

Knowledge – Explicit, implicit and tacit: Philosophical aspects*

MARTIN DAVIES

Abstract Explicit knowledge is knowledge that the knower can make explicit by means of a verbal statement; implicit knowledge is knowledge that is not explicit. Chomskyan tacit knowledge of rules can be understood as requiring a particular structure in the content-involving explanations of transitions between internal representations. States of tacit knowledge are different from beliefs and states of explicit knowledge in that their content is at least partly nonconceptual. Several applications of the notion of tacit knowledge in philosophical theory are reviewed.

Chomsky's notion of tacit knowledge of the rules of a generative grammar provided the theoretical background for research on artificial grammar learning (AGL). Results and disputes in AGL research are reviewed, with particular reference to requirements on testing for conscious *versus* unconscious knowledge. No test is guaranteed to meet the requirements but tests can be validated if they are used in the context of psychological theories about qualitative differences between conscious and unconscious processes.

Explicit knowledge is knowledge that the knower can make explicit by means of a verbal statement: 'Someone has explicit knowledge of something if a statement of it can be elicited from him by suitable enquiry or prompting' (Dummett 1991). Implicit knowledge can then be defined simply as knowledge that is not explicit. On this construal, implicit knowledge corresponds roughly to what Polanyi (1967) calls 'tacit knowing': 'we can know more than we can tell'.

Tacit knowing in Polanyi's sense encompasses a variety of phenomena, such as the ability to recognize something (e.g. a person's face) even though one cannot describe in context-independent terms (e.g. without saying, 'I know that Bob looks like this') how one recognizes it. Polanyi's tacit knowing also includes the holistic diagnostic skills that an experienced and humane psychologist or psychiatrist brings to a clinical encounter. Diagnostic judgment depends on 'content-dependent practical knowledge' (Thornton 2013).

While explicit knowledge can be expressed and communicated linguistically, Polanyi's tacit knowing cannot be captured in language without reliance on context-dependent or demonstrative elements (e.g. I know that to tie my shoelaces I move my fingers *thus* (in this way)'). But tacit knowing can still be communicated and shared. The mechanism of transmission is not telling and understanding, but rather performance and imitation (learning by example). The notion of socially and culturally shared knowledge in the service of solving theoretical and practical problems (perhaps in the form of shared

* I am grateful to Zoltan Dienes, David Lobina and Crysostomos Mantzavinos for their comments and advice.

knowledge about what will happen, and about what to do, in various situations) is of evident importance for the social sciences. Specifically, shared theoretical and practical knowledge can be seen to play a foundational role in the establishment and development of institutions (Mantzavinos 2001, Mantzavinos, North and Shariq 2004).

A different notion of tacit knowledge, introduced by Chomsky (1965), has been of central importance for philosophy and cognitive science over the last half-century, and this will be the focus of the first four sections of this entry. An account of Chomskyan tacit knowledge will be given (section 1) and the relationship between tacit knowledge and other psychological states will be discussed (section 2). Two distinctions that are drawn using the terminology of ‘explicit’ *versus* ‘implicit’ will be explained (section 3): information that is explicitly stored *versus* information that is implicitly represented (implied by explicitly stored information); and rules that are explicitly stored *versus* rules that are implicit in a system (‘wired in’). Four examples of the use of the notion of tacit knowledge in philosophical theory will be reviewed (section 4).

Chomsky’s notion of tacit knowledge provided the theoretical background for empirical research on artificial grammar learning (AGL) and this research programme will be described (section 5). The primary interest of AGL research for this entry is that it makes use of a distinction between conscious and unconscious knowledge – a contemporary rendition of the explicit *versus* implicit knowledge distinction. Theoretical proposals about testing for conscious knowledge will be discussed (section 6) and recent work on conscious (explicit) and unconscious (implicit) knowledge using AGL and related experimental paradigms will be reviewed (section 7).

1. Tacit Knowledge of Rules

In a famous passage, Chomsky (1965) said: ‘Obviously, every speaker of a language has mastered and internalized a generative grammar that expresses his knowledge of his language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them.’ This internalization of the rules or principles of a generative grammar is usually described as tacit knowledge and the claim that each language user tacitly knows a generative grammar has been a core element of Chomskyan linguistics and cognitive science (e.g., Chomsky 1986: ‘the central concern becomes knowledge of language’). Tacit knowledge in this Chomskyan sense counts as implicit according to our initial definition, because a speaker who has tacit knowledge of the rules or principles of a grammar is usually unable to provide a verbal statement of those rules or principles. (A

student of linguistics may be able to state the rules that he tacitly knows. But, even in this case, the tacit knowledge is not simply transformed from implicit to explicit. The linguistics student's state of explicit knowledge will probably be based in part on reading books and listening to lectures and will be somewhat independent from – not constituted or underpinned by, and not a manifestation of – her state of tacit knowledge.)

1.1 Chomsky on the Faculty of Language

We have introduced the notion of tacit knowledge as being Chomskyan. But, before proceeding, it is appropriate to acknowledge that, in more recent presentations (e.g. Chomsky 2000, 2005, Hauser, Chomsky and Fitch 2002), Chomsky describes his project in rather different terms, and from a 'biolinguistic perspective'.

According to these more recent presentations, the object of study is an individual's internal language (I-language), the attained state of the individual's faculty of language. The faculty of language itself is conceived as a component or organ of the mind/brain, with its initial state (studied under the heading 'universal grammar') fixed by genetic endowment. At the heart of the faculty of language are a collection of primitive items (the lexicon or vocabulary) and a procedure that generates complex structures (expressions) from the lexical items. The faculty of language interfaces with sensorimotor (perceptual-articulatory) systems and with conceptual-intentional systems and each expression, E, is a pair <PHON(E), SEM(E)>, in line with the conception of language as a system of sound–meaning pairings.

It seems clear that, on this account, the attained state of the faculty of language is not a state of explicit or tacit knowledge about the I-language; it is the I-language itself. This attained state of the faculty of language is not a state of knowledge about a set of lexical items and about complex expressions built from those items; rather, it is constituted by the lexicon and the generative procedure themselves. And the procedure generates, not representations of complex expressions, but the expressions themselves.

Someone might suggest that the primary function of language is inter-personal communication and that the internally generated expressions of the I-language should be regarded as representations of external tokens of sentences in communicative use. Consequently, the state of the faculty of language should still be regarded as a body of knowledge – knowledge about external language. But Chomsky would utterly reject that suggestion. Following on from some discussion of remarks by François Jacob and Salvador Luria, Chomsky (2005) describes the internally generated expressions as 'the

infinite variety of internal structures that enter into thought, interpretation, planning, and other human mental acts, and that are sometimes put to use in action, including the externalization that is a secondary process if the speculations just reviewed turn out to be correct’.

There is some debate as to whether there has really been a radical change in Chomsky’s substantive position (Collins 2004, Matthews 2006). But, however that may be, many questions arise. For example, even if the attained state of the faculty of language is the I-language, rather than being a body of tacit knowledge about the I-language, subjects still have explicit knowledge about their own languages, and some of this knowledge seems to have a first-person character. If Chomsky is correct, then this first-personal linguistic knowledge cannot be explained in terms of the subject’s drawing on states of tacit knowledge embodied in the faculty of language. So one project for continuing and future research is to provide and evaluate accounts of first-personal linguistic knowledge that do not appeal to tacit knowledge (Smith 2008).

With this acknowledgement of the nature of Chomsky’s recent descriptions of his project, we return to ‘Chomskyan’ tacit knowledge of rules.

1.2 Quine’s Challenge

Quine challenged Chomsky’s introduction of the notion of tacit knowledge by making use of the distinction between behavior that conforms to a rule and behavior that is guided by a rule. As the notion of conformity is used here, conformity to a rule does not require knowledge of the rule. Behavior conforms to a rule if it fits the pattern that the rule requires. As Quine (1972) used the notion of guidance, in contrast, guidance by a rule requires explicit knowledge of that rule: ‘the behavior is not *guided* by the rule unless the behavior knows the rule and can state it’. Thus, a subject can behave in a way that conforms to a rule although his behavior is not guided by the rule.

Chomsky’s tacit knowledge of a rule is supposed to be less demanding than the kind of guidance that requires explicit knowledge of the rule; but it cannot be equated with mere conformity to the rule. In fact, conformity to a rule is neither necessary nor sufficient for tacit knowledge of the rule. Conformity is not necessary for tacit knowledge, since the psychological state of tacit knowledge of a rule does not guarantee effective use of that knowledge to generate behavior that conforms to the rule. In Chomsky’s terms, competence – which is a state of knowledge – must be distinguished from performance. (This distinction between having knowledge and making use of that

knowledge holds good for explicit, as well as for tacit, knowledge. Sometimes – all too often – we possess explicit knowledge that is relevant to the task at hand and yet the knowledge is not put to use in our behavior.)

More importantly for present purposes, conformity to a rule is not sufficient for tacit knowledge, because a tacit knowledge claim is not offered as a summary description of a pattern in behavior but, rather, as a putative explanation of behavior. There will always be alternative sets of rules that require just the same behavior for conformity, and thus provide equally adequate summary descriptions of a battery of conforming behavior. But it is part of the idea of tacit knowledge that a speaker's actual behavior might be correctly explained in terms of tacit knowledge of one set of rules, rather than in terms of tacit knowledge of an alternative set of rules to which it also conforms.

It is at this point that Quine (1972) posed his challenge. He insisted that, if an attribution of tacit knowledge is an empirical claim that goes beyond a summary of conforming behavior, then it should be possible to indicate what kinds of evidence would count in favor of, or against, that attribution. He also insisted that this evidence should involve the subject's own behavior. Quine's demand is challenging because there is no unique set of rules to which a subject's behavior conforms. But there are two points to be made in response. First, there can be no a priori limit on the kinds of evidence that might be relevant to an empirical claim. So it is not legitimate to restrict evidence that could support an attribution of tacit knowledge to the behavior of the very subject to whom the attribution is being made. Second, the more general point about evidence is a fair one. There should be something to be said about what kinds of evidence would support an attribution of tacit knowledge. But more fundamental than the question about evidence is the question what the correctness of such an attribution would consist in – what a state of tacit knowledge is.

1.3 An Account of Tacit Knowledge

We can sketch one answer to this fundamental question by using an example that involves, not rules of syntax, but rules of pronunciation for simple three-letter strings. Suppose that there are five initial consonants, 'B', 'D', 'F', 'H', 'K'; five central vowels, 'A', 'E', 'I', 'O', 'U'; and five final consonants, 'M', 'N', 'P', 'T', 'Z'. Thus, there are 125 three-letter strings. Suppose that each letter is assigned a single phoneme by a letter-sound rule, such as the 'B' → /b/ rule: if a string begins with the letter 'B' then its

pronunciation begins with the phoneme /b/. The fifteen letter–sound rules logically determine pronunciations for all the 125 strings.

