

Lecture 01 : Philosophical Issues in Behavioural Science

Stephen A. Butterfill
< s.butterfill@warwick.ac.uk >

Tuesday, 11th January 2022

Contents

1	Introduction: Why Investigate Philosophical Issues in Behavioural Science?	3
2	Asking Questions	3
2.1	How Many Questions Should I Ask?	4
2.2	How to Use the 'Ask a Question' Feature	4
2.3	Deadline for Questions	4
3	How to Use the Online Lectures	5
4	Assessment and Other Components of This Course	5
4.1	How will I be assessed and given feedback?	6
4.2	Where are the questions?	6
4.3	When are the deadlines?	6
5	Seminar Tasks (yyrama)	6
5.1	Walkthrough	6
6	Instrumental Actions: Goal-Directed and Habitual	7
6.1	Terminology	7
6.2	Main Question	7
6.3	A Standard Answer	7
6.4	A Clue from Animal Learning	8
6.5	What Are Habitual Processes?	9
6.6	Habitual Processes and Instrumental Action	10
6.7	So What?	10
7	Philosophical Theories of Action	10
7.1	The Problem of Action	11

7.2	A Standard Solution	11
7.3	Davidson on Agency	12
7.4	Two Questions, One Answer	14
8	Conclusion	14
9	Question Session 01	14
9.1	The Questions	15
9.2	Identifying a Form of Representation (Tiago's Question) . .	15
9.3	Interpreting the Dual-Process Theory of Instrumental Action	15
	Glossary	16

1. Introduction: Why Investigate Philosophical Issues in Behavioural Science?

Answering broadly philosophical questions about action and about joint action requires reflection on discoveries from the behavioural sciences.

Here's the course in essence. Two questions. One about individual action, the other about collaborative action:

1. Which events in your life are your actions?
2. What distinguishes doing something jointly with another person from acting in parallel with them but merely side by side?

We start with the first question then move on to the second.

We will consider each question from two perspectives, philosophical and scientific.

Although the perspectives are distinct, they are not independent. As things stand, we cannot adequately answer broadly philosophical questions about action, nor about joint action, without reflection on discoveries from the behavioural sciences.

2. Asking Questions

Aim to ask at least three questions during this course. A significant part of your work on this course is to formulate and pose questions in response to the lecture materials (or, if you prefer, in response to the works cited in them).

I sometimes hear people say, 'there's no such thing as a silly question.' This is obviously false. As you know, many questions arise from thoughtlessness, laziness or vanity. (And all three often feature together.)

But question asking is a skill. You cannot improve without practicing it. In asking mostly silly questions, you are attempting to improve your skill with the goal, eventually, of asking better questions.

Genuinely good questions are rare and precious. Identifying and articulating such questions is hard work.

Philosophy is done by asking questions. The questions are not merely a means to learning about philosophy: doing philosophy consists, in part, in asking questions.

As you work through each lecture, you should be attempting to identify and articulate questions. This is a core part of your work. The questions you identify should also be the foundation of your writing.

That's why you will see a section headed 'Ask a Question' on each page of these lecture notes.

When you have a question:

- Discuss it with your lecture buddy or buddies.
- Post it under 'Ask a Question' in the relevant section of the lecture notes.
- Ask it in your seminar.

2.1. How Many Questions Should I Ask?

You don't need to ask a question every week. But you should aim to ask at least three during the course.

2.2. How to Use the 'Ask a Question' Feature

To use this feature, you need to sign up for a github account.¹ (Github is a kind of instaface/whatstok for nerds.) You then need to hit the 'sign in with github' button below. You will be asked to allow access for something called 'utteranc.es' (this is the service that powers the comments). Please do this. You are now ready to ask your first question.

Please do use this feature. Ask your questions at the bottom of the lecture notes from which it arises. That way, your lecturer can collate questions for the live question session and prepare in advance. Also everyone can see and think about the questions.

