

Dennett's dual-process theory of reasoning *

Keith Frankish

1. Introduction

Content and Consciousness (hereafter, C&C)¹ outlined an elegant and powerful framework for thinking about the relation between mind and brain and about how science can inform our understanding of the mind. By locating everyday mentalistic explanations at the personal level of whole, environmentally embedded organisms, and related scientific explanations at the subpersonal level of internal informational states and processes, and by judicious reflections on the relations between these levels, Dennett showed how we can avoid the complementary errors of treating mental states as independent of the brain and of projecting mentalistic categories onto the brain. In this way, combining cognitivism with insights from logical behaviourism, we can halt the swinging of the philosophical pendulum between the “ontic bulge” of dualism (C&C, p.5) and the confused or implausible identities posited by some brands of materialism, thereby freeing ourselves to focus on the truly fruitful question of how the brain can perform feats that warrant the ascription of thoughts and experiences to the organism that possesses it.

The main themes of the book are, of course, intentionality and experience – content and consciousness. Dennett has substantially expanded and revised his views on these topics over the years, though without abandoning the foundations laid down in C&C, and his views have been voluminously discussed in the associated literature. In the field of intentionality, the major lessons of C&C have been widely accepted – and there can be no higher praise than to say that claims that seemed radical forty-odd years ago now seem obvious. In the field of consciousness, the philosophical pendulum continues to swing, and Dennett continues to press the case for his position with clarity and wit. One can only hope that in another forty years, these lessons too will seem obvious.

In this chapter I want to turn aside from these major themes to look at a relatively neglected part of C&C which I believe deserves to be better known by both philosophers and scientists. This is Dennett's discussion of thinking and reasoning in Chapter VIII. In this chapter Dennett distinguishes two senses of “thinking”, parallel to the senses of “awareness” distinguished earlier in the book. In retrospect, this distinction anticipates contemporary “dual-process” theories of reasoning, and I shall show how Dennett's distinction might be developed and argue that it offers an attractive reinterpretation of the dual-process approach.

* This is the author's preprint of a chapter forthcoming in C. Muñoz-Suárez and F. De Brigard (eds.) *Content and Consciousness Revisited*, Springer.

¹ References are to the second edition (Dennett, 1986).

The chapter is structured as follows. Section 2 sets the scene by introducing dual-process theories in psychology. Section 3 reviews Dennett's remarks about thinking in C&C and shows how they suggest a version of dual-process theory conceived in terms of the personal/subpersonal distinction. Section 4 fleshes out this theory, drawing on ideas from Dennett's later work, and Section 5 outlines some attractions and implications of this version of dual-process theory. The final section shows how the proposed theory might be extended by combining it with a dual-attitude theory of belief, also inspired by ideas in Dennett's work.

2. Dual-process theories

In recent decades, researchers studying various aspects of human cognition have proposed dual-process theories. Such theories hold that there are two different processing mechanisms available for problem-solving tasks, usually labelled *Type 1* and *Type 2*, which employ different procedures and may yield conflicting results.² Type 1 processes are typically characterized as fast, effortless, automatic, nonconscious, inflexible, heavily contextualized, and undemanding of working memory, and they are usually held to be responsible for biased and stereotypical responding on problem-solving tasks. Type 2 processes, by contrast, are typically described as slow, effortful, controlled, conscious, malleable, abstract, and demanding of working memory, and they are claimed to be the source of our capacity for normative responding in accordance with logical rules. Theories of this kind have been proposed, largely independently, by researchers on reasoning (e.g., Evans, 1989, 2007, 2010; Evans & Over, 1996; Sloman, 1996; Stanovich, 1999, 2011), decision making (e.g., Kahneman, 2011; Kahneman & Frederick, 2002; Reyna, 2004), social cognition (e.g., Chaiken & Trope, 1999; Smith & DeCoster, 2000), and learning and memory (e.g., Dienes & Perner, 1999; Reber, 1993).³

In the field of reasoning and decision making, dual-process theories were originally proposed to explain conflicts between normative and biased responses on experimental tasks. However, the theories have subsequently been supported by a wide range of other evidence, including, (a), experimental manipulations (including explicit instruction) designed to shift the balance between the two types of processing (e.g., De Neys, 2006; Roberts & Newton, 2001), (b), psychometric studies showing that cognitive ability is differentially linked to performance on tasks where Type 2 thinking (which is demanding of resources) is required for production of the normative response (e.g., Stanovich, 1999; Stanovich & West, 2000), and, (c), neuroimaging studies indicating that responses associated with the different types of processing

² As I use the term, dual-process theories contrast with dual-mode theories, which recognize the existence of two styles of reasoning but regard them as different modes of a single mechanism, or type of mechanism.

³ For surveys of the literature, see Evans, 2008; Frankish & Evans, 2009; Frankish, 2010.

activate different brain regions (e.g., De Neys et al., 2008; Lieberman, 2009; McLure et al., 2004). Theorists disagree about the relations between the two processes and about whether they operate in parallel or in sequence. A popular view is that Type 1 processes generate rapid default responses, which usually control behaviour, but which can, given sufficient resources, motivation, and ability, be intervened upon and replaced with more reflective responses generated by slower, Type 2 processes. Evans calls this view *default-interventionism* (Evans, 2007).

Some dual-process theorists have taken a further step and proposed dual-*system* theories, according to which human cognition is composed of two multi-purpose reasoning systems, usually known as *System 1* and *System 2*, the former supporting Type 1 processes, the latter supporting Type 2 ones (e.g., Epstein, 1994; Evans & Over, 1996; Stanovich, 1999, 2004). Dual-system theorists typically claim that System 1 is an evolutionarily old system, whose performance is unrelated to general intelligence, whereas System 2 is a more recent, uniquely human system, whose performance correlates with general intelligence.

Recently, however, some dual-process theorists have shunned the term “system”, with its implications of unity, discreteness, and functional specialization, and reverted to talk of *types* of processing (e.g., Evans, 2010; Stanovich, 2011). It was always understood that System 1 was actually a suite of subsystems, including domain-specific modules, implicit learning mechanisms, emotional subsystems, and associations and responses learned to automaticity. And some theorists now argue that there are a variety of Type 2 systems as well, unified by their processing characteristics and shared use of working memory (Evans, 2009). It has also been increasingly recognized that Type 2 processing requires supporting Type 1 processing of various kinds.⁴ Theorists have also qualified their descriptions of the two types of process, stressing that many of the features commonly assigned to each should be treated as typical correlates rather than necessary, defining characteristics (Evans & Stanovich, 2013). Thus, on this view, Type 1 processes are often, but not invariably, contextualized, fast, and productive of biased responses, and Type 2 processes often, but not invariably, slow, abstract, and productive of normatively correct responses. At the same time, theorists have highlighted just one or two features as defining of Type 2, as opposed to Type 1, processing. For Evans it is use of working memory; for Stanovich it is reflective control and “cognitive decoupling”, i.e., the capacity to entertain hypotheses and run mental simulations.

