Against Depictive Conventionalism

Catharine Abell

(Published in *The American Philosophical Quarterly*, 42(3), pp.185-197)

Abstract

In this paper, I discuss the influential view that depiction, like language, depends on arbitrary conventions. I argue that this view, however it is elaborated, is false. Any adequate account of depiction must be consistent with the distinctive features of depiction. One such feature is depictive generativity. I argue that, to be consistent with depictive generativity, conventionalism must hold that depiction depends on conventions for the depiction of basic properties of a picture's object. I then argue that two considerations jointly preclude depiction from being governed by such conventions. Firstly, conventions must be salient to those who employ them. Secondly, those parts of pictures that depict basic properties of objects are not salient to the makers and interpreters of pictures.

I Introduction

Language and depiction are the two paradigmatic forms of non-mental representation. Descriptions represent linguistically, and pictures depictively. Depiction distinguishes figurative from abstract paintings. Both abstract and figurative paintings can represent: a slash of red paint may represent sin and a painting of a lamb may represent Christ. However, all and only figurative representations depict. I will use the term "picture" to refer exclusively to representations that depict. A picture represents Christ in virtue of *depicting* a

lamb. The scope of this paper is limited to depiction and to pictures.

The main task of an adequate philosophical account of depiction is to offer an explanation of what it is for a picture to depict a certain object. This explanation must be consistent with pictures having those features which are taken, pre-theoretically, to be distinctive of depiction.[1] Attempts to identify the relation a picture must bear to an object in order to depict that object appeal to such diverse notions as resemblance, perceptual resources, make-believe, information, experience, and convention. According to conventionalism, the only representationally relevant relations between pictures and their objects are conventional. In this respect, then, they are like words. Nelson Goodman, the most influential proponent of conventionalism, is explicit about the similarity between language and depiction. He writes:

> Descriptions are distinguished from depictions not through being more arbitrary but through belonging to articulate rather than to dense schemes; and words are more conventional than pictures only if conventionality is construed in terms of differentiation rather than of artificiality. Nothing here depends upon the internal structure of a symbol; for what describes in some systems may depict in others.[2]

Conventionalism has been enormously influential in both philosophy and art history. Convinced by Goodman's arguments, many art historians have maintained that depiction is conventional.[3]

Criticisms of conventionalism abound. The vast majority of these, however, attack particular conceptions of the conventional relations between pictures and their objects, rather than the general claim that depiction is conventional.[4] Moreover, none of those arguments that attacks conventionalism in general is conclusive.[5] Philosophers usually take depictive conventionalism to be a live, if unpalatable, option. My aim in this paper is to demonstrate that it is not. I argue that, because conventionalism must be consistent with the distinguishing features of depiction, it must construe depictive conventions as conventions for the depiction of basic colour, shape and textural properties. I then argue that two considerations jointly preclude such conventions from governing depiction. Firstly, I argue that, to be conventional, relations between pictures and their objects must be salient to the makers and interpreters of pictures. Secondly, I argue that those parts of pictures that depict basic properties are not salient to them. I conclude that conventionalism is false because it cannot accommodate one of depiction's distinguishing features.

I proceed as follows. In Section II, I examine what a convention is. To this end, I outline David Lewis's account of convention. In Section III, I examine the claim that depiction involves language-like conventions, and distinguish it from two other forms of depictive conventionalism. In Section IV, I describe a distinguishing feature of depiction, *depictive generativity*. To show that it is not easily accommodated by the view that depiction involves language-like conventions, I distinguish it from the generativity of language. In Section V, I

examine how this view could accommodate depictive generativity. I identify some constraints which it must meet in order to do so. In Section VI, I argue that the view cannot meet these constraints. I elucidate the reasons for which it cannot do so and refute some counterarguments.

II What Conventions Are

Conventionalism holds that depiction depends on language-like conventions. To understand this claim, we need an account of convention. Such an account must make sense of our conviction that language is conventional. Not all philosophical accounts of convention are adequate to this task. For example, some accounts construe convention as involving explicit agreement.[6] Because putative linguistic conventions can arise through either tacit or explicit agreement, these accounts preclude us from understanding language as conventional. However, David Lewis's theory of convention is intended specifically to explain the sense in which language is conventional. Since his is the most influential and widely-accepted such theory, I will use it to provide a better understanding of conventionalism.