2.3. Deadline for Questions

Please ask your questions by the day before the question session (Wednesday). As you'll see, I need time to think about them, to order the questions, and to prepare some slides.

I will also sometimes need to do some new research to answer your questions. So sooner is better.

¹ I know people will complain and give me bad feedback for this. 'So much web sites and accounts I need to sign in for this course!' 'Lecturer should just drone over they powerpoint slides and upload them to the moodle!' (This might be true. But if inclined to this view, why are you reading my endnotes?)

3. How to Use the Online Lectures

Watch with a friend, and talk. Take notes. Use the 2x speed option. Skip around. Most importantly, ask questions.

You can use the online lectures however you like, of course. You know how you learn best. But I do have some ideas ...

Watch with a friend, and talk as you watch.

Take notes. (The notes on these pages are intended to provide some key quotes and a list of references to save you some writing; they do not cover everything in the lecture recordings.)

Speed the videos up. I try to speak slowly enough that you may be able to watch at 1.5x or 2x speed.

The lecture notes do not cover everything in the recording; the recording does not cover everything in the lecture notes. Despite this, you are not expected to both read everything and watch everything. Be selective.

If you understand a lecture recording, you probably don't need to read the notes as well.² Alternatively, reading the notes might mean you don't need to watch the recording (if in doubt, use the `slides only` tab to quickly see what is in the recording).

Skip around. If it's too boring, move on. Don't aim to use all the recordings and notes.

Use the search function (top of each page).

Don't feel pressure to complete everything. Do what you find fulfilling.

Most importantly, **ask questions**:

- Put them to your lecture buddy or buddies.
- Ask them in for the next Whole-Class Live Question Session.
- Ask them in your seminar.

4. Assessment and Other Components of This Course

How your assessment breaks down, what the formative (non-assessed) work is, and what the main events each week are for.

² And you are not required to read the footnotes.

4.1. How will I be assessed and given feedback?

This is explained in a handy pdf:

- how you will be assessed and given feedback

4.2. Where are the questions?

The in-term essays can be found on can be found on yyrama.

There are possible titles for your longer essay [pdf]. You may also devise your own question through discussion with me; not that your question will need to be approved.

4.3. When are the deadlines?

Deadlines for assessed work are on tabula.

Deadlines for the weekly seminar tasks depend on which seminar group you are an and can be found on yyrama.

5. Seminar Tasks (yyrama)

The most important work on this course, apart from the assessments, is the weekly seminar tasks. You need to submit some work before your seminar each week. This mostly involves writing, or re-writing, a mini essay as well as peer-reviewing another student's work. The seminars exist for you to discuss your writing.

Sign up on yyrama and let yyrama know which seminar group you are in.

You can find the weekly essays and peer reviews here:

<https://yyrama.butterfill.com/course/view/philosophicalIssuesInBehaviouralScience>

If you attend a different seminar group one week, please update your seminar group on yyrama so that your work goes to your tutor.

5.1. Walkthrough

I rarely find students have a problem using yyrama, but in case you want one there is a step-by-step guide from 4:25 in this video.

6. Instrumental Actions: Goal-Directed and Habitual

An instrumental action is an action that happens in order to bring about an outcome. When you press a lever in order to retrieve a snack, or when you board a bus in order to travel home, you are performing an instrumental action. What grounds the relation between an instrumental action and the outcome it occurs in order to bring about? This section introduces a key distinction between two answers to this question, goal-directed and habitual.

6.1. Terminology

An instrumental action is an action that happens in order to bring about an outcome. We will say that the outcome is a goal of the action,³ and that the action is directed to the outcome.

6.2. Main Question

What is the relation between an instrumental action and the outcome or outcomes to which it is directed?

6.3. A Standard Answer

One standard answer to this question involves intention. An intention specifies an outcome, coordinates your actions, and coordinates your actions in a way that would normally increase the probability of the specified outcome occurring. So if an intention causes you to act, it follows that your action happens in order to bring about the outcome intended. And this implies that your action is instrumental.

