific man on a specific occasion is an act-individual. That two actions fall in the same act-category does not mean they are the same act-individuals. Pulling the trigger when your comrade is in your sights is not the same action as when it is your enemy, though the bare movement may be the same. Indeed, two or more act-individuals might be distinguished *only* by their timing. When one plays music on the piano, the act of hitting the same note, in the same way, at one point in the piece is different

from another, and is different on different occasions of playing the piece at the very least because one has played the piece before. Were it not so, one could never be able to generate a reproducibly timed sequence of identical movements because timing would not differentiate one movement from another.

Now consider a subject performing an experimental task where a movement of some kind is required. The action will fall within an act-category (e.g., button press), but will also necessarily be an act-individual (e.g., a left button press at time t) chosen within a spectrum of possible acts (e.g., left vs. right button press at time t_1 – t_n). The act-individuation will consist *both* in *what* is done and *when* it is done, not just the former. The specificity of any neural correlate of the action therefore must be shown in respect of *both* movement and temporal features. We may be able to show this for the former—for instance, by finding left versus right primary motor activation reflecting a left or right finger movement—but how can we do it for the latter? The timing of a movement cannot be specified dichotomously, and even where the continuity of possible times is constrained by a timing signal of some kind there is still a natural variability of response that is of the same order of magnitude as the temporal antecedence of the *Bereitschaftspotential*. The neural activity differentiating between possible acts differing only in time may itself differ only in its timing, precisely what we are seeking to establish in the experiment. We are therefore caught in a vicious circle with no escape—to estimate specificity we need to know the timing of the correlated activity and to estimate the timing we need to know the specificity.

This theoretical problem has a robust empirical manifestation in the well-known phenomenon of *affordance* (Gibson, 1977; see Sumner & Husain, 2007 for a recent review). It is established that in any set of circumstances the spectrum of possible acts—not just the action performed—may be discernible both behaviorally and neurally (Cisek & Kalaska, 2005; Grèzes & Decety, 2002; Tipper, Paul, & Hayes, 2006). For example, a subject faced with a drinking cup will show activation of the cortical “motor” areas contralateral to the arm with which the position of the handle makes it likely she would reach even if she has no intention to pick it up and never picks it up (because it is someone else’s cup). A great deal of the activity in the brain is related to the recognized occasion of the *possibility* of actions. This ought to be unsurprising—there can be no homunculus in the brain deciding what the agent actually does, and so activity characteristic of a spectrum of action possibilities beyond the action actually made may be observed at any time (Cisek, 2007).

Now *task* affordance—the variety that has been studied—can be easy to detect—in the cup experiment, for example, by the hemispheric laterality of the behaviorally silent neural activity. But *temporal* affordance, where possible acts differ only in time, may not be detected in this way because the associated neural activity may also differ only in time (Figure 2). Furthermore, just as task affordance is greater the less constrained the choice between possible actions, so temporal affordance may be expected to be greater the less sharply defined the range of possible timings of the invited movement—precisely the circumstances sought in experimental demonstrations of the *Bereitschaftspotential*. Indeed, a monotonic modulation of activity in relation to the time of the executed movement is *exactly* what one would expect to see here, for the degree of co-activation of the neural substrates of possible acts would be expected to be proportional to the temporal proximity of their movements to the one actually made. In any event, there is no doubt that temporal affordance will inevitably confound *any* attempt to extract temporal information about the relation between neural activity and the action putatively associated with it. This ineliminable confound makes it impossible to build any strong causal models of voluntary action with such data. Indeed, others have given plausible alternative accounts of the *Bereitschaftspotential* invoking only spontaneous variability in neural activity as its critical foundation (Schurger, Sitt, & Dehaene, 2012).

EARLY ANTECEDENT ACTIVITY

The *Bereitschaftspotential* and its kin typically occur less than one second before the movement. One research group has observed predictive antecedent activity substantially earlier than this, in the context of a free-choice task where the planned action is made differentiable by its multivariate pattern of BOLD activity in the brain (Soon, Brass, Heinze, & Haynes, 2008). Here the evidence seems at first sight stronger than in neurophysiological studies, for the multivariate pattern of focal BOLD activity is potentially much more specific to the task, though, of course, the fatal temporal affordance confound remains. Nonetheless, it is important to consider some crucial features of the approach employed in this study, as the criticisms generalize to other studies that seek to characterize early antecedent activity in this way.