According to Lewis, a convention is a widely adopted solution to a recurrent coordination problem.[7] This account comprises six requirements for conventionality. Lewis identifies three requirements that must be met for something to be a coordination problem. These are:

1. *interdependence*: two or more agents must be involved in a situation of interdependent decision making. That is, the decision each agent makes

must depend on what decision he or she believes the other agent(s) will make.

2. *shared interest*: the interests of all agents must coincide in that situation. That is, the decision each agent makes must be such as to serve the other agents' interests as well as his or her own.

3. *choice*: the situation must admit of two or more *proper coordination equilibria*. That is, there must be at least two combinations of all agents' actions such that, given the other agents' actions, each agent prefers that combination of actions to any other combination he or she could have achieved alone.

Lewis then identifies two further requirements that must be met for a coordination problem to have a conventional solution:

4. *recurrence*: the coordination problem must recur.

5. *wide adoption*: a particular solution to the coordination problem must be widely adopted. That is, a single solution to the coordination problem must be adopted on many of the occasions on which that problem arises. For this requirement to be met, there must be some classification of the solutions that are adopted under which many of them are the same.

However, one can classify any set of solutions so as to construe those solutions, and no others, as the same.[8] This threatens to trivialise the requirement that a single solution be widely adopted. For a solution to be widely adopted, in a non-trivial sense, the sameness of a set of solutions must depend on the conceptions of them as solutions that are *salient* to those who adopt them. A certain conception of something as a solution to a coordination problem is salient if it stands out as being unique "in a preeminently conspicuous respect".[9] If the same such conception is salient to many who adopt solutions to a given coordination problem, then the solutions they adopt will be the same in the required sense. This yields a further requirement for conventionality:

> 5.1 *salience*: for many solutions to a coordination problem, a single conception of them as solutions must be salient to those who adopt them.[10]

This requirement reflects the fact that one can only employ a convention if one knows (explicitly or tacitly) what it is. Imagine that I only ever drive my car to work in the morning and back home at night. I always drive on one side of the road in the morning because, at a certain spot, one gets a beautiful view out to sea. Coming home, I drive on the other side, because I cannot see the view from this direction and don't want to impede the view of those driving in the opposite direction. I happen to act in the same way as those who adopt the convention of driving on the left-hand side of the road. However, I do not adopt that convention myself, since the conception of my solution that is salient to me is not one

according to which I always drive on the left-hand side of the road.

III What Conventionalism Is

If depiction is conventional, pictures must comprise widely-adopted solutions to recurrent coordination problems. The problem of depicting particular objects or object types is the best candidate for the relevant coordination problem. It involves interdependent decision making with a coincidence of agents' interests. The makers and interpreters of pictures seek to exchange information about the (actual or counterfactual) appearances of objects. They succeed only when the information the interpreter derives from the picture matches that which its maker intended to convey. Consequently, they act so as to fulfill one other's interests in effective communication. In addition, there are a number of different ways in which any object can be depicted. This suggests that there are many different combinations of agents' actions in which each agent will prefer to have interpreted a mark as she has, given the way in which the other agents interpret that mark. Moreover, the problem of depicting particular objects or object types recurs. Depiction therefore meets Lewis's first four requirements for conventionality. Whether or not it is conventional, therefore, depends on whether it also meets his fifth requirement, and the salience requirement that this entails.

Whether or not these requirements are met depends on how pictures solve the coordination problem just described. There are three ways in which they could do so, each of which supports a different form of conventionalism. Before proceeding, therefore, I will distinguish the three forms of conventionalism and

identify which is at issue in the question of whether depiction exhibits language-like conventionality.

Pictures may bear some non-conventional relation to their objects. For example, they may either resemble the objects they depict, be experienced as doing so, or engage the same recognitional capacities as those objects. So long as it constrains depiction independently of any recurrent coordination problem, this relation will not itself be conventional. Nevertheless, if at least two different arrangements of marks could exhibit the relevant relation to some object and thus comprise solutions to the problem of depicting it, Lewis's third requirement for conventionality would be met and the wide adoption of one such arrangement would be conventional. Since more than one arrangement of marks can resemble an object, or be experienced as doing so, etc., particular ways of depicting an object may be conventional even if depiction involves non-conventional relations.

If depiction is partly distinguished by non-conventional relations, there will be two ways in which an arrangement of marks could solve the problem of depicting some object. It could do so solely in virtue of bearing the relevant non-conventional relation to the object, or it could do so both by doing so *and* by comprising a widely-adopted solution to that problem. In the former case, whether or not an arrangement of marks comprises a picture will depend entirely on whether it bears the non-conventional relation to the object. It may happen that it both does so and is widely-adopted as a means of depicting some object.