Although these qualifications soften the hard outlines of dual-system theory, they leave intact the core idea that there are two forms of cognition, one that is evolutionarily old, automatic, guided by instinct and habit, and independent of general intelligence, and another that is distinctively human, controlled, flexible,

⁴ Evans, for example, stresses the role of preattentive Type 1 processes in supplying content to Type 2 processing and highlights the need for control processes that allocate resources to the two systems and resolve conflicts between them (Evans, 2009).

dependent on working memory, and linked to general intelligence. This broad picture is well supported. However, many issues remain, in particular about Type 2 processes. Type 2 processing seems capable of some prodigious intellectual feats. Indeed, it seems to occupy the role of something rather like a central executive, which can override instinctive, associative, and emotional responses with rational thoughts and decisions. Now, the positing of such an executive system is, of course, a move which Dennett opposes, as being both unexplanatory and neurologically implausible – a central theme of *Consciousness Explained* (Dennett 1991). Moreover, there are problems in explaining how Type 2 processing could have evolved. It is often claimed that Type 2 processing is evolutionarily recent and even unique to humans (e.g., Evans 2010), but this means it must have developed in a very short span of time on the evolutionary scale. A third issue concerns consciousness. Type 2 thinking is usually characterized as being conscious, but there are reasons for doubting that conscious thought plays any *distinctive* role in guiding behaviour (as dual-process accounts assume it does). From a neural perspective, consciousness seems to be a late-occurring event in the sequence from perceptual input to behavioural output, with conscious awareness of a decision lagging behind the neural processes that initiate it (e.g., Libet, 2004). This suggests that the putative Type 2 processes may be merely side-effects of, or commentaries on, the non-conscious processes that do the real work in guiding behaviour (e.g., Wegner, 2002) – a view which would be a single-process one.⁵

I think there is a solution to these problems, which is compatible with the data and faithful to the spirit of dual-process theory. However, it requires a certain shift of perspective – a shift that can be motivated by looking at Dennett's views on thinking, beginning with those in C&C.

3. Dennett 1969 on thinking

The discussion of thinking and reasoning in C&C comes in Chapter VIII, after Dennett has argued for a view of consciousness as availability to verbal report. On that view, content becomes conscious in virtue of becoming an input to the subject's speech centre – crossing the “awareness line”. (In the terminology of the book, such contents are the objects of *awareness*₁. Contents that are effective in guiding behaviour but are not input to the speech centre are said to be objects of *awareness*₂; C&C, pp.118-9.) This is, of course, in stark opposition to the Cartesian view of consciousness as an internal arena where mental images are observed and mental

⁵ When I talk of *consciousness* in this chapter I mean *access consciousness* – roughly, availability to other central mental processes and to verbal report. The question is whether, in the case of thought, such access is (at least sometimes) associated with a different mechanism of behavioural control. I am not concerned with issues that arise specifically from the role of *phenomenal consciousness*, the putative subjective qualities of experience. (For the distinction between access and phenomenal consciousness, see Block, 1995.)

operations performed, and Dennett's first task in discussing thinking and reasoning is to reject what he calls the “hammer and tongs” view, on which there are agents and objects in consciousness “[o]ne supposes that there are *conscious acts* of reasoning, acts of judgment and acts using concepts, and on the model of public acts we expect some organ, arm or tool to be *acting on* some object or some raw material – all this within the arena of consciousness” (C&C, p.148).

In developing a better view, Dennett invokes the personal/subpersonal distinction. Reasoning, he maintains, is a personal activity – something done by persons, not by brains. But, most of the time, we have no awareness (i.e. awareness₁) of the processes that give rise to our conclusions and judgements, so this personal activity cannot be an operation or process *by which* a result is derived. Rather, it is more like the *reporting* of a result: “[a]s Ryle points out, such quasi-logical verbs as ‘conclude’, ‘deduce’, ‘judge’ and ‘subsume’ do not refer to processes at all, but are used in the presentation of results already arrived at” (op. cit., p.149). Of course, there must *be* operations involved in the production of inferences and judgements, and these must be guided by stored information, but they will be of a subpersonal kind with no awareness₁. The hammer-and-tongs view results from projecting personal categories onto these subpersonal processes, and it involves a confusion of levels.

Dennett notes that this point may be obscured by the fact that we are often introspectively aware of *some* operations associated with problem solving:

[...] while engaged in problem solving we are aware₁ of a series of things prior to arriving at a conclusion, and we can often, on the basis of this awareness₁, divide our problem solving into a sequence of *operations* or *steps*. [...] When one is asked how one figured out the answer, one can often give a list of steps, e.g., ‘first I divided both sides by two, and then I saw that the left side was a prime ...’. What one is doing when one reports these steps is by no means obvious. (C&C, pp.150-1)

However, Dennett stresses, these operations, too, depend on subpersonal processes to fill in missing steps. We may be able to identify some personal-level operations in a problem-solving episode, but when we ask how we carried out *these* operations, we quickly run up against unanalysable personal activities, which must, nonetheless, be the product of complex subpersonal informational processing (C&C, pp.151-2).

Now, the terms “thinking” and “reasoning” could refer either to subpersonal information processing whose effects are manifest in our intelligent behaviour, or to conscious mental acts involved in problem solving (“awareness₁ of an argument sequence leading to a conclusion” or “something like ‘consciously reasoning with concepts’” as Dennett puts it; C&C, p.155). Dennett notes that both usages have firm roots in everyday speech, and he suggests that the best course is to distinguish two senses of the words, parallel to the different senses of awareness distinguished earlier in the book. In the internal-processing sense animals can think, whereas in the conscious-acts sense they cannot (since Dennett identifies consciousness with awareness₁, which requires language). Moreover, thinking in the conscious-acts sense

can be enthymematic, omitting important premises, whereas subpersonal processing cannot (C&C, p.155-6). Subpersonal processes can fill in missing steps in our conscious reasoning, but those subpersonal processes themselves must draw, in some way or another, on all the information required to reach the conclusion. Dennett does not give names to these two types of thinking, but I shall call them *thinking*₁ and *thinking*₂, the former being subpersonal information processing and the latter a process involving conscious mental operations of some kind. (This numbering unfortunately clashes with Dennett's numbering for awareness, where awareness₁ is the conscious form and awareness₂ the behaviourally manifest kind. However, it harmonizes better with the naming conventions in psychological literature.)