To examine the specificity of antecedent neural activity for an action inevitably requires an option of

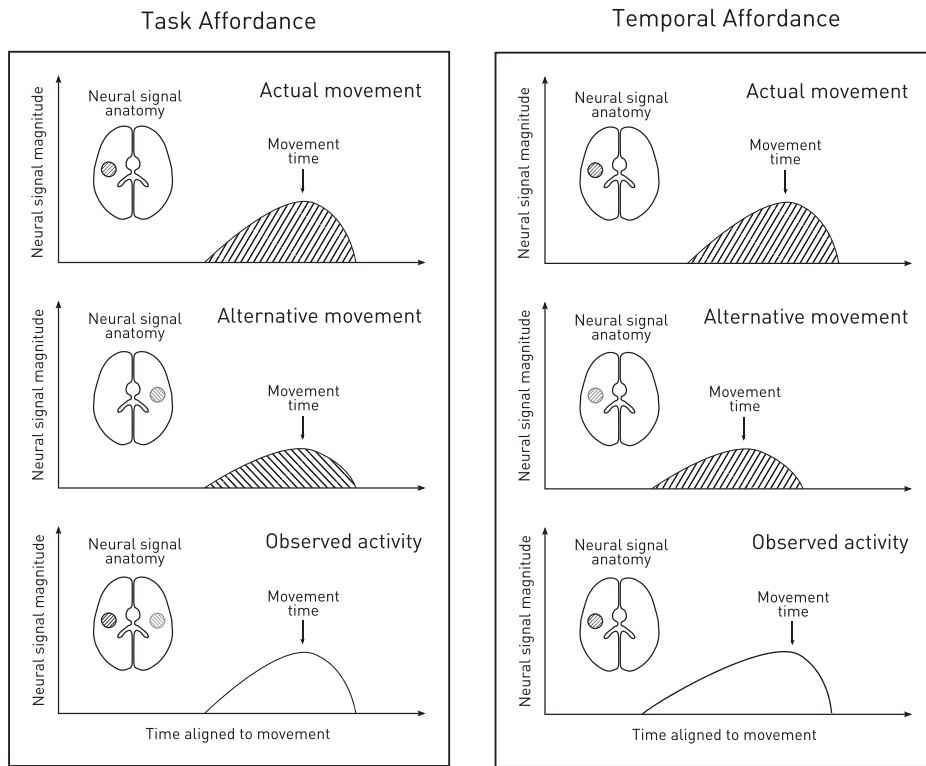


Figure 2. The confound of temporal affordance.

Notes: It is crucial in behavioral experiments to consider not only the movements that are actually performed but also those that *could* have been performed. The neural substrates that make movements possible must also make *alternative* movements possible, and if the former is reflected in observable neural activity so may the latter. Indeed, the affordance literature offers multiple examples of this, both in behavior and in imaging. This literature, however, has examined only one kind of affordance—task affordance—where alternatives of *what* is done are explored. This is relatively easy to study because alternatives for different movements have different behavioral and imaging correlates. For example, hand laterality of response is dissociable by primary motor cortical laterality.

Now if task affordance exists so will *temporal affordance*: Neural activation reflecting the alternatives of *when* something is done, rather than *what*. Such activation will not be easily dissociable with imaging—because the movement itself is the same—and will manifest physiologically only with a widening of the distribution of neural activity time-aligned to the movement—i.e., a *Bereitschaftspotential*. It will also be much more complex because time is not resolvable into a fixed number of discrete categories. It is therefore impossible to tell at what time point the activity is critical to the movement actually executed rather than the alternative movements that are *not* executed, for the two categories are not neurophysiologically dissociable. Temporal affordance therefore sets up an ineliminable confound in interpreting the significance of any time point of the neural activation antecedent to the movement.

at least two actions so that the neural accompaniments of each may be positively identified against the baseline of not acting. To examine the *timing* of such activity also requires two further things. First, the action cannot be guided by an event in the external environment such as the sudden onset of a visual cue because the timing then becomes hijacked by that event, making it impossible to know whether the correlated neural activity is related to the interaction between the cue and the action or to the action alone. The choice must therefore be solely and spontaneously the agent's, without guidance from a temporally limited external cue. Second, the agent cannot know in advance what action she intends to perform on any one trial because then she is free to

decide at any antecedent time—such as the end of the preceding trial—in effect generating an operant marker that ruins the timing in the same way as if the marker were an external cue. The agent must therefore make her choice at the start of the trial and not before.