However, its being widely-adopted would be irrelevant to the fact that it depicts that object. In this case, then, the conventionality of a solution to the problem of depicting some object is representationally irrelevant. Let us call the view that pictures exhibit this form of conventionality *non-representational conventionalism*.

Non-representational conventionalism is not at issue in the debate concerning whether depiction exhibits language-like conventionality. For depiction to do so, depictive conventions would have to be representationally relevant, and non-representational conventionalism denies that they are.

If an arrangement of marks depicts by both bearing some non-conventional relation to an object and being widely-adopted as a solution to the problem of depicting that object, its conventionality will be representationally relevant. Which of several marks that bear the relevant non-conventional relation depicts the object will depend on which of them instantiates the relevant object-depiction convention. This view holds that depiction involves both non-conventional and conventional relations. Because it maintains that the possible conventional solutions to the problem of depicting some object are constrained by the requirement that they bear some non-conventional relation to that object, it holds that depictive conventions are non-arbitrary. Consequently, let us call it *non-arbitrary representational conventionalism*.

Non-arbitrary representational conventionalism is not at issue in the debate over whether depiction depends on language-like conventions. Linguistic

representations are widely held to lack any representationally relevant non-conventional relation to their objects. The plausibility of non-arbitrary representational conventionalism thus lies beyond the scope of this paper.

A final form of conventionalism holds that the only representationally relevant relations between pictures and their objects are conventional. This is the form of conventionalism at issue in this paper. It claims that there are no representationally relevant non-conventional relations between pictures and their objects.

This does not mean that there are no non-conventional constraints on depiction. If we are to use them, *all* systems of signs must meet those non-conventional constraints that are imposed by the limitations to human interpretative capabilities, and by our representational requirements. These constraints preclude us from forming a word from thousands of letters or pictures of different objects from marks indistinguishable to the human eye. They also require us to choose systems of signs that are sufficiently complex for the representational tasks we set for them. These constraints are not representationally relevant because they do not affect what any particular sign in a system represents.

This view holds that, like language, depiction is an *arbitrary* form of representation. There are two senses in which language is arbitrary. Firstly, among those systems of signs that meet the non-conventional constraints described above, any could be used for linguistic representation. Secondly, there

are no constraints on which of the signs in a linguistic system can be used to represent what object.[11] Let us call the view that holds that depiction is arbitrary in these two ways *arbitrary representational conventionalism*. In what follows, I will assess arbitrary representational conventionalism (from now on, simply *conventionalism*).

IV Depictive Generativity

An account of depiction must accommodate its various distinguishing features. One such feature is *depictive generativity*.[12] Roughly characterised, it consists in the fact that the ability to interpret some pictures gives one the ability to go on to interpret other pictures without any instruction about what they depict, so long as one can recognise their objects visually.[13] Furthermore, even if one has acquired this latter ability, one will be unable to interpret a picture unless one is able visually to recognise its object.[14] Let us call the ability to interpret novel pictures without any instruction as to their content *general pictorial competence.*

Depictive generativity involves both a necessary and a sufficient condition:

1. General pictorial competence and the ability visually to recognise an object suffice for the ability to interpret a picture of that object.

2. General pictorial competence and the ability visually to recognise an object are necessary for the ability to interpret a picture of that object.

However, some qualifications are needed. Firstly, the ability to interpret some line drawings may not give one the ability to interpret *any* cubist paintings, even if it gives one the ability to interpret other line drawings. General pictorial competence may not operate across pictorial systems, but may instead be relativised to particular such systems. Secondly, the inability visually to recognise an object under a particular conception need not result in the inability to interpret a picture of that object. Even if I am not able visually to recognise a medieval Greek coin when I see one, I may still be able to recognise coins, and thus partially to interpret a picture depicting such a coin. Depictive generativity holds across all consistent conceptions of a picture's object, whether these conceptions are of the object as a particular, or as type of object.