C&C distinguishes two senses of "thinking", then, but this is not yet a dual-process theory. The core feature of dual-process theories is the claim that there exist two different types of reasoning mechanism with different processing characteristics, as opposed to two different modes of a single mechanism. But, for all that has been said so far, *thinking*₁ and *thinking*₂ might be processes of the same type, differing only in that the latter happen to be conscious. That is, the episodes that are characteristic of *thinking*₂ might simply be episodes of *thinking*₁ that cross the awareness line by becoming inputs to the speech system. Such a view would not amount to a dual-process theory.

There are passages in C&C which might support this interpretation. Dennett stresses that in many cases much the same information processing must go on whichever type of thinking precedes an action. Whether we notice an apple and consciously decide to eat it, or just pick up the apple and start munching, in either case our behaviour is guided by stored information about the edibility of apples, the ownership of this apple, the time to the next meal, and so on (C&C, p.153). This does not count decisively against a dual-process view, however, either as an interpretation of C&C or as the best view of the situation. Dennett's main concern in this chapter is not to explore the nature of *thinking*₂, but to oppose the view that it is the only or core form of thinking and to argue that talk of thinking or reasoning is often simply an idealized intentional characterization of subpersonal information processing operations of which we have no conscious awareness. And while he does hold that the conscious events that are distinctive of *thinking*₂ are the product of *thinking*₁, this is not the same as saying that they are simply portions of *thinking*₁ that happen to be conscious.

Moreover, C&C contains hints of a different view. In discussing the role of awareness, Dennett notes that awareness₁ is associated with enhanced behavioural control: awareness₁ is a central component of attention, and attention improves control. However, he points out, awareness₁ *in itself* could not do this; reportability is not logically related to control. Rather, Dennett suggests, awareness₁ may be a contingent (and not invariable) by-product of a prior shift in control elsewhere in the system.

There seem to be two levels from which we direct our behaviour. At the 'high' level (apparently in the cortex) we correlate information from a variety of sources, the behaviour controlled is versatile and changeable – and not particularly coordinated. Once under control, the behaviour is often made into a routine and the control is packed off into a more automatic and specialized system [...] If 'paying attention' is a matter of dealing with the relevant parts of the environment at the high level, it might also *happen* to be a matter of bringing certain high-level signals across the awareness line, just because that is the way the brain is wired. (C&C, p.124)

Dennett notes that such a contingent connection would be adaptive if it supported the practice of verbal instruction.

Combining these remarks with the distinction of types of thinking yields a dual-process picture, which posits two processes with different characteristics and mechanisms: thinking₁, which is non-conscious, effected by specialized subsystems, and supports fluid, unreflective behaviour, and thinking₂, which is typically conscious, is effected by higher-level mechanisms, and supports more flexible but less fluid behaviour. This outline picture harmonizes well with modern versions of dual-process theory.

C&C, then, offers an anticipation of modern dual-process theory. This in itself is interesting – further evidence of the book's far-sightedness and another instance of the independent emergence of dual-process views different fields (Frankish & Evans, 2009). However, the book also hints at something more. Dennett stresses that reasoning is a personal activity (C&C, p.147-9). In the case of thinking₁ this means simply that verbs of thinking offer “fused” personal-level characterizations of cognitive accomplishments produced by subpersonal processes. However, thinking₂ may be a personal activity in a stronger sense. Dennett suggests that the *process itself* involves the performance of a sequence of personal-level actions; the conscious operations or steps involved are intentional actions, like those involved in, say, talking or drawing, which are purposeful, require effort, and can be done well or badly. This view is hinted at earlier in chapter VIII, where Dennett notes that there is a sense of “thinking” that connotes “[...] purposeful and diligent reasoning, as in the sign on the office wall 'Think!'. [...] In some way or other thinking in this sense, or reasoning, is a process, for it takes time, can leave us exhausted, go astray, be difficult, bog down (C&C, p.147). And it is a view that Dennett endorses explicitly in a latter paper, where he writes that conscious propositional thinking “is a personal level activity, an intentional activity, something we *do*. [...] It is not just something that happens in our bodies. When we think thoughts of this sort, we do, it seems, *manipulate* our thoughts, and it can be difficult or easy work” (Dennett, 1998, p.286).

This points to an alternative way of formulating dual-process theory, on which the core distinction is between reasoning processes that are wholly subpersonal (though the judgements and inferences they generate are ascribed to the person) and thinking that constitutively involves performing intentional, personal-level actions of some kind. (This is not to say that that is *all* that thinking of the second type involves;

subpersonal reasoning processes may be involved in *generating* the actions in question. The claim is that what is distinctive of the second type of thinking is that it involves the performance of *some* intentional actions, of the appropriate kind.)

I believe this is a fruitful approach, which harmonizes well with recent dual-process theories while offering solutions to some of the problems they face. The idea is only hinted at in C&C, and Dennett does not suggest what kind of personal activities might be involved in thinking₂. However, he returns to the subject, from a different perspective, in *Consciousness Explained* (Dennett, 1991), and in the next section I shall show how the ideas there can be used to flesh out the proposed personal/subpersonal approach to dual-process theory.

4. Dennett 1991 on the conscious mind

In *Consciousness Explained* Dennett stresses the relative limitations of the human biological brain. Our brains, he points out, are little different from those of our ancestors 150,000 years ago. Fundamentally, they are collections of specialized but unintelligent subsystems, many innately specified, operating in parallel and competing for control of motor systems. They have, in addition, a high level of plasticity, conferring remarkable capacities for individual learning and adaptation, and they are promiscuous information gatherers across a range of sense modalities. But, considered as bare biological organs, human brains are largely driven by environmental stimuli and have little or no capacity for long-term planning or creative thought. The theoretical task, then, is to explain how such organs could support modern human minds, with their much enhanced powers: “[o]nto this substrate nervous system we now want to imagine building a more human mind, with something like a ‘stream of consciousness’ capable of sustaining the sophisticated sorts of ‘trains of thought’ on which human civilization apparently depends” (Dennett, 1991, p.189). Note that this is essentially the same question as that of how the brain came to support Type 2 processing (or System 2), with its capacity for decontextualized and hypothetical thought and higher-level, reflective behavioural control. And, indeed, what Dennett goes on to propose is in effect a version of dual-system theory.

