

An Ontology of Minimal Creativity

Abstract

A philosophical theory of creativity reasonably begins with an ontology of creativity, with the ultimate goal being an answer to the basic question ‘What is creativity?’. This paper gives an ontology of minimal creativity, arguing, first, that creativity is fundamentally a property *of* processes and, second, that creativity is an artefactual extrinsic property typed coarsely by two features: being the non-accidental product of agency and being novel. This ontology bears implications for a semantics and epistemology of creativity. Creativity statements are always, even when surface language presents otherwise, statements about a process. And one identifies and experiences creativity only when one possesses the relevant background knowledge. In one important sense, creativity is not something we *just see*.

Creativity is a generally important phenomenon, occurring across human culture and context. The potential philosophical importance of creativity mirrors the breadth of the phenomenon. In philosophy of science, an interest in theory construction and change will imply a concern with creativity at group and individual levels. In philosophy of language and linguistics, issues concerning language acquisition and generativity may partly concern creativity. In philosophy of mind and cognitive science, an interest in concepts, problem solving, or action, among other issues, will benefit from an analysis of creativity. And, perhaps most obviously, creativity is a virtue in art making, and so philosophical aesthetics does well to provide a theory of creative art practices. Creativity is thus a widely important *explanandum* for philosophy. It is not, however, widely explained in philosophy. In fact, relative to other phenomena of comparable breadth and interest—concepts, action, perception, to name three—it

receives almost no philosophical attention.¹

There are presumably many factors that contribute to this relative theoretical neglect. One obvious culprit is an over-emphasis on high-level creativity or genius—on the minds and acts of our Picassos and Einsteins. The attraction here is understandable: these are fascinating minds and their acts will leave marks for ages. But if a general phenomenon is the target, then extreme versions of that phenomenon are often not the best places to start. Instead, one might start with the basics, identifying some minimal features of the phenomenon and answering some familiar metaphysical and epistemological questions. The broad goal, then, is to approach a conceptual foundation both for varied philosophical interests and for empirical work that may be done on creative thought and behaviour.

A philosophical theory of creativity ought to provide an ontology of creativity: answers to questions about the kind of thing creativity is. This ontology should be largely consistent with our practices of attributing creativity, even if it provides some normative guidance or constraints on those practices. It thus should inform a semantics of our talk about creativity and provide some epistemic constraints on identifying instances of creativity. A promising order of analysis would go as follows. In attempting to answer the basic question ‘What is creativity?’ one may consult practices of creativity attribution: the contexts and circumstances in which we say ‘that is creative.’ This basic question divides, at least, into two. What, fundamentally, is creativity a

¹ This is not to imply, of course, that there has been *no* philosophical analysis. Some recent examples include: Barsalou and Prinz 2002; Boden 2004; Carruthers 2002, 2007; Gaut and Livingston 2003; Novitz 1999.

property or attribute *of*? Second, what *sort* of property or attribute is creativity? And once one has answered the ontological questions, one can infer some semantic and epistemic features of or constraints on creativity attribution. Thus the metaphysics, semantics, and epistemology of creativity are reflexive: our rational practices will inform our ontology and our ontology will imply meanings and appropriateness conditions for creativity attributions.

Consideration of our attributive practices reveals that creativity is fundamentally attributed to processes. And it is an artefactual extrinsic property that is minimally typed by non-accidental dependence on agency and some degree of novelty. It characterizes things but does not, as sortal properties do, particularize them. This minimal ontology of creativity alone bears semantic and epistemological consequences. Creativity is always attributed to a process and thus attributions whose surface form indicates attribution to persons or products must be interpreted accordingly. Moreover, creativity is not something we just see or hear minus certain background knowledge. That is, creativity is not revealed *merely* by the perceptible properties of an object or event to which we may attribute creativity. Rather, an identification of creativity requires background, contextual knowledge. The basic ontological question is answered with a simple ontology in §I below. The semantic and epistemological implications are teased out in §II that follows.

§I. Ontology

II.1 What is creativity a property of?

We attribute creativity to a variety of things. We call persons creative. We talk of creativity in terms of processes, where the referent is some importantly related set of thoughts and/or action. Finally, we attribute creativity to the things that persons make by such processes—the products—talking of artworks or theories as themselves being creative. A theory of creativity should determine if one of these categories of subject of attribution is primitive. Determining whether one of these subjects is primitive in turn provides insight in answering the basic question: what is creativity?