Language exhibits a feature that is apparently analogous to depictive generativity. Like that of pictures, the interpretation of sentences is generative. This is a consequence of the fact that they have a *combinatorial* semantics: their content is determined, systematically, by the contents of their parts. As Jerry Fodor argues, this means that the the ability to produce and understand some of the sentences in a language is intrinsically connected to the ability to produce and understand others, such that one wouldn't attribute true understanding of English to someone who understands what the sentence "Bob hates Bill" means, but not what "Bill hates Bob" means.[15] Understanding an English sentence in a way that exhibits a true understanding of English entails understanding the semantic rules governing the meanings of its constituent words and understanding the syntactic rules of English, such that one can understand novel

grammatical sentences formed using those words. If we construe general linguistic competence as involving knowledge of the syntactic rules of a language, we get the following characterisation of linguistic generativity:

1. General competence in a language and knowledge of the semantic rules governing the words of a sentence in that language suffice for the ability to interpret that sentence.

2. General competence in a language and knowledge of the semantic rules governing the words of a sentence in that language are necessary for the ability to interpret that sentence.

Language provides the model for the form of conventionality that concerns us here. Consequently, the fact that it exhibits something like depictive generativity might seem to show that conventionalism could readily accommodate depictive generativity. However, there are two important differences between depictive and linguistic generativity. To accommodate depictive generativity, conventionalism must accommodate both differences.

Firstly, the ability visually to recognise an object plays no role in linguistic generativity. Even if I am a competent user of English, I will not be able to interpret the English words for every object I can recognise visually. While I can visually recognise hedgehogs, I may not be able to interpret the word “ilspile”. This case is very different to that of the picture of the medieval Greek coin.

Whereas I am able to interpret that picture under some conception, I cannot get any purchase on what the word “ilspile” represents. Visual recognitional ability plays no role in linguistic generativity because the arrangement of letters that is used to represent one object is not related systematically to those used to represent others. The autonomy of the various conventions governing language means that grasping what some arrangements of letters represent will not confer the ability to grasp what others do, irrespective of whether one can recognise visually the objects they represent. Contrarily, so long as we have pictorial competence, we only need the ability visually to recognise an object under some conception in order to interpret a picture of that object. This distinction seems difficult to accommodate on the assumption that depiction, like language, is governed by a set of unrelated representational conventions.

Secondly, pictorial competence is much more readily acquired than linguistic competence. In order to learn the syntactic rules governing a language, and thus to acquire linguistic competence, we require extensive exposure to the relevant language. Contrarily, while it is unclear exactly how much exposure to pictures we need in order to acquire pictorial competence, it is clearly comparatively little.[16] Again, it is not clear how one might accommodate this difference while holding that depiction, like language, is governed by unrelated representational conventions. I will now consider how conventionalism might accommodate these two distinctive features of depictive generativity.

V Property Conventionalism

A conventionalist account might succeed in accommodating depictive generativity if it holds that depiction is governed by conventions for the depiction of *properties*, rather than conventions for the depiction of *objects*. On such a view, a picture depicts a certain object in virtue of employing conventions for the depiction of various of that object's properties, rather than employing conventions for the depiction of that object, or of that object as instantiating those properties. Because a wide range of objects may share some properties, pictures of quite different objects could employ some of the same conventions. Consequently, the ability to interpret a picture of one object could confer the ability to interpret a picture of some other object, so long as the two pictures employed the same conventions. While individual depictive conventions would operate independently, the conventions governing a particular picture would not be independent of those governing others.

To accommodate depictive generativity, such an account would need to accommodate the fact that the ability visually to recognise something is required to interpret a picture of that thing. It could explain this fact as it applies to properties by holding that the ability to recognise instantiations of the properties a picture depicts is a prerequisite for understanding the representational conventions that govern depiction of those properties. As Robert Hopkins points out, if the conventions governing depiction are limited to conventions for the depiction of colour, shape and textural properties, the ability to recognise instantiations of those properties will indeed be necessary for the ability to understand the relevant conventions.[17] As he argues, in order to understand

representations of colour, shape and textural properties, we need to deploy *observation concepts*, which are so called because the ability to recognise instances of the properties in question is necessary for possession of those concepts.

Such a view can explain why the ability visually to recognise an *object* is required to interpret a picture of that object because it holds that the depiction of any object involves the depiction of various of its basic properties. To interpret a picture depicting some combination of basic properties as a picture of a certain object, it is thus necessary to be able visually to recognise the object in question from those of its basic properties that are depicted in order to interpret the picture as depicting that object rather than any other.