Now consider a person reading these letter strings aloud. Behaviorally, the person produces a spoken pronunciation when a letter string is visually presented; for example, produces the pronunciation /bæm/ (rhymes with ‘ham’) when the letter string ‘BAM’ is presented. Suppose that the person’s behavior – vocally producing pronunciations in response to the visual presentation of letter strings – conforms to the fifteen letter–sound rules. We assume that there are psychological processes underpinning this behavior, and that these processes involve transitions from internal representations of presented letter strings to internal representations of to-be-produced pronunciations. These transitions can also be described as conforming to the letter–sound rules: When an input represents a three-letter string, the output represents the pronunciation that is determined by the rules. In this setting, our fundamental question is: What would it be for the person (or, indeed, for a mechanical reading-aloud system) to have (or embody) tacit knowledge of the fifteen letter–sound rules?

The answer to this question, in outline, is that tacit knowledge of the rules requires a particular structure in the causal explanations of the transitions from internal representations of letter strings (orthographic representations) to internal representations of pronunciations (phonological representations). Consider first the abstract structure of the logical determination of pronunciations by the letter–sound rules. Each rule (e.g. the ‘B’ → /b/ rule) makes a common contribution to determining the pronunciations of 25 three-letter strings (e.g. those beginning with ‘B’); and the pronunciation of each letter string is determined by three letter–sound rules. Tacit knowledge of the fifteen letter–sound rules then requires a causal-explanatory structure that mirrors this structure of logical determination. A state of tacit knowledge of a letter–sound rule makes a common contribution to the causal explanations of 25 transitions from orthographic to phonological representations; and the causal explanation of each transition involves three states of tacit knowledge of letter–sound rules. Clearly, conformity to the rules is not sufficient for tacit knowledge of the rules according to this outlined account. For example, mere conformity to the rules does not guarantee that there is any common factor in the causal explanations of the 25 transitions from orthographic representations of letter strings beginning with ‘B’ to phonological representations of their pronunciations (beginning with /b/).

The kind of structure that tacit knowledge requires, in the causal explanations of transitions between psychological states, is not mysterious. Stepping back from internal representations and tacit knowledge for a moment, consider transitions between ordinary conscious beliefs (or states of explicit knowledge) about presented letter strings and beliefs (or states of explicit knowledge) about to-be-produced pronunciations. In particular, consider these transitions in two subjects, Syst and List. Suppose that Syst has explicit knowledge of the fifteen letter–sound rules and is able to draw on this knowledge to work out – that is, to infer – the pronunciations of all 125 three-letter strings. For each letter–sound rule, explicit knowledge of it figures as a common factor in Syst’s inferences about the pronunciations of 25 three-letter strings and, for each of the 125 three-letter strings, Syst’s inference draws on explicit knowledge of three letter–sound rules. Suppose that List, in contrast, lacks explicit knowledge of the letter–sound rules but has 125 separate pieces of explicit knowledge, one for each three-letter string – from explicit knowledge that string #1, ‘BAM’, is pronounced /bæm/ through to explicit knowledge that string #125, ‘KUZ’, is pronounced /kʌz/ (rhymes with ‘fuzz’). We can say that List has explicit knowledge of 125 pronunciation rules, but each of these rules applies to only one three-letter string.

Syst and List assign the same pronunciations to the 125 letter strings, but there are substantial differences in the structure of the causal explanations of their transitions from beliefs about letter strings to beliefs about pronunciations. For example, there is a single psychological state of Syst (namely, explicit knowledge of the ‘B’ → /b/ letter–sound rule) that makes a common contribution to the causal explanation of each of Syst’s beliefs about the pronunciations of letter strings beginning with the letter ‘B’. But there is no such common factor in the causal explanations of List’s beliefs. Similarly, the psychological states that figure in the explanations of Syst’s beliefs about the pronunciations of some letter strings (e.g. that ‘BAM’ is pronounced /bæm/ and that ‘DEN’ is pronounced /den/) are together sufficient to explain Syst’s beliefs about the pronunciations of other letter strings (‘BAN’, ‘BEM, and four others); but nothing similar is true of List.

These differences in causal-explanatory structure between Syst and List may be manifested in many ways. For example, List may lose (e.g. by forgetting) his explicit knowledge about the pronunciation of just one letter string (e.g. ‘BAM’), while retaining his knowledge about the pronunciations of the other 124 strings. But if Syst loses some of

his relevant explicit knowledge and no longer believes that the letter string ‘BAM’ is pronounced /bæm/, then he will probably also lose his beliefs about several other letter strings (all strings with ‘B’ as initial consonant, or all strings with ‘A’ as central vowel, or all strings with ‘M’ as final consonant). Similarly, List may revise his beliefs about the pronunciations of individual letter strings in a piecemeal way; but Syst is likely to revise his beliefs about the pronunciations of letter strings in a systematic way.

We can use these descriptions of the role of explicit knowledge of rules in Syst’s and List’s transitions between conscious beliefs to clarify the kind of causal-explanatory structure that tacit knowledge of rules requires. A state of explicit knowledge (e.g. of a letter–sound rule, in the case of Syst) is a state with a particular content and it makes a common contribution to the content-involving causal explanations of many inferential transitions between beliefs; the transitions are explained as transitions between beliefs with their particular contents. In a similar way, a state of tacit knowledge is a state with a particular content and it makes a common contribution to the content-involving causal explanations of many inference-like transitions between representations; the transitions are explained as transitions between states with particular representational properties (Peacocke, 2008). These representations might not be beliefs but, rather, representations or states of information of the kind that figure in information-processing psychology. Indeed, states of tacit knowledge are often states of information-processing systems within a subject – systems whose internal operations are not normally accessible to the subject’s consciousness.

In general, a state of tacit knowledge is a state that figures as a common factor in content-involving causal explanations of certain transitions between representations or states of information. Where different sets of rules require the same transitions for conformity (the case that concerned Quine), the correct attribution to a subject of tacit knowledge of one set of rules rather than another requires a particular structure in the causal explanations of transitions between representations within the subject. For example, correct attribution of tacit knowledge of the fifteen letter–sound rules rather than the 125 one-case pronunciation rules requires the Syst-like causal-explanatory structure rather than the List-like structure (Davies 1987, Evans 1981, Peacocke 1989).

Once an account of tacit knowledge in terms of causal-explanatory structure is in place, it is a relatively straightforward matter to give examples of empirical evidence that would confirm the attribution to a subject of tacit knowledge of a particular set of rules, such as rules of a language. Indeed, some of this evidence meets Quine’s additional

requirement of concerning the behavior of the subject to whom the attribution is being made. Relevant behavioral evidence could come from experimental studies of language acquisition, language perception, and language impairment following brain damage, while further evidence would be available from neural imaging.

2. Tacit Knowledge and Other Psychological States

While it is fairly natural to describe tacit knowledge of rules (such as syntactic rules, in Chomsky's example) as a case of 'knowing more than we can tell', it is also possible to query the appropriateness of the use of the term 'knowledge' here. Knowledge is usually taken to require belief and to be connected with reliability of the method by which the belief is produced and with notions of justification, warrant, or entitlement to believe. Information-processing mechanisms whose operation is not accessible to consciousness may well meet requirements of reliability and subjects may have an epistemic entitlement to rely on such mechanisms in forming beliefs (Burge 2003). A belief in whose production states of tacit knowledge are implicated may be justified or warranted. But states of tacit knowledge themselves fall outside the scope of epistemic justification or warrant because they are not beliefs.

2.1 Tacit Knowledge, Beliefs, and Nonconceptual Content

Several considerations may be offered in support of the claim that states of tacit knowledge are different from beliefs. One is that a subject is usually unaware of his or her own states of tacit knowledge and, indeed, unable to become aware of them. But this, by itself, is not clearly decisive because it seems possible to make sense of the notion of an unconscious belief – a belief of which one is not aware, and perhaps even a belief of which one cannot become aware. Second, states of tacit knowledge are limited in their causal role and influence. It is characteristic of beliefs that a belief may figure in many different ways in a subject's thinking, reasoning and planning, depending on the subject's other beliefs and desires. A belief is 'at the service of many distinct projects' (Evans 1981). States of tacit knowledge, in contrast, are not inferentially integrated with the subject's beliefs (Stich 1978); they are inferentially isolated, often with causal roles that are limited to an encapsulated and domain specific information-processing system (a module; Fodor 1983). However, this again may not be quite decisive, because we can make sense of the ideas of fragmented rather than integrated minds, and of beliefs that are compartmentalized, guiding behavior in only some – perhaps very few – contexts (Egan

2008). Some monothematic and circumscribed delusions may be beliefs of this kind (Davies, Coltheart, Langdon and Breen 2001).

A third consideration supporting a distinction between states of tacit knowledge and beliefs is that believing requires the possession of concepts. No one can believe that emus do not fly without having some concept of emus and of flying (and negation). But, in general, the subject who has tacit knowledge of a rule does not need to grasp or possess the concepts that a theorist would use to specify the rule that is tacitly known. This is particularly clear in the case of tacit knowledge of rules of language (e.g. rules of syntax), where every normal speaker has tacit knowledge of rules but only people who have studied linguistics grasp all of the concepts that are needed to specify the rules. Thus, states of tacit knowledge are psychological states with nonconceptual content (or, at least, with content that is partly nonconceptual). The subject of the state does not need to possess (still less, to deploy) all the concepts that would be required to specify the content of the state (Davies 1989). Beliefs, in contrast, are the paradigm example of psychological states with conceptual content.

2.2 Critiques of Tacit Knowledge

Quine is certainly not the only philosopher to have adopted a critical view of the notion of tacit knowledge as it figures in linguistics, psychology, and cognitive science.

Some philosophers have argued that Chomsky's claims about tacit knowledge involve a conceptual confusion because the notion of a rule of language belongs with the idea of a normative practice in which people advert to rules to justify, criticize or excuse their actions (Baker and Hacker 1984). But the notion of a tacitly known rule should be distinguished from the notion of a rule that figures in a normative practice. Once the distinction is made, there remains an important question about the relationship between the rules that are tacitly known by individual speakers and the normative practices in which those speakers participate. But that question does not indicate any conceptual incoherence in the very idea of tacit knowledge.

More recently, Bennett and Hacker (2003) have argued that the information-processing paradigm that is shared by cognitive neuroscience and cognitive science rests on a philosophical mistake. They say that psychological powers (e.g. perceiving, thinking, feeling) that are properly attributed to whole human beings should not be attributed to parts of human beings. According to this critique, neither the brain (in cognitive neuroscience) nor an information-processing module (in cognitive science) is

an appropriate subject of psychological predicates. Here, we should acknowledge that, in the sciences of the mind, systems smaller than a whole human being are credited with psychological states and processes such as inference-like transitions between representations, mediated by states of tacit knowledge. But it is not clear that this is a philosophical mistake, for these psychological attributions can surely be distinguished from attributions of perceptual experiences, thoughts, and feelings.