What is an intention? Although there is much debate about this (Setiya 2014), for our purposes only a widely agreed characteristic is necessary. Intentions are the upshot of beliefs and desires (or are identical to one or both of these). To illustrate:

desire: I fill Zak's glass.

belief: If I pour, I will fill Zak's glass.

intention: I pour to fill Zak's glass.

³ Be careful not to confuse a goal with a goal-state, which is an intention or other state of an agent linking an action to a particular goal to which it is directed. (Some authors use the term 'goal' for goal-states rather than outcomes.) A goal is a possible or actual outcome (such as filling a glass with prosecco). A goal-state is a psychological attribute of an agent (such as an intention to fill a glass with prosecco).

This simplistic example captures a key idea. Behind an intention lie two things. There is a desire to bring an outcome about, and there is a belief about which action will bring the action about.⁴

If you would like more background on action and intention, see Lecture 10 of Mind and Reality.

Our Main Question is about the relation between an instrumental action and the outcome or outcomes to which it is directed. According to the Standard Answer, the relation involves belief, desire and intention:

Background Assumption: Instrumental actions are caused by intentions to bring outcomes about, which are the upshot of desires to bring outcomes about and beliefs that certain actions will bring them about.

Standard Answer: The outcome (or outcomes) to which an instrumental action is directed is that outcome (or outcomes) specified by the intention (or intentions) which caused it.

Does the Standard Answer involving intention provide a full answer to that question? Or are there things other than intentions which might link an instrumental action to an outcome?

6.4. A Clue from Animal Learning

According to Dickinson (2016, p. 177):

‘instrumental behavior is controlled by two dissociable processes: a goal-directed and an habitual process’

He goes on to specify what the ‘goal-directed process’ involves:

‘an action is goal-directed if it is mediated by the interaction of a representation of the causal relationship between the action and outcome and a representation of the current incentive value, or utility, of the outcome in a way that rationalizes the action as instrumental for attaining the goal’ Dickinson (2016, p. 177).

Dickinson’s ‘goal-directed process’ corresponds to the belief–desire model we just considered. The ‘representation of the causal relationship between the action and outcome’ could be a belief about which action will bring an outcome about (e.g. the belief that if I pour, I will fill Zak’s glass). And the ‘representation of the current incentive value, or utility, of the outcome’ could be a desire.

⁴ We will see the same structure when we come to decision theory (in *Expected Utility* in Lecture 06). Preferences correspond to desires and expected utilities to beliefs.

philosophy	animal learning	decision theory
belief	representation of the causal relationship between the action and outcome	subjective probability
desire	representation of the current incentive value, or utility, of the outcome	preference

Table: rough correspondence between terms used for modelling action across three disciplines.

So when Dickinson says that instrumental actions are ‘controlled by two dissociable processes’, he is implying that the Standard Answer about belief, desire and intention cannot fully explain the relation between an instrumental action and the outcome or outcomes to which it is directed. If he is right, we also have to consider something he calls ‘an habitual process’.

6.5. What Are Habitual Processes?

Habitual processes involve connections between stimuli and actions. For example, the presence of an empty glass (a stimulus) may be connected to the action of pouring. These connections are characterised by two features:

1. When the action is performed in the presence of the stimulus, the connection between action and stimulus is strengthened (or ‘reinforced’) if the action is rewarded.
2. If the connection is strong enough, the presence of the stimulus will cause the action to occur.

This is another way of stating *Thorndyke’s Law of Effect*:

‘The presentation of an effective [=rewarding] outcome following an action [...] reinforces a connection between the stimuli present when the action is performed and the action itself so that subsequent presentations of these stimuli elicit the [...] action as a response’ (Dickinson 1994, p.48).