This view can accommodate the fact that pictorial competence is acquired on the basis of comparatively little exposure to pictures. To do so, it need not maintain that any picture employs every depictive convention, or that all such conventions can be learnt from exposure to a single picture. It need only claim that the conventions governing depiction can be learnt through exposure to a limited number of pictures. So long as the conventions governing depiction are restricted to conventions for the depiction of basic, widely-shared properties of objects, this seems plausible. Basic colour, shape and textural properties can be shared by quite disparate objects. Consequently, pictures of very different objects may employ some of the same conventions. It should thus be possible to learn the conventions necessary to interpret a picture of one object from exposure to

pictures of quite different objects. Moreover, so long as depiction is governed only by conventions for the depiction of basic colour, shape and textural properties, it is plausible that the total number of such conventions should be sufficiently limited to be learned, and pictorial competence acquired, after exposure to relatively few pictures.

The form of conventionalism just outlined holds that the depiction of objects depends on conventions governing the depiction of their basic colour, shape and textural properties, rather than that of the objects themselves. Let us call this view *property conventionalism*. It is the only form of conventionalism that might accommodate depictive generativity. If property conventionalism fails, so too will conventionalism as a whole. I will now argue that it does fail.

VI Why Property Conventionalism Is False

According to property conventionalism, depiction is conventional because certain parts of pictures comprise widely-adopted solutions to the problem of depicting colour, shape and textural properties. For the reasons outlined at the beginning of Section III, this problem is a recurrent coordination problem. However, as I will now argue, it is not a problem that admits of a widely-adopted solution. Property conventionalism fails because it cannot meet the requirement that depictive conventions be widely adopted. It cannot do so because it cannot meet Lewis's salience requirement for wide adoption.

Property conventionalism claims that particular parts of a picture are related by

convention to the basic colour, shape and textural properties they depict. For this to be the case, for every relation between a part of a picture and its depictive content, the same conception of that relation must be salient to the makers and interpreters of pictures. That is, there must be unique conceptions both of the part of the picture that is related to depictive content, and of the depictive content to which it is related, under which that part stands out to them as bearing the depiction relation to that content. I will now argue that it is impossible if the relevant depictive content comprises basic colour, shape or textural properties.

Many pictures have parts that are interpretable as depicting things when considered in isolation from the rest of the picture. For example, those parts of a picture of a crowd that depict individual faces may be interpretable in isolation as depicting faces. Let us call such parts *pictorial parts*. Because they are independently interpretable, pictorial parts of a picture are themselves pictures.

The property conventionalist cannot claim that depictive conventions attach to pictorial parts. Because they are themselves pictures, they must exhibit the features common to all pictures. One such feature is that pictures cannot depict properties without depicting some object (be it a particular or merely an object type) as instantiating those properties. For example, a picture cannot depict the property of being a chair without depicting either a particular chair or a type of chair as instantiating that property. Likewise, a picture cannot depict the property of being red without depicting some surface or volume as instantiating that property. The bare depiction of properties is impossible. Consequently, pictorial

parts could only be governed by conventions for the depiction of objects as instantiating properties. However, as I argued above, the property conventionalist must hold that depiction is governed by conventions just for the depiction of basic colour, shape and textural properties if she is to explain depictive generativity. Only then could pictures of quite different objects employ the same conventions.

Not all picture parts are pictorial parts, however. Some are not interpretable as depicting anything when considered in isolation, but are interpretable as depicting something when considered in the context of a picture as a whole. For example, in a rudimentary picture of a face, the eye may be represented by a dot. The dot alone, however, is not interpretable as depicting an eye. It is only interpretable in the context of the picture as a whole. Just as the dot contributes to what the picture can be interpreted as depicting, namely a face with eyes, so too the picture contributes to what the dot depicts. Let us call those parts of pictures that are not interpretable in isolation, but are interpretable as depicting something in the context of other picture parts *sub-pictorial parts*.[18]

In addition to pictorial and sub-pictorial parts, pictures have *non-pictorial parts*. These are parts that are not interpretable as depicting anything either in isolation or in the context of other picture parts. For example, the individual brushmarks I see when I look very closely at the surface of a painting may not depict anything, considered either in isolation or in the context of the rest of the picture. They provide information only about the picture surface, and not about what the picture depicts. Consequently, depictive conventions cannot attach to non-pictorial parts.

They must therefore attach to sub-pictorial parts.[19]

However, a picture's sub-pictorial parts cannot in fact be governed by meaning conventions. As I argued above, the requirement that any convention be salient to its users means that there must be unique conceptions both of the part of the picture that is related to depictive content, and of the depictive content to which it is related, under which that part stands out in some "pre-eminently conspicuous respect" to pictures' makers and interpreters as bearing the depiction relation to that content. However, sub-pictorial parts do not stand out as depicting their contents, because they are not interpretable in isolation.