Dennett argues that the conscious mind is too recent a development for it to be an innately specified biological system with a dedicated neural basis. Rather, he suggests, it is a softwired, or “virtual”, system, which we create for ourselves by engaging in various culturally transmitted behaviours (memes or “good tricks”), which in effect reprogram our biological brains. The most important of these behaviours, he argues, is that of self-directed speech: producing, rehearsing, and rearranging sentences in overt or silent soliloquy (when audible, this is usually called “private speech”, and when inaudible, “inner speech”). This stream of self-directed verbalization transforms the activity of the biological brain, causing its parallel, multi-track hardware to simulate the behaviour of a serial, single-track processor (Dennett 1991, Ch.7).

Of course, the idea that thinking is a sort of inner monologue is not a new one. People often liken thinking to talking to oneself, and Ryle explored the idea in some

depth, struggling with the problem of how to characterize the activity and its purpose in an illuminating way (Ryle, 1979). However, Dennett offers a new slant on the role of self-directed speech. The fundamental idea is that such speech has a *self-stimulatory* effect. Self-generated utterances (questions, commands, reminders, and so on) are “heard” and processed like externally produced ones, and may evoke similar responses, with beneficial effects. Dennett highlights several aspects of this.

First, self-directed speech may promote information access among neural subsystems. A self-generated question may prompt a verbal reply, whose content will then be extracted by the speech comprehension system and made available to other neural subsystems, creating a “virtual wire” through which internally isolated subsystems can communicate. In this way, Dennett suggests, the channel of self-directed speech becomes an “open forum” where stored information can be accessed and applied to any problem (1991, pp.194-7, p.278).

Second, it may enhance behavioural control. Self-generated commands, exhortations, encouragements, and reminders can help to foster focused activity and prevent attention being captured by passing stimuli: “[...] when a task is difficult or unpleasant, it requires ‘concentration,’ something ‘we’ accomplish with the help of much self-admonition and various other mnemonic tricks, rehearsals [...], and other self-manipulations” (Dennett, 1991, p.277). Such manipulations, Dennett suggests, achieve their effects by co-ordinating the activities of specialist subsystems: they serve to “[...] adjudicate disputes, smooth out transitions between regimes, and prevent untimely *coups d'etat* by marshalling the ‘right’ forces” (ibid.).

Third, Dennett stresses that self-directed speech facilitates hypothetical thinking and long-term planning (“producing future”). The idea is that self-generated scenarios or proposals may provoke thoughts of their likely consequences, allowing one to assess courses of action in advance of performing them. Saying to oneself, “What if I did this?” may stimulate thoughts and images of the consequences of doing it, generating positive or negative reactions and so allowing one to evaluate the proposed the action.⁶ Dennett suggests that self-directed speech can also assist planning by reinforcing memory. Commenting on one's actions can make it easier to keep track of the progress one has made and recall the strategies one has used and their success, thus helping one to choose wisely in future (1991, p.278).⁷

Once the trick of self-directed speech had been discovered, Dennett argues, it would have been refined by suppression of overt vocalization and disseminated and elaborated through processes of cultural evolution, and a disposition to master it

⁶ Dennett gives an example using private diagram drawing, which he claims can also be used for self-stimulation; 1991, p.197, pp.220-1.

⁷ Carruthers has proposed a similar account of how self-directed speech supports hypothetical thinking, developed within the context of a massively modular view of the mind (Carruthers, 2006, 2009).

might have been coded into the human genome, thanks to the Baldwin effect.⁸ Other forms of internalized self-stimulation also emerged and spread, including the manipulation of visual imagery (“diagramming to oneself”) as a private substitute for diagram drawing (op. cit., p.275).⁹ As a result, Dennett claims, we have become disposed to develop habits of regular inner speech and other forms of self-stimulation, thereby artificially creating a new level of cognitive activity, which is both serial and heavily language-involving. Dennett calls this softwired system the *Joycean machine* (after James Joyce's 1922 novel *Ulysses*, which records its characters' inner monologues).

This, then, offers an account of the intentional actions involved in thinking₂: they are self-stimulations, typically abbreviated, internalized speech acts. Does this mean that we should, after all, accept a version of the hammer-and-tongs view of thinking derided in C&C, on which there are conscious acts of reasoning, performed upon propositional objects? Only in a very weak sense. On the proposed view there are acts and objects in reasoning, but they are acts of the person, not of an inner, quasi-Cartesian agent, and their objects are sentences, or images of sentences, not wordless propositions.

It is worth stressing that to ascribe these acts to the person is not to ascribe them to something *additional* to the set of subpersonal cognitive mechanisms and other biological subsystems that collectively compose the person. The person is not some extra component or feature with new causal powers, and personal activities are performed in virtue of subpersonal processes. (Note how Dennett puts “we” in inverted commas in the quote above.) Explanations in terms of the person derive from adopting a certain interpretive stance toward a biological entity – viewing it as a unified organism, embedded in, and generally well adapted to, its environment, and with global behavioural dispositions and susceptibilities. It is nonetheless important to make the personal/subpersonal distinction, since personal properties and activities may involve the coordinated activity of many separate subpersonal systems, whose

⁸ The idea is that once an individual has discovered a useful behaviour, such as the knack of making a certain tool, selectional pressure will arise for other members of its community to acquire it too. Those who find it easy to learn the behaviour will be selected for over those who find it hard, and, over time, individuals who are predisposed to learn it will come to predominate in the community (e.g., Dennett, 1991, pp.184-7, 2003).