How do habitual processes differ from those involving belief, desire and intention? Two differences are important for our purposes:

1. The effects of habitual processes do not depend on what you currently desire. This is because the strength of the

stimulus–action connection depends only on what was rewarding for you in the past, not what is rewarding for you now.

2. The effects of habitual processes do not depend on what you currently believe about which outcome the action will have. This is because the strength of the stimulus–action connection depends only on what outcomes the action had in the past, not on which outcomes it will have now.

Because habitual processes have these features, we can be sure that they are genuinely distinct from processes involving belief, desire and intention.

6.6. Habitual Processes and Instrumental Action

Our Main Question is, What is the relation between an instrumental action and the outcome or outcomes to which it is directed? This question can be answered by invoking habitual processes. For if an action is due to an habitual process, then there is a stimulus–action connection which caused it. This stimulus–action connection must have been strengthened in the past because, often enough, some (one or more) rewarding outcomes occurred when the action was performed in the presence of the stimulus. But since habitual processes exist to enable the agent repeatedly bring about such rewarding outcomes, it follows that the action occurs now in order to bring about these (one or more) rewarding outcomes. That is, the action is directed to the outcome; it is an instrumental action.

The Standard Answer therefore fails to provide a full answer to the Main Question about instrumental action. To fully answer it we need not only belief, desire and intention but, minimally, also the kind of stimulus–action connections involved in habitual processes.

6.7. So What?

After this section, you should understand what an instrumental action is, you should understand the Main Question, and you should understand how habitual processes and goal-directed processes differ.

The next step is to investigate possible consequences for philosophical theories of action.

7. Philosophical Theories of Action

Much philosophy of action starts with The Problem of Action: What distinguishes your actions from things that merely happen to you (Davidson,

1971)? According to a standard, widely-accepted solution, actions are those events which stand in an appropriate causal relation to an intention. This is an instance of the Causal Theory of Action, according to which an event is action ‘just in case it has a certain sort of psychological cause’ (Bach, 1978, p. 361). This section explores some of the reasoning supporting the standard solution. Eventually, though, we will have to ask whether discoveries about habitual processes pose any kind of challenge to the philosophers.

7.1. The Problem of Action

Much philosophy of action hinges on the question, What distinguishes your actions from things that merely happen to you (Davidson 1971)?

You trip and fall down a flight of stairs. Falling is something that happens to you, not an action of yours. But watching the sympathetic attention you gain, Buster expertly throws himself down the stairs. Although it looks like another accident, this event is an action.

As Frankfurt (1978, p. 157) put it:

‘The⁵ problem of action is to explicate the contrast between what an agent does and what merely happens to him.’

But is this really a problem? It may be tempting, initially, to suppose that we can answer this question by invoking kinematic features. Perhaps—so the idea—actions are those events which involve some or other patterns in the joint displacements and bodily configurations? Alternatively, it might be tempting to think that we can answer the question by appeal to coordination. Perhaps—so the thought—actions are those events which involve a particular coordination of body parts? If either possibility obtained, the ‘problem of action’ would not be a problem at all. But reflection on the variety of things that count as actions indicates that neither of these initially tempting possibilities is at all likely to obtain. Or so I argue in *Recap: Action* from the lectures on Mind and Reality.

The absence of straightforward answers to the question about what distinguishes actions from things that merely happen to you indicates that it is a genuine problem.

7.2. A Standard Solution

According to a standard, widely-accepted view, actions are those events which stand in an appropriate causal relation to an intention. What distin-

⁵ I dislike this way of stating things. Good philosophers come up with lots of questions. There is insufficient reason to single one of them out as *the* problem.

guishes your falling from Buster's is that his, but not yours, was appropriately related to an intention.

This is an instance of the Causal Theory of Action. According to this view, an event is action 'just in case it has a certain sort of psychological cause' (Bach 1978, p. 361). Proponents of this view may disagree about which states cause actions (Bach is an example of this), or about how to characterise the causal relation (for example, Frankfurt (1978) is concerned, in part, with whether the causes are antecedent to the action or provide ongoing guidance). But they agree that the relation between actions and their psychological causes is what distinguishes your actions from things that merely happen to you.