It has also been argued that the ideas of tacit knowledge of rules, and of unconscious information processing in general, are problematic because of a deep connection between intentionality (the representational nature of thoughts and beliefs) and consciousness. This critique culminates in the claim that ‘there are brute, blind, neurophysiological processes and there is consciousness; but there is nothing else’ (Searle 1990). Specifically, according to this critique, there is no unconscious information processing and there are no states of tacit knowledge that are inaccessible to consciousness. But, once again, a response to the critique can begin from a distinction. Even if we accept that there is an important connection between intentionality and consciousness at the level of the whole human being, states of tacit knowledge and representations within cognitive modules are not beliefs, and a philosophical account of their representational nature may be quite different from an account of the intentionality of beliefs.

There remain questions about the relationship between the psychological properties of whole conscious human beings – which belong at the folk-psychological or personal level of description – and the information-processing properties of component systems of the mind and brain – which belong at the subpersonal level of cognitive and neural mechanisms (Dennett 1969). But there is no reason to assume that answering these questions in a way that highlights what is distinctive about the mental lives of conscious persons will require a radically negative assessment of the cognitive sciences (Davies 2000).

3. Representation and Rules: Explicit and Implicit

Our starting point was that explicit knowledge is knowledge that the knower can make explicit by means of a verbal statement. Thus, explicitness is a matter of the subject being able to present information in linguistic form; it is not a matter of how the information is stored in the subject’s mind or brain when it is not being called upon.

Suppose that someone knows, in the everyday sense of that term, the axioms of some theory. Provided that this knowledge can be verbally articulated, it counts as explicit

knowledge. Now consider some theorem that is derivable from those axioms. The person who knows the axioms may well, with some suitable enquiry and prompting, be able to see that the theorem follows from the axioms and to state it verbally. For example, someone who explicitly knows some elementary arithmetical facts may be able to work out, and to state, that $68+57=125$. On Dummett's (1991) account, this latter piece of knowledge, even though the sum is computed when needed rather than being stored in memory, counts as no less explicit than the stored elementary facts from which it is derived. Indeed, it counts as explicit knowledge even before the knower works it out; for the personal-level notion of explicit knowledge is defined in terms of the possibility of eliciting a verbal statement by enquiry or prompting. Thus, once some propositions are classified as explicitly known, the category of explicit knowledge also includes at least some of the as-yet-undrawn consequences of those propositions.

However, the terms 'explicit' and 'implicit' are also used to mark a number of distinctions that are specifically related to the subpersonal-level storage and processing of information. Two such distinctions are explained in this section.

3.1 Inference and Implicit Representation

Dennett (1983) introduced a quite different distinction from Dummett's when he said: 'Let us have it that for information to be represented *implicitly*, we shall mean that it is *implied* logically by something that is stored explicitly.' On this usage, the as-yet-undrawn consequences of propositions that count as explicitly stored would be classified as implicitly represented. In fact, relative to any given notion of explicit storage of information, it is possible to define a whole family of notions of implicit representation, differing over the inferential resources that can be used in drawing out consequences from the information that is explicitly stored.

There is no uniquely correct account of explicit storage but one natural idea is that explicit storage involves the use of a language-like code (a subpersonal-level echo of the requirement for explicit knowledge). A second idea that is fairly natural in the context of theories of information storage and processing is that explicitly stored information has to be accessed before it can be used. A problem in accessing stored information may result in a failure to use information that is relevant to the task at hand. Putting the two ideas together, we have a picture in which explicit storage of information is a matter of syntactically structured inscriptions on pages that are not always easy to locate. Drawing out the consequences of explicitly stored information could then be conceived as

manipulating inscriptions in ways that are analogous to the formal procedures involved in the derivation of theorems from axioms in a logical system (the computational model of the mind; Fodor 1987).

Beginning from this notion of explicit storage of information, we could consider information that can be arrived at by deductive inference from information that is explicitly stored. If elementary arithmetical facts were explicitly stored in a person's memory then the information that $68+57=125$ would count as being implicitly represented. Thus, the subpersonal-level contrast between explicit storage and implicit representation could distinguish between two cases that would both count as explicit knowledge at the personal level. This deductive notion of implicit representation could be refined by distinguishing, for example, between information that can be arrived at from the explicitly stored information by unrestricted deductive inference, by deductive inference subject to some set of human limitations, or by deductive inference that is within the compass of some specified individual. Alternatively, we could go beyond deductive inference and allow methods of induction (or rule extraction) and abduction (inference to the best explanation). Then, if only facts about particulars were explicitly stored, generalizations or rules that subsumed those facts would count as implicitly represented.

3.2 Rule-Explicit and Rule-Implicit Systems

The notions of explicit storage and implicit representation belong with mechanistic theories about the storage and processing of information. These notions could be applied to any information-processing system and they have no special connection with human subjects. Even in the case of human subjects, we have already seen that subpersonal-level explicit storage is not necessary for personal-level explicit knowledge. Explicitly known information might be either explicitly stored or implicitly represented. It is also the case that explicit storage – even explicit storage within a human subject – is not sufficient for explicit knowledge. Information might be explicitly stored within a human subject yet not available to the subject for verbal report, not accessible to the subject's consciousness, and not conceptualized by the subject.

Tacit knowledge of a rule may involve explicit storage of this kind. An internal representation encodes the rule in a syntactically structured format and a process of searching for and accessing that representation is required before tacit knowledge of the rule can contribute causally to any transitions between input and output representations.

For example, letter–sound rules could be explicitly stored within an information-processing system for reading aloud words and other letter strings aloud (see the implemented computational model of reading aloud described by Coltheart, Rastle, Perry, Langdon and Ziegler 2001). In such a case, the system is said to be rule-explicit, and there is a clear gap between having tacit knowledge of the rules (that is, having the rules represented in the system) and using that tacit knowledge.

While the account of tacit knowledge in terms of causal-explanatory structure allows for tacit knowledge that is a matter of explicit storage, it certainly does not require that tacitly known rules should be explicitly stored. A state of tacit knowledge is a state that figures as a common factor in causal explanations of transitions between representations. The use of an explicitly stored rule can figure as such a causal common factor, but so can a component processor (a ‘transition mediator’; Peacocke 1989) within an information-processing system. The transition corresponding to the rule can be ‘wired in’ to the system, which is then described as rule-implicit (Fodor and Pylyshyn 1988).

The distinction between rule-explicit and rule-implicit systems is quite different from the distinction (in the previous subsection) between explicit storage and implicit representation of information. Rule explicitness is a matter of explicit storage but rule implicitness (that is, a rule that is ‘wired in’ to the system) is not a matter of implicitly represented information (that is, information that is implied by what is explicitly stored). Partial analogies for both contrasts can be found by considering a formal theory in which fifteen proper (that is, not purely logical) axioms state the letter–sound rules, and theorems specifying the pronunciations of the 125 three-letter strings can be derived. If the axioms are analogous to what is explicitly stored in an information-processing system then the theorems that are implied by the axioms are analogous to what is implicitly represented.

Each proper axiom makes a common contribution to the proofs of 25 theorems specifying the pronunciations of letter strings and, in the proof of each such theorem, three proper axioms figure. This structure of derivational common factors would be preserved if the theory were formulated, not with fifteen proper axioms, but instead with fifteen proper rules of inference (e.g. From: String X begins with letter ‘B’, infer: The pronunciation of X begins with the phoneme /b/). Derivationally, a rule of inference is a transition mediator and the distinction between proper axioms and proper rules of inference in a formal theory provides a partial analogy for the distinction between

explicitly stored rules and rules that are implicit in ('wired in' to) an information-processing system.

4. Tacit Knowledge in Philosophy

The starting point for this entry was Dummett's (1991) definition of explicit knowledge as knowledge that the knower can make explicit by means of a verbal statement. Explicit knowledge was contrasted with implicit knowledge, conceived as Polanyi's tacit knowing, and with Chomsky's tacit knowledge. We have given an account of tacit knowledge as having a characteristic causal-explanatory role: a state of tacit knowledge makes a common contribution to content-involving causal explanations of transitions between representations. States of tacit knowledge are different from beliefs and states of explicit knowledge in that the subject of a state of tacit knowledge does not need to possess all the concepts that would be required to specify the content of the state. We have also explained the subpersonal-level distinctions between information that is explicitly stored and information that is implicitly represented and between a tacitly known rule that is explicitly stored and a tacitly known rule that is embodied in a component processor (transition mediator).

In this section, we review four cases where the notion of tacit knowledge figures in philosophical theory.

4.1 Tacit Knowledge of Compositional Semantic Theories

Philosophers of language have an interest in compositional semantic theories for natural languages; that is, theories that reveal how the meanings of complete sentences depend on the meanings of their constituent words and the ways in which those words are put together. In such theories, proper axioms specify the meanings of words and basic modes of combination of words, and theorems that are derivable from those axioms specify the meanings of complete sentences. While it is plausible that ordinary speakers of a natural language have, or can come to have, explicit knowledge of some of these theorems, it is also plausible that ordinary speakers do not need to have, and usually do not have, explicit knowledge of all the axioms of a semantic theory for their language. What, then, is the relationship between an ordinary language user and the axioms of a semantic theory for her language?

One answer to this question is that the epistemological relationship between the language user and the theory extends no further than the theorems that specify meanings

for complete sentences, and that the role of the axioms is merely descriptive. If a semantic theory has just a finite number of proper axioms (one for each word in the language and one for each basic mode of combination), from which meaning specifications for complete sentences can be derived, then those axioms provide a storable summary description of the language user's capacity to understand complete sentences.

A second answer makes use of the notion of explicit knowledge of the axioms, but in an indirect way. Davidson (1984) says that a condition of adequacy for a semantic theory is that what it states should be 'something knowledge of which would suffice for interpreting utterances of speakers of the language to which it applies'. Clearly, if explicit knowledge of the theorems that specify meanings for complete sentences would suffice for interpreting utterances of speakers, then so would explicit knowledge of the axioms from which those theorems can be derived.

Because the explicit knowledge that would suffice for interpretation or understanding is not knowledge that ordinary speakers have, this second answer shares an important property with the first answer. Neither answer proposes a relationship between ordinary speakers and the axioms of a compositional semantic theory that would provide an explanation of the fact that ordinary users of a language are able to understand sentences that they have never encountered before. An explanation of this fact is available, however, if language users have tacit knowledge of the axioms of a compositional semantic theory for their language. If a hitherto unencountered sentence is built out of familiar words in familiar ways then understanding of the new sentence can draw on the language user's tacit knowledge of meaning specifying axioms for those words and modes of combination (Evans 1981, Davies 1987).