⁹ Note that this assumes that we can intentionally generate sensory imagery. This might involve the mental rehearsal of action, as proposed by Carruthers (Carruthers, 2006, 2009). The idea is that when an action schema is activated, an internal efference copy of it is created, which is used to create a “forward model” of the action. This then generates proprioceptive and other sensory representations of the movements involved, which are used to guide the execution of the action and anticipate its consequences. In mental rehearsal, Carruthers argues, action schemata are activated offline, with the muscle commands suppressed but the efference copies still issued. The sensory images produced are then received by input systems (audition, vision, speech comprehension, etc), and the information they carry globally broadcast to modular subsystems. Where the rehearsed action is an utterance, auditory images (inner speech) are produced and interpreted, and their contents broadcast.

importance and explanatory role is visible only from the personal perspective. Thus, personal-level, self-stimulatory reasoning involves many subsystems that are not involved in wholly subpersonal reasoning, including motor systems, working memory, and sensory systems.¹⁰

It may be objected that Dennett's proposal does not achieve its aim of explaining higher-level thought without positing something like a central executive. Intentional actions are motivated by beliefs and desires – desires to achieve ends and beliefs about the means to achieve them. And this might suggest that before engaging in self-stimulation, we must have determined what cognitive and behavioural effects we want to achieve and worked out how subpersonal systems need to be stimulated in order to achieve them. But if so, then all the real problem-solving work would already have been done, presumably by some executive system, and the stimulation would be merely a mechanism for implementing its conclusions – and a clumsy and inefficient one at that (why not pass on the executive's commands directly through internal channels?). Dennett anticipates this objection, of course. He argues that speech production need not be the product of specific intentions formulated by an executive system (a “Conceptualizer” or “Central Meaner”), which figures out what needs to be said in advance. Rather, he suggests, sophisticated speech acts might be generated through a process of quasi-evolutionary competition between numerous unintelligent micro-systems (a “pandemonium” of “word demons”), vying to produce utterances of varying degrees of sophistication and appropriateness (Dennett, 1991, ch.8). We might ascribe a sophisticated communicative intention to the *speaker*, but it need not correspond to any prior *subpersonal* command. Similarly, self-directed speech acts might be generated pandemonium-style, without antecedent calculation of their structure or likely effects. It is true that, if they are to count as intentional, self-stimulations must be susceptible to some intentional characterization, but this need not be in terms of desires for specific cognitive and behavioural effects and beliefs about how to achieve them. The motivating states might simply be a desire to solve some problem and the instrumental belief that doing *this* (uttering the words that spring to one's lips) may help.

¹⁰ It may be asked whether it is legitimate, within the framework of C&C, to talk of personal-level actions *causally affecting* subpersonal information processing. After all, Dennett repeatedly cautions against confusing the levels throughout the book. It is true that strict causal explanations of subpersonal events will be framed wholly in subpersonal terms, but we can talk loosely of token personal events having subpersonal effects, provided the events in question are identical with subpersonal ones. And although Dennett denies that *some* personal events (pains, for example) can be identified with subpersonal ones (C&C, p.94), he does not issue a blanket ban on personal-subpersonal identifications, and suggests that we proceed on a case-by-case basis (pp. 16-18, 96). (At the extreme we can treat the personal descriptions as fused and identify them with descriptions of global physical state; Dennett, 1987, p.57.) In the case of imagistic self-stimulations it is plausible to think that at least rough identifications can be made with sequences of subpersonal events, perhaps involving the offline activation of motor schemata, and causal explanations mentioning them should be understood as shorthand for more rigorous but less perspicuous explanations couched in such terms.

But could pandemonium processes generate the subtle self-stimulations required to support executive control, abstract problem solving, and hypothetical thinking? Where does the intelligence in these acts come from? There are several points that might be made here. First many self-stimulations, verbal and otherwise, are *not* particularly intelligent. Much self-directed speech consists of comments on what is happening, chance associations, whimsy, free-wheeling speculations, and so on (just like the monologues in Joyce's *Ulysses*). Useful queries, exhortations, and ideas might be simply chance products of this continual stream of commentary.¹¹ Second, acts of self-stimulation often form part of a sequence of such acts (trains of thought). Self-generated speech and other imagery may not only stimulate cognitive and affective responses, but also trigger further acts of self-stimulation, shaped by those responses. For example, imagining a course of action may provoke images of the action's likely consequences, which may then suggest images of further actions, and so on. In this way, cycles of self-stimulation may arise, taking unanticipated and creative directions (Carruthers, 2006, Ch.5). Third, self-stimulation may be guided by knowledge imparted by culture. Cultural processes may disseminate, not only the trick of self-stimulation itself, but specific applications of it to particular problems. Think, for example, of mnemonic rhymes, like that for the number of days in the months. If we know the rhyme, we can literally *tell* ourselves how many days there are in each month, even if we cannot recall the information directly. There are countless other problem-solving routines we can learn, involving inner speech or inner diagramming, which embody logical or mathematical principles or heuristics of various kinds. More broadly, we can also learn ways of enhancing self-stimulation through developing habits of self-questioning, self-commentary, and so on – habits sometimes taught under the heading of “metacognition”. As Dennett stresses, the distinctive power of the Joycean mind is due far more to cultural programming than to the underlying biological hardware.

Combining these elements from C&C and *Consciousness Explained*, we have a dual-process theory, which distinguishes subpersonal informational processes (thinking₁) and processes involving personal-level intentional self-stimulation (thinking₂). In the next section I shall compare this approach with standard dual-process theories, indicate some of its theoretical attractions, and suggest how it might be experimentally evaluated.

5. Connections, attractions, and implications

In outline at least, Dennett's dual-process theory harmonizes well with other recent dual-process theories. The features of subpersonal information processing and

¹¹ Note, too, that self-stimulation may not always be beneficial. There can be negative thinking as well as positive, and habits of harmful self-stimulation may contribute to some psychopathologies, such as anxiety, depression, and obsessive-compulsive disorder.

personal self-stimulation coincide with those typically ascribed to Type 1 and Type 2 processes. Typically, subpersonal processes are fast, effortless, automatic, non-conscious, and inflexible, whereas acts of intentional self-stimulation are slow, effortful, controlled, conscious, and malleable. Subpersonal processes are also likely to display far less individual and cultural variation than processes of personal self-stimulation.

Moreover, identifying Type 2 processes with acts of self-stimulation explains *why* Type 2 processes have the features they do. They are slow because they employ serial channels designed for speech production and comprehension or for other forms of mental rehearsal (Dennett, 1991, p.197). They are controlled and effortful because they are intentional actions that demand attentional resources. They are conscious because (as Dennett notes) they are perceived just like external stimuli: inner speech is processed by the auditory system, inner diagramming by the visual system, and so on (Dennett, 1991, pp.225-6). They are malleable because, like other intentional actions, they are responsive to beliefs about how they should be conducted – that is, about what problem-solving routines are normatively warranted (Carruthers, 2009). They exhibit high individual variation because individuals differ in their attentional resources and self-regulatory dispositions, and they exhibit high cultural variation because different cultures inculcate different self-stimulatory habits and different problem-solving strategies. At the same time, however, these features are not essential or unqualified. For example, some well-practised self-stimulatory routines could be relatively swift and effortless. This again is in line with current thinking by dual-process theorists (e.g., Evans and Stanovich, 2013).