7.3. Davidson on Agency

How does Davidson arrive at the view that actions are those events which stand in an appropriate causal relation to an intention?⁶

As background, Davidson notes that the same action can be described in multiple ways. You move your finger, flicking a switch which causes the lights to come on and alerts a prowler (Davidson 1971, p. 53). We have four ways of describing one and the same action: as moving your finger, as flicking a switch, and so on.

Davidson further notes that actions can typically be described both in ways that relate to what you intended (turning the lights on, say) and in ways which do not relate to your intentions (alerting a prowler, perhaps).

This background allows Davidson to distinguish three situations involving someone spilling coffee:

'If [...] I intentionally spill the contents of my cup, mistakenly thinking it is tea when it is coffee, then spilling the coffee is something I do, it is an action of mine, though I do not do it intentionally. On the other hand, if I spill the coffee because you jiggle my hand, I cannot be called the agent. Yet while I may hasten to add my excuse, it is not incorrect, even in this case, to say I spilled the coffee. Thus we must distinguish three situations in which it is correct to say I spilled the coffee: in the first, I do it intentionally; in the second I do not do it intentionally but it is my action (I thought it was tea); in the third it is not my action

⁶ I've heard people who should know say that Davidson does not explicitly commit to this view. But Davidson writes, 'we have discovered no analysis of this relation that does not appeal to the concept of intention' (Davidson 1971, p. 61). And nowhere does he explicitly reject the view that actions are those events which stand in an appropriate causal relation to an intention.

at all (you jiggle my hand).’ (Davidson 1971, p. 45)

In short my spilling the coffee can be caused in three ways:

1. by an intention of mine to spill the coffee;
2. by an intention of mine to spill the tea (where I mistakenly take the coffee to be tea and do not intend to spill coffee);
or
3. by you jiggling my hand (where no intention of mine is directly involved at all).

My spilling the coffee is an action of mine in (1) and (2), but not in (3).

Reflection on (1) and (2) rules out the view that my spilling the coffee is an action of mine only if I intend to spill the coffee.

The contrast between (2) and (3) is what leads Davidson to his view about agency:

‘What is the difference [between (2) and (3)]? The difference seems to lie in the fact that in one case, but not in the other, I am intentionally doing something. My spilling the contents of my cup was intentional; as it happens, this very same act can be redescribed as my spilling the coffee. Of course, thus redescribed the action is no longer intentional; but this fact is apparently irrelevant to the question of agency.

‘And so I think we have one correct answer to our problem: a man is the agent of an act if what he does can be described under an aspect that makes it intentional.’ (Davidson 1971, p. 46)

Suppose we assume, further, that an act can be described under an aspect that makes it intentional only if it stands in an appropriate causal relation to an intention of the agent’s.⁷ Then the Standard Solution mentioned above follows:

Your actions are those events which stand in an appropriate causal relation to an intention of yours.

⁷ Is this assumption true? Bratman allows that actions can be intentional ‘even though [the agent] has no distinctive attitude of intending’ (Bratman 1987, p. 132), and even though the agent lacks the capacity to form intentions altogether (Bratman 2000, p. 51). This view follows from two claims: first, intentions are distinct from any combination of beliefs and desires; and second, beliefs and desires alone may, in certain cases, determine what an agent intentionally does.

7.4. Two Questions, One Answer

We have now encountered intention as providing the standard response to two questions about action:

Question 1: What is the relation between an instrumental action and the outcome or outcomes to which it is directed? (see *Instrumental Actions: Goal-Directed and Habitual* (section §6))

Standard Answer: The outcome (or outcomes) to which an instrumental action is directed is that outcome (or outcomes) specified by the intention (or intentions) which caused it.