One might object to this that some words are not interpretable in isolation, even though their meanings are clearly convention-governed. Two types of words exhibit such interpretative context-dependence: indexicals and ambiguous words. However, they do not show that convention-governed representations can be independently uninterpretable. Moreover, the context-dependence of the meanings of a picture's sub-pictorial parts is not analogous to that of either indexicals or ambiguous words.

In his account of the meaning of indexicals, David Kaplan distinguishes between their *character*: their linguistic meaning, and their *content*: what they contribute to the propositions that sentences containing them can be used to express in a given context.[20] While indexicals' characters are independently interpretable, their contents are not. Their contents depend both on their characters and on the

contexts in which they are used. The character of ‘you’ is something like ‘the person to whom this utterance is addressed’. Its content in a given context is whoever it refers to in that context. Linguistic conventions govern the independently interpretable character of indexicals, rather than their context-dependent content. The context-dependence of the content of indexicals is therefore not a counter-example to the claim that the conventionality of a representation requires its independent interpretability.

Moreover, the property conventionalist cannot construe the context-dependence of sub-pictorial parts as analogous to that of indexicals, and therefore cannot explain it in the same way. An indexical has an indefinite number of possible contents. While its character tells us what features of context are relevant to determining its content, it does not tell us what its possible contents are. For example, the character of ‘you’ tells us to look at the person to whom it is addressed in a given context to determine its content in that context, but it does not constrain to whom it could be addressed. If property conventionalism is to accommodate depictive generativity, it cannot allow sub-pictorial parts to have an indefinite number of possible contents. It requires constraints on what each sub-pictorial part can depict to apply across all possible contexts of interpretation. Otherwise, the ability to interpret such parts in the context of some pictures could not yield the ability to interpret them in the context of novel pictures.

Ambiguous words have at least two different characters. In any disambiguating context, the content of such a word will be identical to one of its characters.

Knowledge of the different conventions governing these characters therefore enables us to tell what the possible contents of an ambiguous word are in any given context. For example, knowledge of the conventions governing ‘bank’ tells us that its characters are ‘financial institution’ and ‘the ground alongside a river’. The word is independently interpretable as having both characters, and therefore as having two possible contents.

Unlike indexicals, therefore, the characters of ambiguous words constrain their possible contents. One might think the property conventionalist could accommodate depictive generativity by construing the context-dependence of sub-pictorial parts as analogous to that of ambiguous words. If a sub-pictorial part had a limited number of possible contents, one could learn each of those contents after exposure to a limited number of pictures, and thereby acquire the ability to interpret that part in the context of other picture parts. However, this suggestion ignores the independent interpretability requirement. As with indexicals, the fact that linguistic conventions govern characters rather than contents means that the context-dependence of the contents of ambiguous words is no threat to this requirement. Consequently, if sub-pictorial parts were analogous to ambiguous words, their different characters, and consequently their possible contents, should be independently interpretable. Some picture parts, such as those depicting Necker cubes, are independently interpretable as having different possible contents. However, because they are independently interpretable, they are pictorial parts and therefore cannot depict basic properties as required by property conventionalism.

One might also object that since, as a matter of fact, we can interpret colour-reversed pictures, my claim that conventions cannot govern the depiction of properties must be false. It is true that, so long as we know which colours stand for what other colours, we can interpret a picture of a red object as representing an object that is, say, green. However, we can only do this because we can apply representational conventions to pictures *on the basis of* their depictive content. That is, we can stipulate that pictorial parts that depict objects of one colour represent objects of some other colour. To interpret pictures using such conventions, we must have prior knowledge of what they depict. Colour-reversed systems of representation are not depictive and therefore do not show that the depiction of properties can be convention-governed.