Consider our practices of creativity attribution. In particular, consider how one justifies an attribution of creativity. Pointing to one of Pollock's action paintings, *White Light*, Maggie says to Phil, "That's creative." Phil receives the attribution with scepticism, and solicits a justification. Maggie might begin by simply invoking features of the work. It is more likely, however, that Maggie's justification will invoke features of Pollock's generative process. She may describe how Pollock would drip, throw and splash paint onto a giant canvass, spread on the floor so he could stand on it, dance across it, "be in it"; or his use of sticks, palette knives trowels, and basting syringes to apply and manipulate paint. She might suggest features of Pollock's thought process: he is often quoted as desiring the work to serve as an expression of the artist's gestures and techniques. He claimed to go into a kind of trance when painting, obliviously leaving handprints, footprints, and cigarette butts in his wake. Maggie may mention the historical context, citing the obvious influences of, but departures from, cubism and surrealism, citing the novelty of Pollock's thoughts and actions in these contexts. This is just one

example, but it is reasonably assumed that this is the kind of account one rationally gives in justifying an attribution of creativity. Given this assumption, one may generalize from features of Maggie's response.

First, note the variety and scope of features invoked: events, thoughts, actions, and contextual factors. This variety motivates the following negative claim: there is no one locus for creativity. Maggie is not saying just of the formal qualities of the object affixed to the gallery wall or just of Pollock's bodily movements or just of Pollock's thoughts or just of the historical context that *that* is the creativity. She has invoked all of these features to explain her initial attribution. Maggie has attributed creativity to a process. One would tell a similar story no matter if the attribution is to a set of actions—say, Bacon's development of the Scientific Method—or to some particular artefact—say, Monet's *Impression, Sunrise*. It seems that attributing creativity to actions and works is an attribution to the generative processes, of which the actions are a part and the works are products (and also, in a sense to be clarified, parts).

Here, then, is a simple and plausible thesis: creativity is, fundamentally, attributed to processes. This process thesis of creativity is motivated by practices of attributing, and justifying attributions of, creativity. The thesis is further motivated by considering arguments against the alternatives: namely, that creativity is fundamentally an attribute of persons, or fundamentally an attribute of products.

Persons

One might take persons to be basic for predication of creativity. Geniuses and masterminds are intriguing and mysterious, and so conversations about creativity will sometimes centre around them. However, our practices indicate that creativity is not solely or even primarily a property of persons. Consider what we value when we value an instance of creativity. It makes little sense to suggest that we value the genius for genius's sake. Geniuses do not, qua geniuses, possess intrinsic value. Rather, one values the brilliant scientist either for the theory she concludes with or the methods by which she comes to that theory, or both. In fact, the point is more fundamental than the one about value. We simply do not attribute creativity to a person if that person does not act in ways that are of some interest or produce some objects or events that are of some interest (to put both points most neutrally). Without a process or a product, there is no creative person.

Products

Perhaps, then, creativity is fundamentally attributed to products: it is a property of products. The strongest version of this view occurs in *anti-intentionalism* in philosophical aesthetics, argued by Monroe Beardsley and others through the middle of the 20th century. Beardsley argued that the value of artworks consists solely in the perceptible properties of the manifest object and the experience of those properties. Thus the value of an artwork for Beardsley is “independent of the manner of production, even of whether the work was produced by an animal or by a computer or by a volcano or by a falling slop-bucket” (Beardsley 1965:

301). The central commitment of anti-intentionalism invites a dilemma. Either creativity is not something we value in artworks (or, put another way, is not relevant to how we value artworks) since it consists in features of production and context, or it is something we value but must be located in the product only. Practices of appreciation and criticism imply that both horns bear undesirable consequences: creativity is something we value and, if the discussion to this point is sound, it is attributed to features beyond the manifest product. The anti-intentionalist acknowledges our practice and so avoids the problems of the first horn, but gets stuck on the second horn in virtue of his anti-intentionalism. But this too is at odds with our practices: attributions of creativity are not exclusive to pure products. The only remaining option for the anti-intentionalist as such is an outright denial of actual practices.² The view generalizes to non-artistic creativity. One might hold that creativity is only attributed to the end result: the scientific theory, the novel public space, the unexpected move in a basketball game. The same dilemma challenges the generalized pure product view. And here again our practices prevail: we do value creativity in science and all other contexts, and not just for some terminating product. A pure product view of creativity—one that claims that creativity is a property solely of products—fails.