More importantly, self-stimulatory processes exhibit the two core features of Type 2 processes identified by Evans and Stanovich: use of working memory and cognitive decoupling. Self-stimulation draws on working memory because it involves attending to and manipulating sensory imagery. (Indeed, it is arguable that working memory just *is* the set of resources involved in manipulating sensory imagery in the service of self-stimulation; Carruthers, 2006.) And self-stimulation supports cognitive decoupling since inner speech and other self-generated sensory images can represent non-actual scenarios.

Given these similarities, evidence that supports standard dual-process theories, with the features mentioned, also supports the Dennettian version. And, since the subpersonal/personal distinction implies the others, including the putative core ones, it has a claim to be the truly fundamental distinction, from which all the others follow.

It may be objected that this is not a genuine dual-process theory: although Type 2 thinking involves mechanisms of self-stimulation not involved in Type 1 thinking, all the real *inferential* work is done by the subpersonal reasoning processes that first generate and then respond to sensory imagery. Now, it is true that, on the proposed view, Type 2 reasoning constitutively involves passages of Type 1 reasoning (together with the activity of other motor systems, working memory, and sensory systems), but that is not to say that it is simply a *mode* of Type 1 thinking. There are two points to make, concerning the form and content of Type 2 thinking. Concerning form, on the

proposed view, dual-process theory picks out different levels of functional organization, one partially realized in the other, and processes at different levels may have very different formal characteristics. Just as parallel connectionist processes can be implemented on a suitably programmed serial computer, so (the claim is) slow, controlled, and serial Type 2 processes are (partially) implemented in fast, automatic, and parallel Type 1 processes. The facts about implementation do not impugn the reality of the processes implemented. Second, concerning content, the reasoning at the two levels will be directed to subtly different problems. Personal-level processes will be directed to some real-world problem, whereas the supporting subpersonal activity will be devoted to the problem of producing and evaluating sensory imagery relevant to that problem. For example, suppose I have to work out how to fit various differently sized objects inside a box, and that I do this by imagining various possible arrangements to find one that works. Here *I* am thinking about the objects and the box, whereas the underlying subpersonal processes are concerned with what sensory images to produce, the interpretation of these images and the scenarios they represent, and what further images to produce. These processes become relevant to solving the box problem only when viewed as part of an extended process of problem solving involving the box and my purposes concerning it – a perspective that requires a shift to the personal level.¹² (Of course, this is not to say that subpersonal reasoning processes never engage directly with the real-world problems. Much of the time they do just that, guiding behaviour without the occurrence of any self-stimulatory processes. They have the character described only when they are supporting Type 2 reasoning.) I conclude that Dennett's approach qualifies as a genuine form of dual-process theory, albeit an unorthodox one. The approach has, moreover, several theoretical attractions, as I shall now explain.

First, as we have seen, the approach offers an explanation of how Type 2 processing could be implemented without a dedicated executive subsystem. On this view, executive functions are performed by temporary coalitions of specialist subsystems, formed under the influence self-stimulatory habits:

In our brains there is a cobbled-together collection of specialist brain circuits, which, thanks to a family of habits inculcated partly by culture and partly by individual self-exploration, conspire together to produce a more or less orderly, more or less effective, more or less well-designed virtual machine, the *Joycean machine*. By yoking these independently evolved specialist organs together in common cause, and thereby giving their union vastly enhanced powers, this virtual machine, this software of the brain, performs a sort of

¹² It is true (as mentioned earlier) that the results of thinking₁ (the decisions, conclusions, and actions) and the beliefs and desires that explain these results, are ascribed to the person too. Thus in a sense there are two levels of personal activity here, and I solve the box problem *by* solving (in a constitutive sense) the problem of how to stimulate myself in relevant ways. But the inferential operations involved in solving the latter problem are wholly subpersonal.

internal political miracle: It creates a *virtual captain* of the crew, without elevating any one of them to long-term dictatorial power. (Dennett, 1991, p.228)

Putting it another way, on this view, the executive system is nothing smaller than the person – the whole system acting upon itself through self-stimulation.

Second, and relatedly, Dennett's approach explains how Type 2 thinking could have evolved in such a relatively short space of time. The key biological developments were that of a language system together with a capacity for the mental rehearsal of action, and supporting working memory resources. It is plausible to think that all of these evolved independently and for other purposes, and that Type 2 thinking involved their collective exaptation for cognitive purposes (Carruthers, 2006). The other developments required were cultural, not biological, and included the invention, dissemination, and refinement of personal reasoning strategies, perhaps reinforced by the emergence of some innate dispositions to master such strategies, fostered by the Baldwin effect. These developments could have been extremely rapid on the evolutionary scale.

Third, the approach explains how there can be a distinct role for conscious thought, even if consciousness is a late event in neural processing. The key point is that the last event in one cycle of processing can become the first in a new cycle, via the loop of self-stimulation. To forget this, Dennett remarks, “is like forgetting that the end product of apple trees is not apples – it's more apple trees” (1991, p.255). Indeed, understood as imagistic self-stimulations, conscious thoughts become cognitively effective in virtue of the very same fact that makes them conscious: namely that they are received and processed by sensory input systems.

Thus, Dennett's version of dual-process theory harmonizes well with existing ones and has some distinct advantages as well. However, the theory also differs in important ways from standard dual-process accounts, in particular in its implications for the way Type 2 thinking is implemented in the brain, and I shall conclude this section by highlighting some of these differences.

The first concerns the relation between the neural bases for Type 1 and Type 2 thinking. Psychologists tend to think of the two processes as associated with distinct neural structures. However, on Dennett's view, this will not be the case. On this view, Type 2 thinking does not engage distinct neural mechanisms but involves the higher-level exploitation and coordination of Type 1 mechanisms. As Dennett puts it, the installation of the Joycean machine “is determined by myriad microsettings in the plasticity of the brain, which means that its functionally important features are very likely to be invisible to neuroanatomical scrutiny in spite of the extreme salience of the effects” (Dennett, 1991, p.219). It is true that Type 2 thinking will engage some *additional* mechanisms, such as working memory, that are not involved in Type 1 processing, and the theory thus predicts that there will be some salient differences in the patterns of neural activity associated with each type of thinking (as the evidence indicates there are), but many processing resources will be shared between them.