Now the first requirement is easy, but what about the second? To be confident that the subject chose freely at the start of a trial the choice on any trial must show no relation to the history of preceding choices: i.e., the sequence of choices must be random. But it has been repeatedly shown that human beings are poor at generating random sequences—there is inevitably contamination from the preceding history (Brugger, 1997; Towse & Valentine, 1997; Tune,

1964; Wagenaar, 1972). Indeed, far from ignoring the history, a human will naturally monitor the sequence—both past and projected into the future—so as to attempt to *pseudorandomize* it, whether she admits to it or not (cf. Ayton & Fischer, 2004). Note that this problem is not overcome by not giving subjects any instruction about what to do, and selecting, post hoc, those who did not respond in a readily identifiable pattern (as Soon et al., 2008 did)—indeed, the more complex the pattern the more elaborate the subject's pseudorandomization scheme is likely to have been. Thus, there are grounds here naturally to expect a response bias *at any time during the experimental run*, making it impossible to interpret the significance of the timing of any correlated neural activity. This problem undermines all “free choice” experimental designs—the very framing of the action and its context makes it impossible to obtain an interpretable result. Even if one could avoid a sequence by testing each subject on only one trial, the critical activity could occur at any time after the subject is instructed about how to do the task, unless the task instruction is immediately followed by a movement, in which case the interval of antecedence reduces to the range of choice reaction times, and the question of early antecedence cannot be tested.

Specificity for the action that follows is also uncertain here. The predictive power in Soon et al.'s study was no higher than 60%, barely above chance. Crucially, prediction here does not even imply specificity for the action, but only some specificity for the *outcome on each trial*. The earliest predictive activity may have been related to the subject's consistently thinking at this point about the *opposite* action (following an inversion of the gambler's fallacy). No causal model can be plausibly built on such data.

In short, there are currently neither cognitive nor neural antecedent events that can be confidently brought into play in causal models of voluntary action. Let us now turn to the second question of reclassifying all or some voluntary actions as covertly automatic.

Covert automaticity in voluntary action

We have seen that a key index of the voluntariness of an action is the degree of control the agent has over it, fully voluntary action being the kind where the agent has two-way power to act or abstain from acting as he pleases. But neither our observation of someone's control, nor her sincere avowal of a capacity for it, need be infallible. Where an antecedent or coincident somatic event of some kind is highly correlated with

an action, the possibility arises that the movement is directly caused by a component marked by the event, whatever the subject might say. In such circumstances, a movement that may *appear* to be voluntary, and is reported to be voluntary by the agent, may in fact be covertly automatic. That no marker is visible does not mean that none exists—new technologies such as functional brain imaging may disclose hidden somatic markers that render actions previously thought to be voluntary as *covertly* automatic. We now need to consider the criteria that must be satisfied for such re-descriptions to be valid.

Identifying a highly correlated antecedent is only the beginning of testing for covert automaticity. For example, blushing may always precede a voluntary gesture of embarrassment, but it does not causally necessitate the performance of such gestures—one is always free to make a gesture or not to make it—and so the gesture remains within the spectrum of voluntary action. Observing a correlation merely opens a *possibility* that needs to be explicitly tested by investigating whether or not the agent is able to do *otherwise* than the correlate suggests, i.e., whether she retains control and is able voluntarily to dissociate the action from its correlate.

Such a test requires an experimental behavioral paradigm where, in its most basic form, on each trial the agent is invited to make a movement for which an antecedent has been identified, and on a subset of trials is unpredictably cued—just before she acts—to *change* what she does to something else or indeed to no movement at all. The degree of voluntariness is indexed by the proportion of trials on which the subject can successfully change her action in response to the change cue. Until an antecedent–action pairing has been investigated in this way, no inference to covert automaticity can be made.