This result is of little surprise. If the opening example teaches anything, it is that in distinguishing and discussing a creative x where x is an object or event (such as an artwork), we invoke properties not intrinsic to x : relational features like the means of producing x , the context of x , and the maker of x . Perhaps a better product view says this: creativity is fundamentally

² See XXXX for extended discussion of an anti-intentionalist, product theory of creativity and how it implies trouble for anti-intentionalism generally.

attributed to objects, which as products, possess a package of relational or extrinsic properties.³ Thus in attributing creativity to *White Light*, Maggie is attributing creativity to the particular in the gallery, where that particular possesses the following properties (in virtue of which it is creative): being the result of Pollock's paint dripping, flinging, and splashing; being walked, crouched, and perhaps danced on by Pollock; being the object of various Pollock desires, beliefs, imaginings and other thoughts; being cubist-influenced and surrealist-influenced; being a departure from cubism and surrealism; being novel relative to a broad art-historical context (and many more narrow contexts). Perhaps all of this is plausible: once one accepts relational properties into her ontology, it is difficult to draw a principled line between legitimate relational properties and dubious or gerrymandered ones.

However, any theory of creativity, including a refined product view, must also accommodate attributions that are not (ostensibly) made to one particular object or event. In order to accommodate the varied features of our attributive practice, one will need to include a number of additional properties and some of them may be problematic (at least as being properties *of* a single object). Intuitively, in justifying an attribution of creativity, like Maggie's attribution to *White Light*, one might mention many events and properties that pre-exist the finished work (in fact, this is trivially true). For example, Pollock had thoughts and plans about his next painting before beginning it, he may have made preliminary sketches, he may have

³ This approach comports well with certain ontologies of art, for example, those that take works of art to be partly constituted by the creative process that results in a finished product. Gregory Currie takes artworks to be event-types (Currie 1989); David Davies takes artworks to be event-tokens (Davies 2004).

chosen the canvass out of a number of possibilities. These seem to be constituents of the relevant creativity. That is, these features are plausibly invoked or mentioned in the context of the Pollock work as an explanation of Pollock's creativity (and this is true even if such features are not recognized as such by all appreciators: more on this below in §II). If on the product view, creativity is a property only of the finished product inclusive of its extrinsic or relational properties, then one must somehow accommodate the constituents just mentioned and presumably by specifying properties *of the product* as follows: being preceded by a number of (*de dicto*) thoughts and intentions, being preceded by preliminary sketches, being the result of a selection of canvasses. Indeed, the product theorist will also need to specify intrinsic properties to accommodate the temporal disparity between the generating process and the finished product, such that *White Light* has the *property of* being started before it is finished.

Some of these putative properties may seem awkward. And one may question whether one object—a painting, a sculpture, a theory—can be reasonably supposed to have all of these properties and more. However, commitment to these properties and their possession by products may be ontologically acceptable, most especially if such commitments are necessary to accommodate our practices and achieve a workable ontology. The refined product view is not the only option, however. The process thesis offers another alternative, and one that both more naturally accommodates our practices of creativity attribution and achieves a more conservative ontology. It begins by denying the premise of a product view: it denies that there is a singular, spatio-temporal locus of creativity. Up to this point, this denial has been motivated by

consideration of attributive practices. Brief consideration of the ontology of processes and events further motivates the process thesis.

Processes

Metaphysical theories of processes are scarce. Least controversially, processes are understood to be composed of parts and to be temporally indexed. Conjoining these two features implies that processes are roughly structured in phases or stages. Moreover, processes are particulars (or, at least, there are particular process tokens). Metaphysical theories of events abound. And a few standard features of different categories of events help to characterize creative processes.

We appreciate creative processes both in terms of the activities that make up proper parts of the process and as resulting in certain accomplishments. Metaphysicians since Aristotle recognize both general categories of event.⁴ Both accomplishment and activity terms, as contrasted with achievement and state terms, possess continuous tenses. *Accomplishments*, as contrasted with *activities*, culminate in a terminus that justifies predication of the accomplishment term. We do not say that ‘S has baked a cake’ until the cake is baked. We may, however, predicate the corresponding activity term: we may say that ‘S is baking a cake’ at any point during the activity. Accomplishments are, by additional contrast with activities, non-homogeneous: accomplishment terms do not appropriately apply to any sub-part of the whole event. Running is an homogeneous, non-culminating activity. To run a mile, by contrast, is a non-homogeneous, culminating accomplishment; it is not until the mile mark is reached that the

⁴ Aristotle, *Metaphysics*, Book IX, Chapter 6.