A second consequence concerns language. It is often claimed that Type 2 thinking is linked to language, in that it is directly responsive to verbal instruction in a way that Type 1 thinking is not. However, if Dennett is right, there will be a much stronger link between Type 2 thinking and language. For the resources of the language system, including both language production and comprehension, will be constitutively involved in Type 2 thinking. As Dennett puts it, “a large portion – perhaps even the lion's share – of the activity that takes place in adult human brains is involved in a sort of word processing: speech production and comprehension, and the serial rehearsal and rearrangement of linguistic items, or better, their neural surrogates” (1991, p.225). Of course, Type 2 thinking may involve non-linguistic forms of self-stimulation, too, such as inner diagramming, but the representational powers of language will make it the dominant form. Note, too, that though Dennett talks of speech here, it would be better to say *language*. I assume that Joycean processes could be implemented in sign language, instead of speech, either with overt signing, or covertly, using proprioceptive or visual imagery.

Further consequences follow from the status of Type 2 processes as intentional actions. First, as *actions*, they will involve the activation of brain regions associated with behavioural control, such as the motor and premotor cortex. Second, as *problem-solving* actions, they will draw on metarepresentational resources. As noted earlier, self-stimulatory activities will typically be motivated by a desire to solve some problem and instrumental beliefs about the strategies that may work. That is, engaging in Type 2 thinking involves thinking not only about the first-order problem one faces but also about the meta-problem of how to solve this problem, and it will therefore draw on metarepresentational and metacognitive resources. This meta-level thinking will, of course, usually be of the subpersonal, Type 1 kind.

The consequences mentioned suggest ways in which Dennett's version of dual-process theory may be experimentally tested, using techniques such as deficit studies, neuroimaging, and dual-task methodologies (in which a subject is required to perform two tasks simultaneously, in order to determine if they share processing responses). The fate of the theory will ultimately depend on such investigations. For the present, however, I shall conclude by highlighting a further theoretical attraction of Dennett's approach, which lies in the way it can be extended.

6. From dual processes to dual attitudes

Dual-process theories are often combined with what I shall call *dual-attitude* theories, according to which each type of processing has its own memory system, with a distinct set of propositional attitudes (e.g., Reber, 1993). This view is supported by social-psychological work on persuasion and attitude change, which has led several theorists to distinguish two memory systems: an implicit system, which is non-conscious, automatic, fast-access, and slow-learning, and an explicit system, which is conscious, effortful, slow-access, and fast-learning (e.g., Wilson et al., 2000; Smith and DeCoster, 2000; Smith and Collins, 2009). Again, Dennett's work contains the seeds of

an alternative approach to dual-attitude theory, also rooted in the subpersonal/personal distinction first introduced in C&C. There is no space here to develop the approach in detail, but I shall sketch the outline (for more details, see Frankish 2004).

The key ideas appear in a chapter in *Brainstorms* (Dennett, 1978, pp.300-9), which looks at the operation of changing one's mind. Here Dennett endorses a suggestion by Ronald de Sousa (de Sousa, 1971) that we have two levels of belief: graded, nonverbal belief, which is common to humans and animals, and binary, verbalized belief, which results from assenting to a natural language sentence. De Sousa likens the act of assent to a "*bet on truth alone*, solely determined by epistemic desirabilities" (quoted in Dennett, 1978, p.304). The product of such an epistemic bet, Dennett suggests, is not so much a belief as a state of commitment or ownership. Assenting to a sentence involves metaphorically putting it in a box marked "True" and committing oneself to asserting it in appropriate contexts. Dennett calls these commitments "opinions" and notes that a person's opinions may diverge from their nonverbal beliefs, as manifest in their behaviour. To make up, or change, one's mind about something, Dennett proposes, is to form or revise an opinion. (Dennett adds, however, that not all opinions are the product of deliberation; some sentences are such sure bets that we add them to our collection of opinions without thinking.)

Dennett occasionally invokes the belief/opinion distinction in his later work, but he does not build on it and does not connect it with his account of the Joycean machine. However, it is natural to make such a connection. The key move is to suppose that the commitment involved in opinion formation extends not only to asserting the endorsed sentence in public, but also to holding it true in one's private self-stimulatory activities. This might involve telling oneself that it is true, taking it as a premise when constructing explicit arguments, rejecting sentences that conflict with it, and so on. (We might say that if conscious reasoning is an exploration of a theoretical landscape – a metaphor used by Ryle (2009) – then opinions are signposts we erect along the way.) On this view, a person's opinions will shape the course of their self-stimulatory activities and the cognitive and behavioural effects that result, functioning very much as beliefs are supposed to do. Indeed, on this view, it becomes attractive to redescribe the belief/opinion distinction as one between types of belief, Type 1 and Type 2, the former associated with Type 1 reasoning and the latter with Type 2. We might also identify a parallel Type 2 form of desire, which involves committing oneself to taking a sentence as a statement of a goal and treating it as a fixed point in our self-stimulatory activities.

This, then, offers a Dennettian dual-attitude view to complement the dual-process one. The theory retains the overall characteristics of standard accounts. Type 1 beliefs are formed slowly, through exposure to environmental regularities, but they influence behaviour rapidly and without conscious thought. Type 2 beliefs, on the hand, can be formed rapidly through one-off acts of assent, but they influence action only through slow and effortful self-stimulation. Moreover, they are activated only when the agent

is engaged in self-stimulation, or prompted to engage in it. In contexts where there are no prompts to self-stimulation they remain inert.