The appeal to tacit knowledge here has an additional theoretical advantage of assigning a determinate meaning even to sentences that will never be used. An alternative solution to the 'meaning without use' problem appeals, not to tacit knowledge, but simply to extrapolation from the used sentences to unused sentences: 'Use determines some meanings, those meanings determine the rules, and the rules determine the rest of the meanings' (Lewis 1992). However, this alternative solution to the problem has an unattractive consequence.

Suppose that a speaker PB uses (or equally, suppose that a population of speakers use) a finite set S of sentences from a language (for which there is a compositional semantic theory), but that PB has learned the meanings of these sentences by dipping into

a phrase book. Let us stipulate that PB treats each sentence in S as an unstructured string of letters or sounds, the meaning of which he needs to memorize. Consequently, PB has a separate piece of tacit or explicit knowledge about the meaning of each sentence in the set. We can say that PB has tacit or explicit knowledge of a non-compositional semantic theory with an axiom for each sentence in S. But PB's tacit or explicit knowledge about meanings is not adequate for understanding any sentence beyond the set S.

A semantic theorist could fit a compositional semantic theory to the sentences in S and extrapolate to sentences outside that set. But it is plausible that sentences outside S do not belong to the phrase-book 'language' that PB has learned, and that sentences outside S have no determinate meaning for PB – no determinate meaning in PB's language. Consequently, the appeal to tacit knowledge is to be preferred over the appeal to extrapolation (Schiffer 1993, 2006). As Loar (1981) says: 'the Chomskyan idea of the internalization of the generative procedures of a grammar has got to be invoked to ... make sense of literal meaning'.

Near the beginning of this entry (section 1.1), we noted that the notion of tacit knowledge does not figure in Chomsky's recent descriptions of his project (Chomsky 2000, 2005, Hauser, Chomsky and Fitch 2002). Chomsky's account of the attained state of the faculty of language does not absolutely exclude the possibility that the subject has tacit knowledge of linguistic rules or principles – perhaps, tacit knowledge of a compositional semantic theory for her own language – embodied in some component of the mind/brain other than the faculty of language. But it is a significant question for future research whether appeal to a generative procedure can do the same work in philosophy of language as the appeal to tacit knowledge of a compositional semantic theory.

A more general question in the same area is whether Chomsky's distinction between competence and performance can be retained if the attained state of the faculty of language is not to be conceived as a body of knowledge. Collins (2007) argues that the distinction does not depend on an epistemological – or even on a representational – understanding of competence. Linguistic cognition – and thence, linguistic behavior – depends on many component systems of the mind/brain. As Hauser, Chomsky and Fitch (2002) say: 'a neuroscientist might ask: What components of the human nervous system are recruited in the use of language in its broadest sense? Because any aspect of cognition appears to be, at least in principle, accessible to language, the broadest answer to this question is, probably, "most of it" ... this conception is too broad to be of much use.'

Thus, reflecting normal scientific practice, the study of the faculty of language focuses on one component of the mind/brain ‘under idealization of its integration with other systems’. The notion of competence ‘is intended to abstract one system [the faculty of language] out of an ensemble of systems whose integration underlies performance’ (Collins 2007).

4.2 Tacit Knowledge and Concept Possession

The notion of tacit knowledge also figures in philosophical theories about possession of concepts. One possible proposal in this area would be that possession of some concepts – particularly, concepts of logical operations such as conjunction and disjunction (alternation), but not only those – consists in tacit knowledge of logical or proper rules of inference that codify the inferential potential that is lodged in those concepts. Another possible proposal would be that possession of some concepts consists in tacit knowledge of definitions of those concepts.

Peacocke (2008) calls a state of tacit knowledge that is required for possession of a concept an ‘implicit conception’ and he proposes that, at least in some cases, an implicit conception is or includes a state of tacit knowledge of the contribution to truth conditions that is made by the concept. For example, in the case of the logical concept of alternation, expressed by ‘or’, this would not be tacit knowledge of rules of inference, but tacit knowledge that any proposition (thought content) of the form ‘*A or B*’ is true if and only if either *A* is true or *B* is true. Such a state of tacit knowledge makes a common contribution to the explanations of the subject’s evaluation for truth of many specific thought contents involving the concept (in the case of the example, a common contribution to the subject’s truth-evaluation of many specific contents of the ‘*A or B*’ form). The state of tacit knowledge of truth conditions can also contribute to the explanation of the subject’s evaluation of a principle or rule of inference for validity. This proposal helps to explain the fact that even a primitive (non-derived) principle or rule of inference for a concept is susceptible of rational justification, a kind of justification that draws on the subject’s grasp of the concept in question.

Here it is important to distinguish between principles personal-level acceptance of which can be justified by drawing on one’s grasp of a concept and principles personal-level acceptance of which is necessary for possession of the concept. Many principles of the first kind are not of the second kind and, in fact, some philosophers doubt that there are any principles at all of the second kind (Williamson 2006). Suppose, though, that we

were to allow (as conceptual-role theories of concept possession maintain) that, for a given concept, there are principles or rules of inference such that personal-level acceptance of, or commitment to, those principles or rules is necessary for possession of the concept. Peacocke (2008) argues that this could not provide a complete account of concept possession.

His argument appeals to ‘the phenomenon of new principles’. There are principles involving the concept in question, which have two properties. First, the principles are new, in the sense that they do not follow from the principles or rules already mentioned in the putative conceptual-role account of possession of the concept. Second, a subject can rationally justify acceptance of the principles by drawing on his grasp of the concept. These properties are together problematic for the conceptual role account because, according to that account, possession of the concept consists in personal-level acceptance of or commitment to principles or rules from which the new principles do not follow. But rational justification of new principles has a ready explanation if concept possession consists (in whole or in part) in tacit knowledge of the concept’s contribution to truth conditions. Peacocke (2008) draws a parallel between explaining rational justification of new principles and explaining understanding of new sentences: ‘I am inclined to think that the Phenomenon of New Principles is as decisive an argument against personal-level conceptual role theories as the phenomenon of understanding sentences one has never encountered before is decisive against theories of meaning which do not proceed compositionally.’

4.3 Implicit Rules in Connectionist Networks

A third point at which the notion of tacit knowledge figures in philosophical theory concerns the information processing that takes place in connectionist networks. A network consists of units, organized in layers, and weighted connections between units in successive layers. Input and output representations are patterns of activation across units in the input and output layers respectively. Processing in a network is a matter of activation being passed from the input units to the output units, usually via a layer of hidden (i.e. neither input nor output) units. The activation of units in each successive layer is determined by the activation of units in the previous layer and the weights on the connections between the two layers. Thus, the precise weights on the connections are crucial for the network’s performance of its task (e.g. generating an output representation of the correct pronunciation, given an input representation of a letter string). In this sense,

a network's 'knowledge' about its task domain is embodied in the weights on its connections.

In a pioneering contribution to the development of connectionist models of cognitive processes, Rumelhart and McClelland (1986) highlighted a difference between connectionist models and earlier models by suggesting that connectionist networks 'may provide a mechanism sufficient to capture lawful behavior, without requiring the postulation of explicit but inaccessible rules'. In another influential paper, McClelland, Rumelhart and Hinton (1986) said, 'Using knowledge in processing is no longer a matter of finding the relevant information in memory and bringing it to bear; it is part and parcel of the processing itself'. These (and many similar) passages draw attention to significant features of connectionist networks. But it is important to see that they do not, by themselves, reveal a difference between processing in networks and processing that is guided by tacitly known rules.

Earlier (section 3.2), we distinguished between rule-explicit and rule-implicit systems. Tacit knowledge of a rule may involve explicit storage: a representation that encodes the rule must be searched for and accessed before the tacit knowledge can be used. We also noted earlier that explicit storage is not sufficient for explicit knowledge. Tacit knowledge of rules may involve explicit storage of information that is not available to the subject for verbal report and not accessible to the subject's consciousness. This kind of tacit knowledge ('finding the relevant information in memory and bringing it to bear'; 'explicit but inaccessible rules') can, indeed, be contrasted with the way in which knowledge about a task domain is embodied in a connectionist network.

But a rule may be both tacitly known and 'wired in' to a system, because a component processor, no less than an explicitly stored representation, can make a common contribution to the content-involving explanations of many transitions between representations. Peacocke (2008) makes the same point about implicit conceptions. At the subpersonal level, the content of an implicit conception might be 'explicitly formulated in a language of thought' or 'grounded in the operation of a processor'. When a tacitly known rule is implicit or 'wired in' it is well described as 'part and parcel of the processing'. Thus, the connectionist architecture of units and activations, connections and weights, does not exclude the possibility that tacit knowledge of rules may be embodied in networks.

If we consider again the very simple example of transitions from representations of three-letter strings to representations of their pronunciations, it is certainly possible to

describe a connectionist network in which tacit knowledge of each of the 15 letter–sound rules is embodied in a particular set of weighted connections, so that the causal-explanatory structure is like that in Syst. And it is possible to describe a different network in which tacit knowledge of each of the 125 one-case pronunciation rules is embodied in a particular set of weighted connections, so that the causal-explanatory structure is like that in List (Davies 1995). With more realistic tasks (e.g. when a proportion of the letter strings have irregular or exception pronunciations), and with networks that are more typical of contemporary connectionist research (e.g. Plaut, McClelland, Seidenberg and Patterson 1996), it is usually not possible to identify component processors (sets of weighted connections) that embody letter–sound rules. Nevertheless, it may be possible to find theoretically significant commonalities in the processing of similar items, and the pattern of commonalities may be preserved across different networks (e.g. with different numbers of units and connections) that perform the same task (Shea 2007).

4.4 Tacit Theory and Simulation in Everyday Psychological Understanding

A fourth example is provided by the debate between the theory theory and mental simulation approaches to everyday psychological understanding. On both sides of the debate, it is agreed that we have a capacity to predict, interpret, and explain the mental lives of other people. What is the basis of this capacity? According to the theory theory, we draw on a rich body of knowledge about psychological matters – a theory, in a suitably inclusive use of that term. According to the simulation theory, in apparent contrast, we identify with the other person in imagination and replicate or re-enact aspects of the other person’s mental life. If the other person is thinking about interest rates and the stock market then we think about those same financial matters (we ‘co-cognize’ with the other person) rather than drawing on a rich body of psychological principles concerned with how people think about financial matters (Heal 1986, 1998).