Now this “change” paradigm (and its relative the “stop-signal” paradigm) (Band & van Boxtel, 1999; Logan & Burkell, 1986; Logan & Cowan, 1984; Nachev, Rees, Parton, Kennard, & Husain, 2005), not only allows us to test any specific case but also gives us a ceiling for the maximal antecedence of any *general* correlate of action. It is so because the mean interval between the onset of a change cue and successfully refraining from making the original movement is fairly consistent across individuals and basic tasks. This value is in the region of 200 ms (Bedard et al., 2002; De Jong, Coles, Logan, & Gratton, 1990; Hanes & Carpenter, 1999). A correlated antecedent earlier than this can therefore *always* be overridden because that is how

long it maximally takes for someone to change to another action in response to an unpredictable event. Note that since an external event may be made wholly unpredictable, it is impossible for the agent consistently to anticipate it (unless she has supernatural powers).

It may be objected here that a change in response to an external event implies a “stimulus-driven” action of necessarily lesser voluntariness than spontaneous “internally-driven” action. Leaving the problems of the notion of “internal initiation” aside (see Nachev & Husain, 2010; Obhi, 2012; Schüür & Haggard, 2011 for discussion), it is false to argue that an action prompted by a stimulus must be less voluntary than a spontaneous one. Vomiting may be preceded by a wholly spontaneous experience of nausea, yawning by boredom, blushing by an embarrassing thought, etc., and yet none of the associated manifestations is thereby rendered voluntary. Conversely, a complex action such as writing may be prompted by hearing the bell signaling the start of a written examination, a football kick by the referee’s whistle, a heckle by a bad speech, etc. In direct comparison, if I ran over a pedestrian in response to a red light would I be less culpable than if I ran them over without looking at the lights? The primary criterion of voluntariness remains control, within a spectrum of possible actions the agent knows she can make, not the presence or absence of any coincident events, either in the external environment or the body of the agent.

Indeed, far from being an underestimate of the maximum priority of a causal antecedent, change reaction times in response to an external event are likely to be *longer* than those in response to a spontaneous change, for in the former case there is the added burden of detecting a visual or other sensory event. The minimum change time more accurately corresponds to the “point of no return”—the time before the actual execution of a movement after which the process is ballistic. The precise magnitude of this is controversial (see De Jong et al., 1990; Osman, Kornblum, & Meyer, 1990; Allen Osman, Kornblum, & Meyer, 1986; Verbruggen & Logan, 2009) but it is obviously less than the value above.

Now let us suppose that we identify a neural correlate within a plausible window of antecedence. If the agent can (and does) act contrary to the anticipated consequent, a causal link between the antecedent and the movement must be discounted and the action cannot be covertly automatic. If the agent’s action is always congruent with the antecedent, a causal link between the antecedent and the action remains *possible* but since no *dissociation*

from what the agent wishes to do has been observed, there are no grounds for overruling the natural interpretation of the action as voluntary.

Nonetheless, such a finding would seem to imply backward causation, for how else could an event (say) 100 ms before a voluntary action be a consequence and not a cause of it? There are two aspects to the answer. First, we cannot time an action to that temporal precision—we cannot rely on the onset of the movement, for we know that it is necessarily preceded at least by activation of the motor neurons immediately upstream of the muscle, and we have seen that there is no cognitive correlate we can fall back on. Second, the antecedent needs to be empirically proven to have temporal specificity for the action or else the significance of its timing will be extinguished by the temporal affordance confound just as with the *Bereitschaftspotential*. In neurophysiological studies where no such definitive proof of specificity is available, the observed antecedence is very short anyway: Of the order of 40 ms or less (Hanes, Patterson II, & Schall, 1998; Paré & Hanes, 2003, see Schall, 2004 for a detailed discussion). No problem of backward causation therefore arises.

It is striking that none of the extant empirical work arguing for covert automaticity has tested the question directly in this way. We have seen that Libet-style experiments depend on identifying a conceptually questionable temporal event that is not even identifiable within the temporal scale under study (“When did you decide?” often has an answer, but never in milliseconds—any more than “When did you start lunch?”). Instead of testing control—the crucial aspect—the experimenter here artificially asks the subject to suspend it temporarily, and measures the predictive power of an antecedent prior to the purported “intention time,” inferring that the antecedent must be superior to the agent if its predictive power is greater than chance at that point. Not only is this irrelevant to the key question, it unjustifiably assumes that the agent’s inclination—at *the same time as the predictive antecedent*—shows no bias toward one response or another. Thus if, in Soon et al.’s experiment subjects reported an inclination towards their subsequent responses with >60% accuracy at the same time as the predictive antecedent (i.e., exceeding its predictive power), then no superiority of the neural marker could have been asserted. Naturally, one could not easily set up such an experiment because the subject’s report of an inclination in a free-choice task would inevitably be a self-fulfilling prophecy. But that the problem is empirically intractable does not relieve the

experimenter of the burden of producing an explicit comparison if the claim is to stand.