Finally, one might argue that my argument contradicts psychologists' claims that the interpretation of line drawings involves the application of interpretative rules. John Willats argues that, just as the marks on a page of writing represent letters, so too depictive marks represent "picture primitives".[21] As letters form words that represent aspects of an object, picture primitives represent specific aspects of visible scenes called "scene primitives". As words combine to form descriptions of objects, these primitives then combine to represent whole objects. This account may seem to vindicate conventionalism. However, two considerations show that it does not. Firstly, according to Willats, the interpretation of any picture primitive is context-dependent, since it depends on the other picture primitives with which it is combined.[22] Secondly, Willats denies that the representational

content of pictures is arbitrary. Instead, he claims that pictures bear a non-conventional relation to their objects: they depict them by exploiting the human visual system's mechanisms for detecting features of those objects from the light they reflect.[23]

Property conventionalism is false. To be true, the depictive conventions it posits would have to attach to either pictorial or sub-pictorial parts. They cannot attach to pictorial parts since such parts cannot depict basic colour, shape or textural properties alone, but instead must depict some object as instantiating those properties. Moreover, although sub-pictorial parts can depict basic colour, shape and textural properties, depictive conventions cannot govern such parts since they are not independently interpretable. They are therefore not salient to pictures' makers and interpreters as bearing the depiction relation to such properties.

VII Conclusion

An adequate account of depiction must be consistent with pictures having those features that are distinctive of depiction. One such feature is depictive generativity. Conventionalism could accommodate depictive generativity only by holding that pictures are governed by conventions for the depiction of basic colour, shape and textural properties. However, pictures cannot be governed by such conventions. For depictive conventions to meet Lewis's salience requirement, the picture parts to which they attach must be independently interpretable. Any independently interpretable part of a picture is itself a picture.

Therefore, it cannot simply depict colour, shape or textural properties, but must always depict some object as instantiating those properties.[24]

---

[1] See Hopkins, R. (1995). "Explaining Depiction." Philosophical Review **104**: 425-455.He argues that an account of depiction must *explain why* pictures have these features. I agree, but will retain the weaker consistency requirement for current purposes.

[2] Goodman, N. (1976). Languages of Art. Indianapolis, Hackett Publishing Company. pp230-31.

[3] See, for example, Bryson, N. (1983). Vision and Painting: The Logic of the Gaze. New Haven, Yale University Press. Although some philosophers deny that Goodman is in fact a conventionalist (see, for example, Lopes, D. (1996). Understanding Pictures. Oxford, Clarendon Press. p65), his commitment to the arbitrariness of depiction makes it clear that he is. Goodman himself denies conventionalism regarding both depiction *and language* because he believes it entails the claim that they are rule-governed. He argues that "(t)he way pictures and descriptions are thus classified into kinds [according to what they represent], like most habitual ways of classifying, is far from sharp or stable, and resists codification" (Goodman, op. cit. p23. The text in square brackets is mine). He also writes that understanding a picture "is not a matter of bringing to bear universal rules that determine the identification and manipulation of its component symbols" (Goodman, N. and C. Z. Elgin (1988). Reconceptions in Philosophy and Other Arts and Sciences. London, Routledge. p110). However, the claim that depiction is conventional does not entail the claim that it is rule governed. I am interested to determine whether depiction is conventional in the way that language is. Goodman's claims about the arbitrariness of depiction suffice to show that he thinks it is.

[4] For example, the following criticise Goodman's characterisation of depictive conventions, rather than his claim that depiction is conventional: Wollheim, R. (1970). "Nelson Goodman's Languages of Art." Journal of Philosophy **LXVII**(16): 531-539, Savile, A. (1971). "Nelson Goodman's 'Languages of Art': A Study." British Journal of Aesthetics **11**: 3-27, Walton, K. (1974). "Are Representations Symbols?" Monist **58**: 236-254, Robinson, J. (1979). "Some Remarks on Goodman's Language Theory of Pictures." British Journal of Aesthetics **19**: 63-75.

[5] See, for example, Hopkins, R. (2003). "Perspective, Convention and Compromise." Looking Into Pictures. H. Hecht, M. Atherton and R. Schwartz. Cambridge, Massachusetts, MIT Press**:** 145-164.

Hopkins argues against the conventionality of depiction. However, he assumes that depiction depends on non-conventional relations between picture and object (on the preservation of "visible figure"). He therefore fails to address the claim that depiction, like language, depends on *arbitrary* conventions. Crispin Sartwell construes the fact that pictures are generative (see Section III) as an argument against conventionalism (Sartwell, C. (1991). "Natural Generativity and Imitation." British Journal of Aesthetics **31**: 58-67. p63). However, as I will argue, the generativity of depiction does not automatically preclude its conventionality. Gregory Currie rejects the conventionalism of film (See Chapter 4 of Currie, G. (1995). Image and Mind: Film, Philosophy and Cognitive Science. Cambridge, Cambridge University Press). However, as I argue in Note 19, his argument fails for both film and depiction.