The theory theory faces objections if its claim is that ordinary people with a normal capacity for everyday psychological understanding have a rich body of explicit knowledge about psychological matters. Ordinary people are generally not able to articulate many principles of psychological theory: ‘[W]hy, one wonders, should it be so difficult to articulate laws if we appeal to them all the time in our interpretive practice?’ (Goldman 1989). The obvious response on behalf of the theory theory is to say that the rich body of knowledge is tacit knowledge rather than explicit knowledge. Thus, Stich and Nichols (1992) draw attention to the ‘dominant explanatory strategy’ in cognitive

science, which is to appeal to ‘an internally represented “knowledge structure” ... which serves to guide the execution of the capacity to be explained’. On their account, the theory theory says that everyday psychological understanding is explained by the presence of a ‘largely tacit psychological theory’.

A concern about the shape of the debate may then arise. The theory theory and the simulation theory were supposed to be competing accounts of the basis of our capacity for everyday psychological understanding. But, depending on what it takes for a theory to be tacitly known, might not the simulation theory turn out to be consistent with – even a version of – the tacit theory theory? Blackburn (1992) gives expression to this concern that the tacit theory theory may be so undemanding as to subsume the supposed simulation alternative: ‘If we are good at something ... then we can be thought of as making tacit (very tacit) use of some set of principles that could, in principle, provide a description of a device ... that is also good at it.’

The concern would be justified if tacit knowledge were to be conceived in the way that Blackburn envisages. The debate between the tacit theory theory and the mental simulation theory would collapse if the attribution of tacit knowledge of principles to a person with a particular capacity required no more than that explicit knowledge of the principles should suffice for having the capacity, or no more than that behavioral manifestations of the capacity should conform to the principles, or no more than that it should be as if the person knew the principles.

The account of tacit knowledge given in this entry (section 1.3) is certainly more demanding than those trivializing accounts. But Heal (1994) has argued that, even on that account of tacit knowledge, the mental simulation theory turns out to be a version of the tacit theory theory. In fact, she argues more generally that, if a mechanism is used to simulate the operation of other mechanisms of the same kind so as to permit predictions about them, then the mechanism embodies tacit knowledge of theoretical principles about how mechanisms of that kind operate.

Suppose (to take an extremely simple example) that a subject, S, begins with the information about another person, O, that O believes the proposition that interest rates are falling (p). Suppose, too, that S wonders whether, if the question of the truth of the disjunctive proposition that interest rates are falling or the stock market is rising (p or q) were to arise, O would believe that it is true. In order to generate a prediction about O, S co-cognizes with O and, in suppositional or imaginative mode, thinks that p and considers whether the proposition p or q is true. Drawing on her own tacit knowledge of

the inference rule of ‘*or*’-introduction or her own tacit knowledge of the contribution made to truth conditions by the concept of alternation, S concludes that *p or q* is true. Then, on the basis of this simulation exercise, S predicts that O believes *p or q*.

Clearly, S could proceed in the same way in a multitude of actual or potential cases. Beginning from the information that a subject X believes a proposition *A*, S would arrive at the prediction that X will believe a proposition of the ‘*A or B*’ form, if the question of that proposition’s truth arises. Furthermore, S’s transition would have the same causal explanation in each case. Thus, Heal (1994) argues, S has tacit knowledge of the general psychological principle that if X believes *A* then X also believes *A or B* – just as the tacit theory theory would say.

Davies and Stone (2001) respond to Heal’s argument with two points. First, the attribution to S of tacit knowledge of the psychological principle that if X believes *A* then X also believes *A or B* is unmotivated. Simulation theorists, including Heal, agree that the use of mental simulation to arrive at psychological predictions about other people depends on tacit or explicit knowledge or acceptance of some principle linking the simulating subject’s own thinking with other people’s thinking. One example of a principle of this kind is (Sim–Other): If, within the scope of a simulation exercise that begins from the thought *C*, I arrive at the thought *D* then, if O believes *C* then *ceteris paribus* O will also believe *D*. The transition in S’s simulation exercise from the thought that interest rates are falling (*p*) to the thought that interest rates are falling or the stock market is rising (*q*), together with S’s acceptance of the principle (Sim–Other), already provides an explanation of S’s proceeding from the information that O believes *p* to the prediction that O will believe *q*. No additional knowledge, tacit or explicit, about psychological matters is required.

The second point that Davies and Stone (2001) make is that the attribution to S of tacit knowledge of the psychological principle that if X believes *A* then X also believes *A or B* is not licensed by the account of tacit knowledge. Heal allows that some component of S’s subpersonal-level reasoning system embodies tacit knowledge about the inferential or truth conditional properties of alternation and she proposes that this same information-processing component also embodies tacit knowledge of the principle that if X believes *A* then X also believes *A or B*. But states of tacit knowledge figure in content-involving explanations and the component in question does not play the right causal-explanatory role to embody tacit knowledge of that principle. A state of tacit knowledge of the psychological principle would explain transitions between

representations under their content-involving descriptions. These would have to be representations with contents of the forms ‘X believes *A*’ and ‘X believes *A or B*’. But the information-processing component in question mediates transitions between states whose contents do not concern believing at all but instead (in our example) concern financial matters. The transition between representations with contents ‘interest rates are falling’ and ‘interest rates are falling or the stock market is rising’ is, of course, entirely appropriate to the causal-explanatory role of a state of tacit knowledge about the inferential or truth conditional properties of alternation. But it is not appropriate to the causal-explanatory role of a state of tacit knowledge of a psychological principle.

5. Tacit Knowledge in Psychology: Artificial Grammar Learning

Half a century ago, Chomsky’s notion of tacit knowledge of the rules of a generative grammar provided the theoretical background for empirical research on artificial grammar learning (AGL). Chomsky’s early work on formal language theory also provided the main tool that has been used in this research.

In a typical AGL experiment, the grammar in question is a regular or finite state grammar. Chomsky (1956) described finite state grammars, phrase structure grammars, and transformational grammars, and showed that finite state grammars cannot capture the syntax of English. Chomsky and Miller (1958) provided a more extensive formal account of finite state grammars and the languages that they generate. Although finite state grammars do not include a grammar for English, they illustrate well enough the core notions that Chomsky associates with the faculty of language: a finite vocabulary and a procedure that generates infinitely many finite complex structures (expressions) from the vocabulary items.

A regular or finite state grammar can be represented graphically by a state diagram (a diagram for a finite state automaton); an example is provided in Figure 1. There is a finite set of states – in the example, four states (S_1, S_2, S_3, S_4) apart from the start and end states (S_0 and S_0'); and arrows indicate possible transitions between states. The vocabulary is a finite set of items – in the example, the set of letters $\{P, S, T, V, X\}$. Beginning in the start state, S_0 , each successive transition adds a letter to the string. The state that the system is in determines which letters can be added (e.g. in state S_2 , S or X can be added) and no memory for earlier transitions, or for letters already in the string, is required.

A string of letters is an expression of the language generated by the grammar just in case it corresponds to a path through the state diagram, from the start state to the end

state. Thus, the strings TTS and VVS are expressions of the language generated by the grammar corresponding to the state diagram in Figure 1, as are TPPTS, TPPPTS, TPPPTS, and so on, and TTXVPS, TTXVPXVPS, TTXVPXVPXVPS, and so on.




Figure 1. Diagram of a finite state grammar

A grammar is canonically specified by a set of rules and the rules of a regular grammar have the form: $A \rightarrow a B$, where A denotes a state other than the end state, a denotes a vocabulary item, and B denotes a state. The rules of a regular grammar correspond to the possible transitions in the state diagram that represents the grammar graphically. Each rule specifies the state in which the transition applies, the vocabulary item that is added to the string, and the state following the transition. Thus, the rules of the grammar represented by the state diagram in Figure 1 include: $S_0 \rightarrow T S_1$, $S_0 \rightarrow V S_3$, $S_1 \rightarrow P S_1$, $S_1 \rightarrow T S_2$, $S_2 \rightarrow X S_3$, $S_2 \rightarrow S S_4$, and so on. The rules can be read as saying that the symbol on the left of the arrow can be replaced by the symbols on the right (they are 'rewrite' rules). A string, w , is said to be generated by the grammar if and only if it is possible to begin with the symbol for the start state and reach w (followed by the symbol for the end state) by a finite sequence of rule applications. Thus, for example, a derivation of the string TTS would proceed from S_0 to $T S_1$ (by the first rule above), then to $TT S_2$ (by the fourth rule), and finally to $TTS S_4$ (by the sixth rule). For the purpose of displaying the derivations of strings, a more perspicuous notation would use the null

symbol to denote the end state, rules would have the form $A \rightarrow a B$ (as before) or $A \rightarrow a$, and the sixth rule would be simply: $S_2 \rightarrow S$.

5.1 Learning and Classification

In a typical AGL experiment, subjects learn to classify letter strings as being grammatical or ungrammatical. In general, there are many ways in which subjects may learn to classify items as belonging or not belonging to some category. In one kind of case, subjects may be provided with a principle or rule of classification. That is, a teacher may make a rule explicit and the subjects may acquire explicit knowledge of the rule and use that explicit knowledge in making their classifications.

In other cases, subjects may simply be provided with a set of exemplars of the category to study or memorize. After a period of exposure to the exemplars, subjects may achieve performance on the classification task that is significantly better than would be predicted if they were guessing (or tossing a coin). In principle, their classifications might even be reliably correct. A subject who has attained this level of performance may also be able to state a rule to which her classification performance reliably conforms. That is, even without being explicitly told any rule of classification, the subject may acquire explicit knowledge of a rule. In such a case, it may be (though it need not be) that it is because the subject has acquired explicit knowledge of the rule that her classifications are reliably correct.

It may, however, happen that, following the period of exposure to exemplars, the subject achieves classification performance that substantially, or even reliably, conforms to a rule of classification, although the subject is quite unable to state the rule. Indeed, the subject may be unable to provide any illuminating account of how she is making the classification. Such a case fits Polanyi's (1967) conception of tacit knowing – knowing more than one can tell. The subject is able to – has learned to – make the classification, but cannot tell – cannot make explicit – how to make the classification.

Cases of this kind are sometimes described as examples of implicit learning: the subject 'learns about the structure of a fairly complex stimulus environment, without necessarily intending to do so, and in such a way that the resulting knowledge is difficult to express' (Dienes and Berry 1997). Implicit learning is an unconscious and unintentional process (Shanks 2005) – in contrast to learning by testing or evaluating explicitly formulated hypotheses, for example. Because the description of implicit learning seems to fit learning (or, more generally, acquisition) of our first language, it is

natural to ask whether implicit learning results in Chomskyan tacit knowledge of a rule (or a set of rules) of classification.