Question 2: What distinguishes your actions from things that merely happen to you? (The Problem of Action, this section)

Standard Solution: Your actions are those events which stand in an appropriate causal relation to an intention of yours.

On Question 1, the existence of habitual processes demonstrates that the Standard Answer to Question 1 is at best incomplete (see *Instrumental Actions: Goal-Directed and Habitual* (section §6)).

Our next issue concerns the second question. Does the existence of habitual processes also pose any kind of challenge to how philosophers standardly answer The Problem of Action?

8. Conclusion

In this lecture we have begun to think about instrumental action from the point of view of theories of animal learning, distinguishing habitual from goal-directed processes. And we have considered action from the point of view of philosophy of action, focussing on *The Problem of Action* and the notion of intention.

9. Question Session 01

If available (no promises), recordings of the live whole-class lecture will be here, together with slides and references. They are usually available on the day after the session. (You may need to refresh this page to make them appear.)

The recording is not on youtube, only streams.

9.1. The Questions

This session covers three questions:

- [Abiopa] Do habitual processes not still occur with the desire/intention of bringing about the same outcome as in the past? (Two answers are offered: no; and sort of.)
- [Tiago] How should we understand the term ‘representation’? (More on this below.)
- [Jan] Will there be a reading list for this module? (I take you through the various ways of construction one, whether you need the bare minimum to pass or can dedicate some hours to research for the course each week.)

9.2. Identifying a Form of Representation (Tiago’s Question)

On a widely accepted view, representations involve subjects having attitudes toward contents. Possible attitudes include believing, wanting, intending and knowing. The content is what distinguishes one belief from all others, or one desire from all others. The content is also what determines whether a belief is true or false, and whether a desire is satisfied or unsatisfied.

There are three main tasks in specifying a form of representation. The first task is to identify its subject (a person, perhaps; but not necessarily).

The second task is to characterise some attitudes. This typically involves specifying their distinctive functional and normative roles.⁸

The third task is to find a scheme for specifying the contents of mental states. This typically involves one or another kind of proposition, although some have suggested other abstract entities including map-like representations.⁹

There may be reasons to postulate further aspects of representations; later in the course we will encounter an argument for the view that representations can differ in format as well as in content.

9.3. Interpreting the Dual-Process Theory of Instrumental Action

In formulating the dual-process theory of instrumental action, Dickinson (2016, p. 177) mentions representations but does not explicitly identify sub-

⁸ For examples, see Bratman (1987) on intention or Velleman (2000, chapter 11) on belief.

⁹ See Braddon-Mitchell & Jackson (1996, p. 163): ‘what is inside our heads should be thought of as more like maps than sentences.’

ject, attitude or scheme for specifying content. How should we do this?

One possibility would be to identify the representations with beliefs and desires (as hinted at in the table in *Instrumental Actions: Goal-Directed and Habitual* (section §6)). In this case we incur commitments related to features associated with beliefs and desires, such as the inferential integration of belief. If the representations involved in goal-directed processes lack these features, the identification of them with beliefs and desires would fail (and I think Klossek et al. (2011)'s findings, discussed in *Goal-Directed and Habitual: Some Evidence* in Lecture 02, could provide grounds to suspect this).

A different possibility would be to take Dickinson's characterisation of goal-directed processes as providing an implicit functional role, and then use this to characterise the attitude.

Glossary

Causal Theory of Action According to this view, an event is action 'just in case it has a certain sort of psychological cause' (Bach 1978, p. 361). 12

directed For an action to be *directed* to an outcome is for the action to happen in order to bring that outcome about. 7, 10, 14

dual-process theory of instrumental action instrumental action 'is controlled by two dissociable processes: a goal-directed and an habitual process' (Dickinson 2016, p. 177). 15

goal A *goal* of an action is an outcome to which it is directed. 7

goal-directed process A process which involves 'a representation of the causal relationship between the action and outcome and a representation of the current incentive value, or utility, of the outcome' and which influences an action 'in a way that rationalizes the action as instrumental for attaining the goal' (Dickinson 2016, p. 177). 10