[6] See, for example, Quine, W.V.O. (1936). "Truth by Convention" in Philosophical Essays for A.N. Whitehead, edited by O.H. Lee. New York, Longmans. pp90-124.

[7] Lewis, D. (1969). Convention. Cambridge, MA, Harvard University Press.

[8] For example, each may be classified as a solution that has been adopted to the relevant coordination problem.

[9] Lewis, D. (1969). Convention. Cambridge, MA, Harvard University Press. p38

[10] Neither the salience requirement, nor any of Lewis's other requirements entails that conventions are rule governed in the sense to which Goodman objects. He accepts that depiction and language involve habitual plans of correlation according to which referents are ascribed to descriptions and depictions. There need be no rule, independent of the plans of correlation to which individual agents adhere, that determines what solution is adopted. For a solution to a recurrent coordination problem to be conventional, it is sufficient for there to be habitual practices according to which the different agents involved in a decision-making situation always knowingly adopt the same solution to the problem. Since Goodman maintains that depiction is governed by such practices, or "plans of correlation", he is a conventionalist in Lewis's sense (Goodman, N. op. cit. p39-40).

[11] We may find it expedient to use certain words within a language to refer to certain objects. For example, it might be useful to choose shorter words for common objects. However, there is no reason why a language could not, in principle, use long words for common objects.

[12] The definitive discussion of this feature is in Schier, F. (1986). Deeper Into Pictures. Cambridge, Cambridge University Press. pp 43-55.

[13] I use the term 'interpret' to refer solely to the process of determining what a picture depicts.

[14] It follows from the centrality of the ability visually to recognise objects to picture interpretation that only visible objects can be depicted. This feature is noted in Hopkins, R. (1995) op. cit.

[15] Fodor, J. A. (1987). Psychosemantics. Cambridge, Massachusetts, The MIT Press. p149.

[16] See Hochberg, J. and V. Brooks (1962). "Pictorial Recognition as an Unlearned Ability: A Study of One Child's Performance." American Journal of Psychology **75**: 624-628. This study found that a nineteen month old child whose access to pictures had been severely restricted, and who had never received any instruction in interpreting pictures was nonetheless able spontaneously to interpret pictures of a wide range of objects.

[17] Hopkins, R. (1997). "El Greco's Eyesight: Interpreting Pictures and the Psychology of Vision." Philosophical Quarterly **47**: 441-458. p450.

[18] The distinction between pictorial and non-pictorial parts is based on Schier's distinction between iconic and sub-iconic parts. See Schier, F. op. cit. p70.

[19] Failure to recognise the distinction between pictorial, sub-pictorial and non-pictorial picture parts undermines one argument against conventionality. Gregory Currie denies that film is conventional on the basis that, unlike language, it is not comprised of independent, meaningful parts (see Chapter 4 of Currie, G. (1995). Image and Mind: Film, Philosophy and Cognitive Science. Cambridge, Cambridge University Press). Because it has no such parts, he argues, there can be no rules of composition enabling the representationally relevant parts of a film to be combined to form further meaningful parts in the way that words are combined to form sentences. He argues that film has no independent, meaningful parts because every temporal and spatial part of a film image is meaningful "to the limits of visual discriminability" (Ibid. p 130). He holds that this is true even of parts of an image that lack a discriminable internal structure, since they represent their objects as uniform.

Currie's argument is applicable to both film and depiction. However, it fails in both cases. All pictures are ultimately comprised of parts that have no meaning at all: their non-pictorial parts. Consequently, they are not meaningful in the way he claims. The smallest meaningful parts of any picture are its sub-pictorial parts. While such parts do not have meaning when instantiated in isolation, they nonetheless have a meaning that is independent of that of the picture's other parts. Currie provides no reason why sub-pictorial parts could not be governed by meaning conventions, and be combined using rules of composition to form

independently meaningful pictorial parts.

[20] Kaplan, D. (1989) "Demonstratives", in J. Almog et. al. (eds) Themes from Kaplan. Oxford. Oxford University Press.

[21] Willats, J. (1997). Art and Representation: New Principles in the Analysis of Pictures. Princeton, Princeton University Press.

[22] Ibid. p112-119.

[23] Ibid. p146.

[24] I am extremely grateful to Tim Bayne for discussion of these issues and for his comments on earlier versions of this material. I am also grateful to Jordi Fernandez and to the anonymous referees of this journal for their comments and criticisms.