5.2 Implicit Learning and Tacit Knowledge of Rules

At the beginning of a seminal paper in which the term ‘implicit learning’ was introduced, Reber (1967) wrote: ‘In recent years, the model of the verbal organism as an imitative and generalizing mechanism has been largely replaced by a model that characterizes him as a “sentence generating machine” who has learned a “generative grammar” in some implicit fashion.’ In the learning phase of his first AGL study, Reber (1967, Experiment 2) presented subjects with 20 strings generated by the finite state grammar in Figure 1. In this phase, there was no mention of grammars or rules and the subjects were required simply to memorize the presented strings. At the beginning of the test phase of the study, subjects were told that the 20 strings had been generated by a set of grammatical rules, but were provided with no further information about the rules. Subjects were then asked to classify new strings as either grammatical or ungrammatical. The striking finding was that 79% of the subjects’ classification responses were correct.

A separate experiment (1967, Experiment 1) showed that, following an extensive learning phase (in which, again, there was no mention of grammars or rules), subjects were unable to answer any questions about what the rules governing the strings might be. Unless they were given considerable prompting, subjects were unable even to say which letters the strings could begin with (T or V), or end with (only S). Reber concluded that the learning phase of the experiment provided subjects with abstract information about the rule-governed structure of the presented strings; and that this information was used to classify new strings in the test phase, although no verbal account of the rules could be elicited from subjects. In short, Reber’s view was that, through a process of implicit learning, tacit knowledge of rules was acquired.

5.3 Tacit Knowledge of Rules or Explicit Knowledge of Fragments?

It has been argued, however, that it would be possible for subjects to achieve the better-than-chance classification performance that is typically demonstrated in the test phase of an AGL experiment without acquiring knowledge of the rules of the grammar used to generate the strings presented in the learning phase. One proposal is that information about pairs of adjacent letters (bigrams) that occur in the presented strings (perhaps along with information about allowable first and last letters) would be sufficient to account for

subjects' classification performance. Furthermore, the results of a recognition memory test support the hypothesis that, in the learning phase, subjects acquire explicit knowledge about which bigrams occurred in the presented strings. Thus, an alternative to Reber's view is that what is acquired during the learning phase and used during the test phase is not tacit knowledge of abstract rules but explicit knowledge of specific fragments (Perruchet and Pacteau 1990). In a review of research on implicit learning, Shanks (2005) supports this alternative proposal: 'participants merely focus on and learn about small letter chunks. In the test stage they endorse (i.e. call grammatical) any item composed mainly of familiar letter pairs'.

Here it is relevant that there are findings of transfer of classification performance to strings that are generated by the same finite state grammar, but from a new set of letters (Reber 1969), or even to strings of stimuli presented in a different modality from those in the learning phase (e.g. letters *versus* tones; Altmann, Dienes and Goode 1995). Reber explains transfer in terms of abstract knowledge (even if not knowledge of the basic rules of the grammar) acquired during the learning phase. An alternative proposal is that transfer can be explained in terms of abstraction or analogical reasoning in the test phase from explicit knowledge, about specific fragments (e.g. bigrams or trigrams) or about specific complete exemplars, acquired in the learning phase. Shanks (2005) argues that the balance of evidence favors the alternative proposal and concludes: 'In a nutshell, performance in AGL experiments appears to be fully explained without the need to impute unconscious rule-learning'.

5.4 Artificial Grammar Learning in Patients with Amnesia

In comparing Reber's view with alternatives, we have contrasted implicit abstract knowledge with explicit specific knowledge. But the implicit *versus* explicit distinction and the abstract *versus* specific distinction are not to be conflated.

Patients suffering from amnesia cannot (by definition) retain new explicit knowledge but, in AGL experiments, they show very similar classification performance to that of age-matched controls (Knowlton and Squire 1994, 1996). This finding is enough to rebut any suggestion that acquisition of explicit knowledge in the learning phase is essential for better-than-chance classification performance in the test phase of AGL experiments. But the fact that patients with amnesia cannot rely on explicit knowledge acquired in the learning phase to classify new strings presented in the test phase does not settle the question whether the patients rely on abstract or specific information or both.

Knowlton and Squire showed that, in both patients and controls, classification performance was influenced by the occurrence in test strings of bigrams and trigrams (chunks) that occurred frequently in strings presented in the learning phase, despite the fact that the patients with amnesia scored only marginally above chance on a recognition memory test for chunks that had appeared in the learning phase. Classification performance was also influenced by whether the test strings conformed to the rules of the grammar and there was some transfer of performance to test strings composed from a new set of letters (though at a lower level than with test strings composed from the original set of letters). At least in the case of patients with amnesia, it is not plausible that transfer of classification performance depends on abstraction or analogy in the test phase, based on explicit knowledge about chunks or complete strings retained from the learning phase. Rather, it seems that patients acquire specific information about chunks and complete strings presented in the learning phase, and also more abstract information that supports transfer.

Knowlton and Squire's (1994, 1996) findings do not exclude the possibility that subjects with unimpaired memory – and particularly, perhaps, healthy younger subjects – may acquire in the learning phase, and use in the test phase, information that is available for verbal report. But the findings, and the similar performance of patients and aged-matched controls, certainly suggest that, even if in normal healthy subjects some of the relevant information is available for verbal report, not all of it is. As Reber (1989) puts the point: 'the implicitly acquired epistemic contents of mind are always richer and more sophisticated than what can be explicated'.

6. Conscious and Unconscious Knowledge

It will be clear from the previous section that the dialectical context of AGL experiments is different from that of typical appeals to tacit knowledge in philosophy. Philosophical appeals to tacit knowledge typically occur in contexts where it is not in dispute that relevant explicit knowledge is lacking. In AGL research, in contrast, the claim that what has been acquired in the learning phase is implicit or tacit knowledge, rather than explicit knowledge, is highly contested. It is also a feature of AGL research that the question of what has been acquired in the learning phase is often formulated in terms of conscious knowledge *versus* unconscious knowledge – a contemporary version of the explicit knowledge *versus* implicit knowledge contrast. Thus, considerable importance attaches to the empirical assessment of conscious *versus* unconscious knowledge.

6.1 Two Criteria for a Test of Conscious Knowledge

Shanks and St John (1994) argue that, in studies that aim to demonstrate that unconscious knowledge (e.g. Chomskyan tacit knowledge of rules) has been acquired through implicit learning, two criteria must be met in order to exclude the possibility that what subjects have acquired is really conscious knowledge. First, if the subjects' awareness of a particular kind of information is tested then this must be the same kind of information that is drawn on in the subjects' performance (the information criterion). For example, in Reber's (1967) study, subjects' conscious knowledge of the rules governing the presented strings was tested. So it is important that it should be information about those rules that was drawn on in subjects' classification of new strings as grammatical or ungrammatical. Clearly, no tacit knowledge or implicit learning is demonstrated if one kind of information (e.g. about bigrams that occurred in strings presented in the learning phase) is drawn on in subjects' performance but subjects are found to be unaware of some other kind of information (e.g. about the rules of the finite state grammar).

The second criterion is that tests of subjects' conscious knowledge should be sufficiently sensitive to reveal all of the relevant conscious knowledge (the sensitivity criterion). Here it is relevant that, in research on memory, both recall and recognition have been regarded as explicit tests of memory (Schacter 1989), with recognition being generally more sensitive than recall. Thus, if subjects' performance on a classification test were to draw on information about grammatical rules then simply asking subjects what the rules are (a test of free recall) would meet the information criterion, but might be relatively insensitive to subjects' actual conscious knowledge of rules. For example, subjects might not volunteer a statement of conscious knowledge about which they were not confident. If a test of conscious knowledge used in an AGL experiment does not meet both the information criterion and the sensitivity criterion then task-relevant conscious knowledge acquired during the learning phase may be underestimated.

6.2 Exhaustiveness and Exclusiveness

Lessons for research on implicit learning have also been drawn from Reingold and Merikle's (1990) discussion of research on unconscious perception. They note that one requirement for an adequate measure of conscious awareness is that it should be exhaustive (a requirement very similar to the sensitivity criterion). Given a finding that performance apparently outruns conscious perception (in Reingold and Merikle's case) or

conscious knowledge (in the AGL case), it cannot be concluded that there is perception or learning without awareness if the test of conscious awareness is not exhaustive.

A second requirement for a test of conscious awareness is that it should be exclusive; that is, results on the test should be influenced only by conscious, and not also by unconscious, perception or knowledge. Given a finding that performance apparently does not outrun conscious perception or knowledge, it cannot be concluded that there is no perception or learning without awareness if the test of conscious awareness is not exclusive (Reingold and Merikle 1990). In the AGL case, specifically, results on a test of conscious awareness of what has been learned should be influenced only by conscious knowledge, and not contaminated by unconscious knowledge.

If a test of conscious knowledge is not exhaustive then it may underestimate conscious knowledge and overestimate implicit learning; if a test is not exclusive then it may overestimate conscious knowledge and underestimate implicit learning. Different theorists put more stress on the exhaustiveness (sensitivity) requirement or on the exclusiveness requirement and this may reflect, in part, their view as to whether conscious or unconscious processing enjoys theoretical primacy. One influential view is that, in human beings, conscious learning is the default case and ‘unconscious learning has not yet been satisfactorily established’ (Shanks and St John 1994; see also Shanks 2005). Reber, on the other hand, defends the primacy of the implicit (Reber 1990).

6.3 Validating a Test of Conscious Knowledge

Unless a test of conscious knowledge is, and is known to be, both exhaustive and exclusive it cannot, by itself, be used to adjudicate evenhandedly a putative case of implicit learning. We have already seen that a test of conscious knowledge based on free recall is unlikely to be exhaustive. Tests based on recognition memory (rather than recall), which we have mentioned as being more sensitive to conscious knowledge, may not be exclusive because responses may be contaminated by unconscious processes. For example, recognition judgments are influenced by perceptual fluency. Subjects are more likely to ‘recognize’ a test word falsely as having occurred on a previously studied list if their perception of the test word is facilitated by very brief prior presentation of the same word under conditions (masking) that preclude conscious perception (Jacoby and Whitehouse 1989).

The fact that a recognition memory test of conscious knowledge is not exclusive illustrates the more general point that distinctions between psychological tasks

(e.g. between so-called explicit and implicit tests of memory) do not coincide with distinctions between psychological processes (Jacoby 1991). Recognition has been regarded as an explicit test of memory, but the results of a recognition test may be influenced by unconscious, as well as by conscious, processes. More generally, we should not assume that an experimental task (particularly, a task used as a test of conscious knowledge) is ‘process pure’, that is, that it involves only conscious processes or else only unconscious processes. There is no a priori guarantee that any particular test of conscious knowledge will be exhaustive and exclusive.