In short, there are currently no grounds for reclassifying any voluntary action as covertly automatic based on antecedent predictors.

CONCLUSIONS

We have seen that the study of the antecedents of voluntary action is contaminated by conceptual unclarity that make the hypotheses being tested incoherent and therefore untestable. To make progress here we must take heed of the following points.

First, voluntary action is not generally preceded by any identifiable occurrent mental event such as an urge, a “felt intention,” or an “act of will.” Urges aside, the antecedents of voluntary action do not have a “volume of experience” on which sharp temporal parameters can be pinned. Urges included, if the subsequent movement is causally linked to the antecedent, then the action becomes not voluntary but automatic, and so not the kind of action we are trying to explain. If these antecedents are not causally linked to the action, then their neural underpinning need not be causally linked to the action either. The neural processes here tell us little about the causation of voluntary action. What needs to be investigated are not the neural causes of voluntary action, but rather the neural conditions that enable (make possible) voluntary action.

Second, temporal priority does not equal causal priority. The cause of an event cannot occur after the event but it can occur *any* time before or concurrently with the event. Within this interval, time can impose a decisive hierarchy on the contribution of any component only if there is a serial chain of events. But the brain is not a serial deterministic system, like a mechanical clock mechanism. If we know anything about the brain, it is that it is vastly parallel in its organization and that the causal connections between any serial neural components are complex. Imagine two neural components: An early component whose activity predicts an action with 10% accuracy, and a late component whose activity predicts with 90% accuracy—is the early component properly referred to as the ultimately causal agent because it is earlier? Simply by noting that a creature is born with a primate brain—an event so antecedent it precedes cognitive experience—we can predict with at least 50% certainty that it will seek to avoid spiders—does that rob it of freedom to decide what to do when it sees a spider? Clearly, our index here has to be primarily predictive power, not antecedence.

Third, the occurrence of neural activity antecedent to action robs the action of its voluntariness only if the predictions based upon it hold irrespective of the agent’s wanting to act otherwise. We have yet to find a neural marker that provides the basis for such prediction. Hence, we have no grounds for reclassifying any kind of avowed voluntary action as covertly involuntary.

Fourth, even if a voluntary action were to be found to be covertly involuntary in this way, there is a hard limit on the maximal interval between the neural marker and the execution of the movement given by observed behavior. If one can voluntarily and unpredictably change from one action to another at latencies well under 200 ms, it is not possible for any involuntary component to take longer than this. To be sure, in certain special situations there may be actions or classes of action that can be predicted *at any time* such as compulsive behavior in association with lesions of the brain, say. But an interval true of voluntary action *generally* cannot be longer than this—the neural here cannot possibly override the observed behavior.

Fifth, rare discrepancies between a subject’s actions and her sincere report—before, during, and after an action—have no great physiological or philosophical significance. Acting and reporting on one’s actions are things that human beings have the powers to do. There is no necessity for one to be causally linked to the other, and hence, that there should be departures from their agreement in special circumstances, is unproblematic. Equally, there is no necessity for the neural structures on which acting depends to be shared with the neural structures on which reporting on one’s actions depends. Indeed, given the functional specialization of the brain it would be very odd for these different things to be dependent on the same substrate.

Original manuscript received 24 August 2013

Revised manuscript received 9 June 2014

First published online 30 June 2014

REFERENCES

- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness? *Memory & Cognition*, *32*(8), 1369–1378. doi:10.3758/BF03206327
- Band, G. P. H., & van Boxtel, G. J. M. (1999). Inhibitory motor control in stop paradigms: Review and reinterpretation of neural mechanisms. *Acta Psychologica*, *101*(2–3), 179–211. doi:10.1016/S0001-6918(99)00005-0
- Banks, W. P., & Isham, E. A. (2009). We infer rather than perceive the moment we decided to act. *Psychological*