However, like the companion dual-process view, this view differs from standard ones in that it is rooted in the subpersonal/personal distinction. Although both types of belief are ascribed to persons, the processes involved in their formation and behavioural manifestation are located at different levels. Type 1 beliefs are formed by subpersonal processes, and they manifest themselves directly in behaviour without intervening personal activity. Type 2 beliefs, on the other hand, are formed through personal acts of assent, and they influence behaviour via personal-level self-stimulation. As with the companion account, this has consequences for the neural basis of the higher-level states. On the proposed view, Type 2 beliefs are commitments, and commitments can be analysed as complexes of beliefs and desires. Simplifying somewhat, to be committed to doing something is to *believe* that one has committed oneself to doing it and to desire to honour one's commitments. So, to have the Type 2 belief that *p* is to believe that one has committed oneself to holding true a sentence with content *p* and to want to honour this commitment.¹³ (A similar analysis holds for Type 2 desire.) These constituting attitudes need not themselves be Type 2 ones, of course, and typically will not be. (And if they are, the attitudes constituting *those* attitudes will surely not be.) Thus, on this view, Type 2 attitudes are ultimately constituted by a set of metarepresentational Type 1 attitudes. The upshot is that, like the Joycean machine, the Type 2 memory system is a softwired, virtual one, realized in Type 1 states. Again, this view has attractions from an evolutionary perspective and means that Type 2 attitudes will not have a separate neural basis.

Dennett's dual-process theory thus naturally extends to give a picture of the human mind as a two-level structure composed of a lower level of subpersonal informational states and processes and a higher "virtual" level of personally constructed mental attitudes and processes, which is constitutively dependent on the former. This is, I suggest, another reason to prefer it.

7. Conclusion

One of Dennett's aims in C&C concerns was, I take it, to correct the error of philosophers who overlooked the existence of subpersonal reasoning processes or mischaracterized them in personal terms. But if the account sketched here is correct, then some contemporary psychologists are making the complementary error of overlooking personal reasoning processes or mischaracterizing them as subpersonal.

¹³ This is argued in detail in Frankish, 2004.

Highlighting this possible error is one of the many salutary tasks C&C still performs in the fifth decade after its publication.¹⁴

References

- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227-87.
- Carruthers, P. (2006). *The architecture of the mind: Massive modularity and the flexibility of thought*. Oxford: Oxford University Press.
- Carruthers, P. (2009). An architecture for dual reasoning. In J. St. B. T. Evans and K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 109-27). Oxford: Oxford University Press.
- Chaiken, S., and Trope, Y. (Eds.). (1999). *Dual-process theories in social psychology*. New York: Guilford Press.
- De Neys, W. (2006). Dual processing in reasoning: Two systems but one reasoner. *Psychological Science*, 17(5), 428-33.
- De Neys, W., Vartanian, O., and Goel, V. (2008). Smarter than we think: When our brains detect that we are biased. *Psychological Science*, 19(5), 483-9.
- De Sousa, R. B. (1971). How to give a piece of your mind: Or, the logic of belief and assent. *The Review of Metaphysics*, 25(1), 52-79.
- Dennett, D. C. (1978). *Brainstorms: Philosophical essays on mind and psychology*. Montpelier, VT: Bradford Books.
- Dennett, D. C. (1986). *Content and consciousness* (2nd ed.). London: Routledge and Kegan Paul. (First edition 1969.)
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1991). *Consciousness explained*. Boston: Little, Brown and Co.
- Dennett, D. C. (1998). Reflections on language and mind. In P. Carruthers and J. Boucher (Eds.), *Language and thought: Interdisciplinary themes* (pp. 284-94). Cambridge: Cambridge University Press.
- Dennett, D. C. (2003). The Baldwin effect: A crane, not a skyhook. In B. H. Weber and D. J. Depew (Eds.), *Evolution and learning: The Baldwin effect reconsidered* (pp.69-79). Cambridge, MA: MIT Press.
- Dienes, Z., and Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22(5), 735-808.
- Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *The American psychologist*, 49(8), 709-24.

¹⁴ I thank the editors of this volume for their helpful comments on an earlier draft of this chapter. Thanks are also due to Jonathan Evans, Eileen Frankish, Liz Irvine, and Maria Kasmirli for their comments and advice.

- Evans, J. St. B. T. (1989). *Bias in human reasoning: Causes and consequences*. Hove: Lawrence Erlbaum Associates.
- Evans, J. St. B. T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. Hove: Psychology Press.
- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-78.
- Evans, J. St. B. T. (2009). How many dual-process theories do we need? One, two, or many? In J. St. B. T. Evans and K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 33-54). Oxford: Oxford University Press.
- Evans, J. St. B. T. (2010). *Thinking twice: Two minds in one brain*. Oxford: Oxford University Press.
- Evans, J. St. B. T., and Over, D. E. (1996). *Rationality and reasoning*. Hove: Psychology Press.
- Evans, J. St. B. T., and Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3).
- Frankish, K. (2004). *Mind and supermind*. Cambridge: Cambridge University Press.
- Frankish, K. (2010). Dual-process and dual-system theories of reasoning. *Philosophy Compass*, 5(10), 914-26.
- Frankish, K., and Evans, J. St. B. T. (2009). The duality of mind: An historical perspective. In J. St. B. T. Evans and K. Frankish (Eds.), (pp. 1-29). Oxford: Oxford University Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Kahneman, D., and Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, and D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). Cambridge: Cambridge University Press.
- Libet B. (2004). *Mind time: The temporal factor in consciousness*. Cambridge, MA: Harvard University Press.
- Lieberman, M. D. (2009). What zombies can't do: A social cognitive neuroscience approach to the irreducibility of reflective consciousness. In J. St. B. T. Evans and K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 293-316). Oxford: Oxford University Press.
- McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695), 503-7.
- Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.
- Reyna, V. F. (2004). How people make decisions that involve risk: A dual-processes approach. *Current Directions in Psychological Science*, 13(2), 60-6.
- Roberts, M. J., and Newton, E. J. (2001). Inspection times, the change task, and the rapid-response selection task. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology*, 54(4), 1031-48.

- Ryle, G. (1979). *On Thinking*. Oxford: Blackwell.
- Ryle, G. (2009). The thinking of thoughts: What is “Le Penseur” doing? In G. Ryle, *Collected Papers Volume 2: Collected Essays 1929-1968* (pp. 494-510). Abingdon: Routledge.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3-22.
- Smith, E. R., and Collins, E. C. (2009). Dual-process models: A social psychological perspective. In J. St. B. T. Evans and K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 197-216). Oxford: Oxford University Press.
- Smith, E. R., and DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4(2), 108-31.
- Stanovich, K. E., and West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, 23(5), 645-726.
- Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Stanovich, K. E. (2004). *The robot’s rebellion: Finding meaning in the age of Darwin*. Chicago: University of Chicago Press.
- Stanovich, K. E. (2011). *Rationality and the reflective mind*. New York: Oxford University Press.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wilson, T. D., Lindsey, S., and Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107(1), 101-26.