Reingold and Merikle (1990) argue for an approach on which a putative test (or ‘operational definition’) of conscious awareness is validated by converging empirical evidence. In their view, the theoretical importance of the distinction between conscious and unconscious processes lies in the existence of qualitative differences between the two. Sometimes, indeed, conscious and unconscious processes are not just qualitatively different, but opposed to each other. For example, the recognition memory experiment briefly mentioned two paragraphs back involved two conditions. In one (unaware condition, described above), presentation of a test word was preceded by a very brief masked presentation of the same word. In the aware condition, in contrast, presentation of a test word was preceded by a presentation of the same word that was long enough to allow conscious perception. As described above, subjects in the unaware condition were more likely to ‘recognize’ the test word falsely as having occurred on a previously studied list (by comparison with an appropriate baseline condition in which there was no prior presentation of the test word). But subjects in the aware condition were less likely (than in the baseline condition) to ‘recognize’ the test word falsely as having occurred on the previously studied list.

Thus, there is the prospect that a putative test of conscious knowledge can be validated in the context of psychological theory to the extent that the test correctly predicts processing of qualitatively different – perhaps opposed – kinds. In the next section, we describe some AGL research that adopts the approach recommended by Reingold and Merikle.

7. Conscious Knowledge in the Context of Psychological Theory

In the learning phase of an AGL experiment, subjects may acquire knowledge about the structure of the presented strings. This may be knowledge of rules or knowledge about individual strings. Possible examples of knowledge of rules include knowledge of the

rules of the finite state grammar that was used to generate the strings, but also knowledge of simpler rules such as that strings can begin with T or V, that all strings end with S, and that P can be repeated. Possible examples of knowledge about individual strings presented in the learning phase include knowledge that the string VXXVS was presented, that a string including the bigram TS was presented, and that a string including the trigram VPX was presented. (The distinction between the two kinds of knowledge is not sharp. If a string beginning with VX was presented in the learning phase then strings can begin with VX.) The knowledge acquired in the learning phase may be drawn on in the test phase as subjects make judgments about test strings, to the effect that they are grammatical (that is, conform to the rules that were used to generate the strings presented in the learning phase) or ungrammatical.

7.1 Conscious and Unconscious Judgment Knowledge

Dienes and Scott (2005) distinguish between structural knowledge acquired during the learning phase and what they call ‘judgment knowledge’ arrived at and deployed in the test phase – for example, the knowledge that the string VVPXXVS, presented in the test phase, is grammatical. The aim of their study was to assess each kind of knowledge as conscious or unconscious, in the context of psychological theories about conscious and unconscious processes. Dienes and Scott commend two confidence criteria for unconscious knowledge, the guessing criterion and the zero-correlation criterion (see also Dienes 2008, Dienes 2012). A subject who is asked to judge whether strings of letters are grammatical or not may, on some trials, believe herself to have no basis for the judgment and, consequently, she may have no confidence in her judgment; she may describe herself as guessing. (Strictly speaking, she may have confidence of 50% in her judgment – she might as well be tossing a coin.) If the subject’s guessing responses are correct at a rate that is better than chance then this is evidence (that is, *prima facie* but defeasible evidence) of unconscious judgment knowledge (according to the guessing criterion).

On other trials, the subject may believe herself to be knowledgeable to some extent and may have some degree of confidence in her judgment. If the subject’s guessing responses and these more confident responses together reveal a correlation between accuracy and confidence, then this is evidence of some conscious judgment knowledge. If there is no correlation between accuracy and confidence but the subject’s overall performance is better than chance then this is, again, evidence of unconscious judgment knowledge (according to the zero-correlation criterion). Clearly, the use of the guessing

criterion and the zero-correlation criterion may provide evidence of both unconscious and conscious judgment knowledge. For example, a subject's guessing performance may be better than chance and her more confident judgments may be more likely to be correct than her guesses.

Dienes and Scott (2005) investigated judgment knowledge by asking subjects in the test phase of an AGL experiment to classify each string as grammatical or ungrammatical and then to indicate their degree of confidence in the classification judgment (from 50% to 100% confident). Overall, subjects' confidence ratings provided evidence of both unconscious and conscious judgment knowledge.

7.2 Conscious and Unconscious Structural Knowledge

In principle, it would seem that conscious or unconscious judgment knowledge and structural knowledge could be combined in three ways. In the case of guessing, there is no conscious judgment knowledge or conscious structural knowledge. But conscious judgment knowledge can be combined with either unconscious or conscious structural knowledge. When a subject's judgment knowledge is conscious but her structural knowledge is unconscious, the phenomenology may be that of intuition.

Dienes and Scott (2005) investigated structural knowledge by asking subjects to indicate the source of their judgment knowledge, choosing between four options: guess (it seemed to the subject that the judgment had no basis whatsoever); intuition (the subject had some confidence in the judgment but had no idea why it was right); rules (the subject felt that she based the judgment on some rule or rules acquired in the learning phase – rules that she could state if asked); memory (the subject felt that she based the judgment on memory for particular items or parts of items from the learning phase). The 'guess' and 'intuition' responses provided evidence that the subject's structural knowledge was unconscious, while the 'rules' and 'memory' responses provided evidence that it was at least partly conscious. Some more recent studies include an additional potential source of subjects' judgment knowledge: familiarity (the subject felt that she based the judgment on a feeling of familiarity but had no idea what the familiarity itself was based on). The 'familiarity' response, like the 'intuition' response, provides evidence of unconscious structural knowledge (Scott and Dienes 2008; and see Dienes, Scott and Wan 2011, for a review).

It is independently plausible on theoretical grounds that the acquisition and use of conscious structural knowledge would make demands on working memory and executive

function. Thus, having subjects perform an executive working memory load task (e.g. generating random numbers) during the learning phase of an AGL experiment should disrupt the acquisition of conscious structural knowledge by comparison with unconscious structural knowledge. If ‘rules’/‘memory’ responses indicate conscious structural knowledge then one can predict that subjects who performed a load task during the learning phase should be less likely (than subjects who did not perform the load task) to attribute the source of their judgments in the test phase to rules or memory. This prediction was confirmed in Dienes and Scott’s (2005) study and, to that extent, the use of the ‘rules’/‘memory’ *versus* ‘guess’/‘intuition’ contrast to distinguish between conscious and unconscious structural knowledge was validated.

The study by Dienes and Scott (2005) investigated the distinction between conscious and unconscious structural knowledge in the context of psychological theory about working memory and executive function. But the load task manipulation did nothing to validate the distinction between conscious and unconscious judgment knowledge.

7.3 The Serial Reaction Time Task

Converging evidence to support the distinction between conscious and unconscious judgment knowledge has come from a different implicit learning paradigm, the serial reaction time (SRT) task (Destrebecqz and Cleeremans 2001). In the learning phase, subjects see a sequence of stimuli on a computer screen. Each stimulus appears at one of several (e.g. four) possible locations on the screen and the subject’s task is simply to press a key corresponding to the stimulus location as quickly as possible. In fact, the sequence of stimuli follows a repeating pattern, in which the locations are determined by second-order conditional rules. That is, each position is determined by the previous two positions (e.g. positions 4 and then 2 are always followed by position 3). Learning is revealed in the pattern of reaction times, which decrease as the learning phase progresses, increase sharply when a new sequence is introduced, and then drop back to the lower level when the original sequence returns.

In the test phase, there are two conditions, motivated by Jacoby’s (1991) ‘process dissociation framework’. The leading idea is that the use of conscious knowledge can be controlled, whereas the use of unconscious knowledge is automatic. In one condition of the test phase (inclusion), subjects are asked to generate a sequence that resembles the training sequence as much as possible. In this condition, automatic use of unconscious knowledge is facilitatory. Subjects can perform the task by drawing on either conscious

or unconscious knowledge about which position would be a legal continuation of the sequence (e.g. that 3 would be a legal continuation of ...42). Consequently, legal continuations in the inclusion condition do not provide evidence that distinguishes between conscious and unconscious knowledge acquired during the learning phase.

In a second – theoretically crucial – condition (exclusion), subjects are asked to generate a sequence that avoids the regularities of the training sequence. In this condition, controlled use of conscious knowledge and automatic use of unconscious knowledge are set against each other. To perform the task, subjects need to make controlled use of their conscious knowledge about which position would be a legal continuation. Specifically, they need to continue in a way that avoids that position. In this condition, automatic use of unconscious knowledge interferes with, rather than facilitates, performance of the exclusion task. Thus, legal continuations in the exclusion condition provide evidence of unconscious knowledge acquired during the learning phase. (Relevant evidence would include the proportion of legal continuations being above a suitable baseline and there being no fewer legal continuations in the exclusion than in the inclusion condition.)

The results of experiments using the SRT task do not all point in the same direction. But, given appropriate conditions in the learning phase, evidence from the exclusion condition in the test phase does support the hypothesis that unconscious knowledge about legal continuations is acquired (Destrebecqz and Cleeremans 2001, Fu, Fu and Dienes 2008). Furthermore, it can be argued that the relevant distinction for the application of Jacoby's (1991) process dissociation framework to sequence generation in the test phase of the SRT task is not between conscious and unconscious structural knowledge, but between conscious and unconscious judgment knowledge. Specifically, the relevant distinction is between intuition and guessing, because conscious judgment knowledge that is based on unconscious structural knowledge can still be controlled (Fu, Dienes and Fu 2010).

7.4 Summary and Continuing Research

When the distinction between explicit and implicit knowledge is interpreted as that between conscious and unconscious knowledge, empirical research using the AGL task can be seen as a chapter in the scientific study of consciousness. In this research, attributions of tacit knowledge of rules, and of unconscious knowledge more generally, have been contested and various requirements for a test of conscious knowledge have been proposed.

There is no test of conscious knowledge for which we have a guarantee that it will meet the demanding requirements of exhaustiveness and exclusiveness – revealing all the relevant conscious knowledge while remaining uncontaminated by unconscious knowledge. But putative tests may be validated by converging evidence in the context of psychological theory. Dienes and colleagues have proposed tests for conscious judgment knowledge (based on subjects' confidence ratings) and for conscious structural knowledge (based on subjects' source attributions for their judgments). As briefly described in this section, they have taken steps to validate the tests by appeal to converging evidence from (respectively) the exclusion condition of an experiment based on Jacoby's (1991) process dissociation framework and an executive working memory load task (see Dienes 2012, for a review).

Recent and continuing research on implicit learning, using the AGL task, is investigating new measures of conscious awareness such as wagering (Persaud, McLeod and Cowey 2007, Dienes and Seth 2010), and comparing different measures (Wierzchon, Asanowicz, Paulewicz and Cleeremans 2012). Research also continues on feelings of familiarity and their relationship to perceptual fluency (Scott and Dienes 2008, 2010, Dienes, Scott and Wan 2011) and the AGL paradigm has been adapted to investigate cross-cultural differences in unconscious processes (Kiyokawa, Dienes, Tanaka, Yamada and Crowe 2012). In other work, researchers have developed implemented computational models of implicit learning tasks, including the AGL and SRT tasks (see Cleeremans and Dienes 2008, for a review).