- Science*, 20(1), 17–21. doi:10.1111/j.1467-9280.2008.02254.x
- Banks, W. P., & Isham, E. A. (2010). Do we really know what we are doing? Implications of reported time of decision for theories of volition. *Conscious Will and Responsibility*, 47–60. doi:10.1093/acprof:oso/9780195381641.003.0006
- Bedard, A. -C., Nichols, S., Barbosa, J. A., Schachar, R., Logan, G. D., & Tannock, R. (2002). The development of selective inhibitory control across the life span. *Developmental Neuropsychology*, 21(1), 93–111. doi:10.1207/S15326942DN2101_5
- Brugger, P. (1997). Variables that influence the generation of random sequences: An update. *Perceptual and Motor Skills*, 84(2), 627–661. doi:10.2466/pms.1997.84.2.627
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1585–1599. doi:10.1098/rstb.2007.2054
- Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, 45(5), 801–814. doi:10.1016/j.neuron.2005.01.027
- De Jong, R., Coles, M. G., Logan, G. D., & Gratton, G. (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), 164–182. doi:10.1037/0096-1523.16.1.164
- Desmurget, M., Reilly, K. T., Richard, N., Szathmari, A., Mottolese, C., & Sirigu, A. (2009). Movement intention after parietal cortex stimulation in humans. *Science*, 324(5928), 811–813. doi:10.1126/science.1169896
- Eimer, M. (1998). The lateralized readiness potential as an on-line measure of central response activation processes. *Behavior Research Methods, Instruments, & Computers*, 30(1), 146–156. doi:10.3758/BF03209424
- Fried, I., Katz, A., McCarthy, G., Sass, K. J., Williamson, P., Spencer, S. S., & Spencer, D. D. (1991). Functional organization of human supplementary motor cortex studied by electrical stimulation. *The Journal of Neuroscience*, 11(11), 3656–3666.
- Gibson, J. J. (1977). *"The theory of affordances," in perceiving, acting, and knowing. Towards an ecological psychology.* Hoboken, NJ: John Wiley & Sons Inc.
- Grèzes, J., & Decety, J. (2002). Does visual perception of object afford action? Evidence from a neuroimaging study. *Neuropsychologia*, 40(2), 212–222. doi:10.1016/S0028-3932(01)00089-6
- Hacker, P. M. S. (2007). *Human nature: The categorical framework.* Oxford: Wiley.
- Haggard, P. (2008). Human volition: Towards a neuroscience of will. *Nature reviews. Neuroscience*, 9(12), 934–946. doi:10.1038/nrn2497
- Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126(1), 128–133. doi:10.1007/s002210050722
- Haggard, P., & Libet, B. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies*, 8(11), 47–64.
- Hanes, D. P., & Carpenter, R. H. S. (1999). Countermanding saccades in humans. *Vision Research*, 39(16), 2777–2791. doi:10.1016/S0042-6989(99)00011-5
- Hanes, D. P., Patterson II, W. F., & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology*, 79(2), 817–834.
- Hazelhoff, F. F., & Wiersma, H. (1924). Die Wahrnehmungszeit. *Zeitschrift Für Psychologie*, 96, 171–188.
- Jackson, S. R., Parkinson, A., Kim, S. Y., Schürmann, M., & Eickhoff, S. B. (2011). On the functional anatomy of the urge-for-action. *Cognitive Neuroscience*, 2(3–4), 227–243. doi:10.1080/17588928.2011.604717
- Kornhuber, H. H., & Deecke, L. (1965). Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. *Pflüger's Archiv für die gesamte Physiologie des Menschen und der Tiere*, 284(1), 1–17. doi:10.1007/BF00412364
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-Potential) the unconscious initiation of a freely voluntary act. *Brain*, 106(3), 623–642. doi:10.1093/brain/106.3.623
- Lim, S. H., Dinner, D. S., Pillay, P. K., Lüders, H., Morris, H. H., Klem, G., ... Awad, I. A. (1994). Functional anatomy of the human supplementary sensorimotor area: Results of extraoperative electrical stimulation. *Electroencephalography and Clinical Neurophysiology*, 91(3), 179–193. doi:10.1016/0013-4694(94)90068-X
- Logan, G. D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation: A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, 12(4), 549–563. doi:10.1037/0096-1523.12.4.549
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91(3), 295–327. doi:10.1037/0033-295X.91.3.295
- Mackie, J. L. (1974). *The cement of the universe.* Oxford: Clarendon Press.
- Nachev, P. (2011). Urges, inhibition, and voluntary action. *Cognitive Neuroscience*, 2(3–4), 247–248.
- Nachev, P., & Husain, M. (2010). Action and the fallacy of the "internal": Comment on Passingham et al. *Trends in Cognitive Sciences*, 14(5), 192–193. doi:10.1016/j.tics.2010.03.002
- Nachev, P., Kennard, C., & Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature reviews. Neuroscience*, 9(11), 856–869. doi:10.1038/nrn2478
- Nachev, P., Rees, G., Parton, A., Kennard, C., & Husain, M. (2005). Volition and conflict in human medial frontal cortex. *Current Biology: CB*, 15(2), 122–128. doi:10.1016/j.cub.2005.01.006
- Obhi, S. S. (2012). The troublesome distinction between self-generated and externally triggered action: A commentary on. *Consciousness and Cognition*, 21(1), 587–588. doi:10.1016/j.concog.2011.09.014
- Osman, A., Kornblum, S., & Meyer, D. E. (1986). The point of no return in choice reaction time: Controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and*