goal-state an intention or other state of an agent which links an action of hers to a particular goal to which it is directed. 7

habitual process A process underpinning some instrumental actions which obeys *Thorndyke's Law of Effect*: 'The presentation of an effective [=rewarding] outcome following an action [...] reinforces a connection between the stimuli present when the action is performed and the action itself so that subsequent presentations of these stimuli elicit the [...] action as a response' (Dickinson 1994, p.48). 10, 11, 14

inferential integration For states to be *inferentially integrated* means that: (a) they can come to be nonaccidentally related in ways that are approximately rational thanks to processes of inference and practical reasoning; and (b) in the absence of obstacles such as time pressure, distraction, motivations to be irrational, self-deception or exhaustion, approximately rational harmony will characteristically be maintained among those states that are currently active. 16

instrumental action An action is *instrumental* if it happens in order to bring about an outcome, as when you press a lever in order to obtain food. (In this case, obtaining food is the outcome, lever pressing is the action, and the action is instrumental because it occurs in order to bring it about that you obtain food.)

You may find variations on this definition of *instrumental* in the literature. Dickinson (2016, p. 177) characterises instrumental actions differently: in place of the teleological ‘in order to bring about an outcome’, he stipulates that an instrumental action is one that is ‘controlled by the contingency between’ the action and an outcome. And de Wit & Dickinson (2009, p. 464) stipulate that ‘instrumental actions are *learned*’. 7, 8, 10, 14, 16

outcome An outcome of an action is a possible or actual state of affairs. 7, 8, 10, 14

problem a question that is difficult to answer. 11

stimulus A *stimulus* is just a situation or event. Typically, ‘stimulus’ is used to label things which do, or might, prompt actions such as the presence of a lever or the flashing of a light. 9

The Problem of Action What distinguishes your actions from things that merely happen to you? (According to Frankfurt (1978, p. 157), ‘The problem of action is to explicate the contrast between what an agent does and what merely happens to him.’) 14

References

- Bach, K. (1978). A representational theory of action. *Philosophical Studies*, 34(4), 361–379.
- Braddon-Mitchell, D. & Jackson, F. (1996). *Philosophy of Mind and Cognition*. Oxford: Blackwell.

- Bratman, M. E. (1987). *Intentions, Plans, and Practical Reasoning*. Cambridge, MA: Harvard University Press.
- Bratman, M. E. (2000). Valuing and the will. *Noûs*, 34(supplement 14), 249–265. Reprinted in Bratman, M. (2007) *Structures of Agency*. Oxford: Oxford University Press (pp. 47-67).
- Davidson, D. (1971). Agency. In R. Binkley, R. Bronaugh, & A. Marras (Eds.), *Agent, Action, and Reason*, (pp. 3–25). Toronto: University of Toronto Press. Reprinted in Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Oxford University Press.
- de Wit, S. & Dickinson, A. (2009). Associative theories of goal-directed behaviour: A case for animal–human translational models. *Psychological Research PRPF*, 73(4), 463–476.
- Dickinson, A. (1994). Instrumental conditioning. In N. Mackintosh (Ed.), *Animal Learning and Cognition*. London: Academic Press.
- Dickinson, A. (2016). Instrumental conditioning revisited: Updating dual-process theory. In J. B. Trobalon & V. D. Chamizo (Eds.), *Associative learning and cognition*, volume 51 (pp. 177–195). Edicions Universitat Barcelona.
- Frankfurt, H. (1978). The problem of action. *American Philosophical Quarterly*, 15(2), 157–162.
- Klossek, U. M. H., Yu, S., & Dickinson, A. (2011). Choice and goal-directed behavior in preschool children. *Learning & Behavior*, 39(4), 350–357.
- Setiya, K. (2014). Intention. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 ed.). Metaphysics Research Lab, Stanford University.
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford: Oxford University Press.