Finally, the methods of AGL are being used, in conjunction with formal language theory, to investigate learning – particularly, implicit learning – of grammars with greater expressive power than regular (that is, finite state) grammars (Rohrmeier, Fu and Dienes, 2012). (Recall that Chomsky (1956) showed that finite state grammars cannot capture the syntax of English.) Some of this research extends to AGL in pre-verbal infants and non-human animals and addresses the questions whether the human capacity to acquire supra-regular grammars is biologically distinctive and what its neural substrate might be (Fitch, Friederici and Hagoort 2012, Fitch and Friederici 2012).

Bibliography

- Altmann, G. T. M., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 899–912.
- Baker, G. P., & Hacker, P. M. S. (1984). *Language, Sense and Nonsense: A Critical Investigation into Modern Theories of Language*. Oxford, UK: Blackwell.
- Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical Foundations of Neuroscience*. Oxford: Blackwell.
- Blackburn, S. (1992). Theory, observation and drama. *Mind & Language*, 7, 187–203 [Reprinted in Davies & Stone (Eds.), 1995].
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, 67, 503–48.
- Chomsky, N. (1956). Three models for the description of language. *IRE Transactions on Information Theory*, 2(3), 113–24.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger.
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2005). Three factors in language design. *Linguistic Inquiry*, 36, 1–22.
- Chomsky, N., & Miller, G. A. (1958). Finite state languages. *Information and Control*, 1, 91–112.
- Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology*. Cambridge: Cambridge University Press
- Collins, J. (2004). Faculty disputes. *Mind & Language*, 17, 503–33.
- Collins, J. (2007). Linguistic competence without knowledge of language. *Philosophy Compass*, 2, 880–95.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–56.
- Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

- Davies, M. (1987). Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96, 441–462.
- Davies, M. (1989). Tacit knowledge and subdoxastic states. In A. George (Ed.), *Reflections on Chomsky*. Oxford, UK: Blackwell.
- Davies, M. (1995). Two notions of implicit rules. *Philosophical Perspectives*, 9: *AI, Connectionism, and Philosophical Psychology*, 153–83.
- Davies, M. (2000). Interaction without reduction: The relationship between personal and sub-personal levels of description. *Mind and Society*, 1, 87–105.
- Davies, M., & Stone, T. (2001). Mental simulation, tacit theory, and the threat of collapse. *Philosophical Topics*, 29, 127–73.
- Davies, M., & Stone, T. (Eds.) (1995). *Folk Psychology: The Theory of Mind Debate*. Oxford: Blackwell.
- Davies, M., Coltheart, M., Langdon, R., & Breen, N. (2001). Monothematic delusions: Towards a two-factor account. *Philosophy, Psychiatry, & Psychology*, 8, 133–58.
- Dennett, D. C. (1969). *Content and Consciousness*. London : Routledge and Kegan Paul.
- Dennett, D. C. (1983). Styles of mental representation. *Behavioral and Brain Sciences*, 83, 213–226 [Reprinted in: *The Intentional Stance*. Cambridge, MA: MIT Press, 1987].
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, 8, 343–350.
- Dienes, Z. (2008). Subjective measures of unconscious knowledge. *Progress in Brain Research*, 168, 49–64.
- Dienes, Z. (2012). Conscious versus unconscious learning of structure. In P. Rebuschat, & J. N. Williams (Eds.), *Statistical Learning and Language Acquisition*. Berlin: de Gruyter Mouton.
- Dienes, Z., & Berry, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin and Review*, 4, 3–23.
- Dienes, Z., & Seth, A. (2010). Gambling on the unconscious: A comparison of wagering and confidence ratings as measures of awareness in an artificial grammar task. *Consciousness and Cognition*, 19, 674–81.
- Dienes, Z., & Scott, R. (2005). Measuring unconscious knowledge: Distinguishing between structural knowledge and judgment knowledge. *Psychological Research*, 69, 338–51.

- Dienes, Z., Scott, R. B., & Wan, L. (2011). The role of familiarity in implicit learning. In P. A. Higham, & J. P. Leboe (Eds.), *Constructions of Remembering and Metacognition: Essays in Honour of Bruce Whittlesea*. Basingstoke: Palgrave Macmillan.
- Dummett, M. (1991). *The Logical Basis of Metaphysics*. London: Duckworth.
- Egan, A. (2008). Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140, 47–63.
- Evans, G. (1981). Semantic theory and tacit knowledge. In S. H. Holtzmann, & C. M. Leich (Eds.), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul [Reprinted in: *Collected Papers*. Oxford University Press, Oxford, 1985].
- Fitch, W. T., & Friederici, A. D. (2012). Artificial grammar learning meets formal language theory: An overview. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367, 1933–55.
- Fitch, W. T., Friederici, A. D., & Hagoort, P. (2012). Pattern perception and computational complexity: Introduction to the special issue. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367, 1925–32.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Appendix, ‘Why there still has to be a language of thought’. Cambridge, MA: MIT Press.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.
- Fu, Q., Dienes, Z., & Fu, X. (2010). Can unconscious knowledge allow control in sequence learning? *Consciousness and Cognition*, 19, 462–474.
- Fu, Q., Fu, X., & Dienes, Z. (2008). Implicit sequence learning and conscious awareness. *Consciousness and Cognition*, 17, 185–202.
- Goldman, A. I. (1989). Interpretation psychologized. *Mind & Language*, 4, 161–85 [Reprinted in Davies & Stone (Eds.), 1995].
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569–79.
- Heal, J. (1986). Replication and functionalism. In J. Butterfield (Ed.), *Language, Mind and Logic*. Cambridge: Cambridge University Press [Reprinted in Davies & Stone (Eds.), 1995].

- Heal, J. (1994). Simulation vs. theory theory: What is at issue? In C. Peacocke (Ed.), *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind (Proceedings of the British Academy, vol. 83)*. Oxford: Oxford University Press.
- Heal, J. (1998). Co-cognition and off-line simulation: Two ways of understanding the simulation approach. *Mind & Language*, 14, 477–98.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, 30, 513–41.
- Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, 118, 126–35.
- Kiyokawa, S., Dienes, Z., Tanaka, D., Yamada, A., & Crowe, L. (2012). Cross cultural differences in unconscious knowledge. *Cognition*, 124, 16–24.
- Knowlton, B. J., & Squire, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 79–91.
- Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 169–81.
- Lewis, D. (1992). Meaning without use: Reply to Hawthorne. *Australasian Journal of Philosophy*, 70, 106–10.
- Loar, B. (1981). *Mind and Meaning*. Cambridge: Cambridge University Press.
- Matthews, R. J. (2006). Knowledge of language and linguistic competence. *Philosophical Issues*, 16, *Philosophy of Language*, 200–20.
- Mantzavinos, C. (2001). *Individuals, Institutions, and Markets*. Cambridge: Cambridge University Press.
- Mantzavinos, C., North, D. C., & Shariq, S. (2004). Learning, institutions, and economic performance. *Perspectives on Politics*, 2, 75–84.
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Peacocke (1989). When is a grammar psychologically real? In A. George (Ed.), *Reflections on Chomsky*. Oxford, UK: Blackwell.

- Peacocke, C. (2008). *Truly Understood*. Chapter 4, 'Implicit conceptions'. Oxford: Oxford University Press.
- Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, 119, 264–75.
- Persaud, N., McLeod, P., & Cowey, A. (2007). Post-decision wagering objectively measures awareness. *Nature Neuroscience*, 10, 257–61.
- Plaut, D.C., McClelland, J.L., Seidenberg, M.S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.
- Polanyi, M. (1967). *The Tacit Dimension*. London: Routledge and Kegan Paul.
- Quine, W. V. O. (1972). Methodological reflections on current linguistic theory. In D. Davidson, & G. Harman (Eds.), *Semantics of Natural Language*. Dordrecht: Reidel.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6, 855–63.
- Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology*, 81, 115–9.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219–235.
- Reber, A. S. (1990). On the primacy of the implicit: Comment on Perruchet and Pacteau. *Journal of Experimental Psychology: General*, 119, 340–2.
- Reingold, E. M., & Merikle, P. M. (1990). On the inter-relatedness of theory and measurement in the study of unconscious processes. *Mind & Language*, 5, 9–28.
- Rohrmeier, M., Fu, Q., & Dienes, Z. (2012). Implicit learning of recursive context-free grammars. *PLoS One*, 7(10), e45885.
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2: Psychological and Biological Models*. Cambridge, MA: MIT Press.
- Schacter, D. L. (1989). Memory. In M. I. Posner (Ed.), *Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Schiffer, S. (1993). Actual-language relations. *Philosophical Perspectives*, 7: *Language and Logic*, 231–58.

- Schiffer, S. (2006). Two perspectives on knowledge of language. *Philosophical Issues, 16: Philosophy of Language*, 275–87.
- Scott, R. B., & Dienes, Z. (2008). The conscious, the unconscious, and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1264–88.
- Scott, R. B., & Dienes, Z. (2010). Fluency does not express implicit knowledge of artificial grammars. *Cognition*, 114, 372–88.
- Searle, J. (1990). Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences*, 13, 585–595.
- Shanks, D. R. (2005). Implicit learning. In K. Lamberts, & R. L. Goldstone (Eds.), *Handbook of Cognition*. London: SAGE Publications Ltd.
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367–395.
- Shea, N. (2007). Content and its vehicles in connectionist systems. *Mind & Language*, 22, 246–69.
- Smith, B. C. (2008). What remains of our knowledge of language? Reply to Collins. *Croatian Journal of Philosophy*, 8, 57–76.
- Stich, S. (1978). Beliefs and subdoxastic states. *Philosophy of Science*, 45, 499–518.
- Stich, S., & Nichols, S. (1992). Folk psychology: Simulation or tacit theory? *Mind & Language*, 7, 35–71 [Reprinted in Davies & Stone (Eds.), 1995].
- Thornton, T. (2013). Clinical judgment, tacit knowledge, and recognition in psychiatric diagnosis. In K. W. M. Fulford, M. Davies, R. G. T. Gipps, G. Graham, J. Z. Sadler, & G. Stanghellini (Eds.), *The Oxford Handbook of Philosophy and Psychiatry*. Oxford: Oxford University Press.
- Wierchoń, M., Asanowicz, D., Paulewicz, B., & Cleeremans, A. (2012). Subjective measures of consciousness in artificial grammar learning task. *Consciousness and Cognition*, 21, 1141–53.
- Williamson, T. (2006). Conceptual truth. *Proceedings of the Aristotelian Society Supplementary Volume*, 80, 1–41.