- Performance*, 12(3), 243–258. doi:10.1037/0096-1523.12.3.243
- Osman, A., Kornblum, S., & Meyer, D. E. (1990). Does motor programming necessitate response execution? *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), 183–198. doi:10.1037/0096-1523.16.1.183
- Paré, M., & Hanes, D. P. (2003). Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *The Journal of Neuroscience*, 23(16), 6480–6489.
- Rothwell, J. C., & Edwards, M. J. (2011). An urge to act or an urge to suppress?. *Cognitive Neuroscience*, 2(3–4), 250–251.
- Schall, J. D. (2004). On building a bridge between brain and behavior. *Annual Review of Psychology*, 55, 23–50. doi:10.1146/annurev.psych.55.090902.141907
- Schopenhauer, A. (1859). *Die Welt als Wille und Vorstellung*. (3rd edition). FA Brockhaus: Leipzig.
- Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, 109(42), E2904–E2913. doi:10.1073/pnas.1210467109
- Schüür, F., & Haggard, P. (2011). What are self-generated actions? *Consciousness and Cognition*, 20(4), 1697–1704. doi:10.1016/j.concog.2011.09.006
- Shapiro, A. K., Shapiro, E. S., Gerald, J., & Feinberg, T. E. (1988). *Gilles de la Tourette syndrome* (2nd ed., Vol. xxiv). New York, NY, England: Raven Press, Publishers.
- Soon, C. S., Brass, M., Heinze, H. -J., & Haynes, J. -D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, 11(5), 543–545. doi:10.1038/nn.2112
- Sumner, P., & Husain, M. (2008). At the edge of consciousness: Automatic motor activation and voluntary control. *The Neuroscientist*, 14(5), 474–486. doi:10.1177/1073858408314435
- Tehovnik, E. J. (1996). Electrical stimulation of neural tissue to evoke behavioral responses. *Journal of Neuroscience Methods*, 65(1), 1–17. doi:10.1016/0165-0270(95)00131-X
- Tipper, S. P., Paul, M. A., & Hayes, A. E. (2006). Vision-for-action: The effects of object property discrimination and action state on affordance compatibility effects. *Psychonomic Bulletin & Review*, 13(3), 493–498. doi:10.3758/BF03193875
- Towse, J. N., & Valentine, J. D. (1997). Random generation of numbers: A search for underlying processes. *The European Journal of Cognitive Psychology*, 9(4), 381–400. doi:10.1080/713752566
- Tune, G. S. (1964). A brief survey of variables that influence random-generation. *Perceptual and Motor Skills*, 18, 705–710. doi:10.2466/pms.1964.18.3.705
- Verbruggen, F., & Logan, G. D. (2009). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews*, 33(5), 647–661. doi:10.1016/j.neubiorev.2008.08.014
- von Wright, G. H. (1963). *Norm and action: A logical enquiry*. Vol. 15. New York: Routledge and Kegan Paul.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of literature. *Psychological Bulletin*, 65–72. doi:10.1037/h0032060
- Walter, H. (2001). *Neurophilosophy of free will: From libertarian illusions to a concept of natural autonomy*. Cambridge, MA: MIT Press.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.