“climax casts its shadow backwards, giving a new color to all that went before” (Vendler 1957: 146).⁵

Processes fall neatly into neither category of event. They may or may not involve a terminus, and with respect to homogeneity, they are describable by both event terms. Building a tree fort is a process. On the one hand, we cannot say of any part of the process that ‘S has built a tree fort’; we say this only at the time of accomplishment, when we have a tree fort. On the other hand, we can say of any part of the process that ‘S is building a tree fort.’ Prior to the output, S is engaged in the activity of tree fort building. In a process like building a tree fort, the accomplishment is important to the activity. Planning structural details, gathering and binding materials, selecting tools, and so on, constitute building a tree fort only if a tree fort is the target outcome. Otherwise, they are just a disjointed collection of activities.

The activity and accomplishment event categories are illuminative while not definitive of processes. These categories provide some distinctions useful for characterizing *creative* processes in particular. A creative process is continuous (in the sense that it spans some segment of time) and generally proceeds towards some terminus, namely, creation of some artefact or event. The end thus “casts its shadow backwards” on to the generative process. By the same token, we appreciate the activities—the thoughts and actions—of the process and how those components are structured or organized when attributing creativity, even when we do so *through* the terminus. This is not to suggest that creativity is appreciated through or in the product only. This is the mistake of a product view. Consider again the Pollock example. Maggie may have

⁵ More carefully, ‘running’ is a term for an homogeneous, non-culminating activity; ‘to run a mile’ is a term for a non-homogeneous, culminating accomplishment. The former activity term does not appropriately apply to the latter accomplishment; the latter accomplishment term does not appropriately apply to the former activity. See also Ryle 1949.

initially pointed to *White Light* in uttering ‘that is creative’. But appreciation and attribution of creativity do not depend upon the presence of some finished artefact. Two claims need to be distinguished here. First, we do attribute creativity without attributing it to some finished product; creativity attributions can be made independent of a singular, terminating artwork, theory, or other artefact. Maggie described Pollock’s techniques, motives, bodily actions, and the surrounding art-historical context. This description equally supports an attribution of creativity to Pollock’s process or some proper part thereof as it does an attribution to a particular work. Now of course in this example, Maggie’s report that ‘that is creative’ is one where the demonstrative refers (partly) to the artwork. But the features invoked—namely the various features of Pollock’s generative process—imply that the property is attributed to more than just the artwork. This is the second claim: in situations such as these, the process is appreciated through what it accomplishes, the finished artwork. In all creativity attributions, creativity is a property of the product only insofar as the product is a part of the process from which it results.

Creativity is, fundamentally, a property of processes. This thesis best explains the ways and variety with which we attribute creativity. Most especially it explains the reasons and justifications we offer in support of an attribution, be it an attribution (ostensibly) to a work of art or a scientific theorem, a person, or a series of thoughts and actions. It also commits to a relatively conservative ontology: processes are understood (at least for a start) in terms of categories of events. There is room here to finesse the characterization of processes: as concrete particulars, as event-types, as event-tokens, and so on. However one’s metaphysics dictates these choices, it is a safe assumption that processes, as minimally characterized above, are a legitimate category of thing. This is the only ontological commitment the process thesis needs.

II.2 What type of property is creativity?

The question about types must be parsed into two. First, how is creativity typed as a property? What are the conditions for possession of creativity? And second, with that answer in hand, what sort of property is creativity? These questions are taken in turn. As with the ontological question from the previous section—what is creativity fundamentally a property of?—questions regarding the nature of creativity as a property will be approached through analysis of our linguistic and conceptual practices. This strategy is justified in either or both of the following ways.

Plausibly, creativity is an artefactual kind. There are no creative *F*s without human intentions and, likely, social practices surrounding those intentions.⁶ A number of persuasive arguments have been made that artefactual and social kinds *qua* artefactual and social kinds are ones to which we have a privileged epistemic access: since our concepts are partly constitutive of artefactual kinds we are in a better position to provide theories of those kinds (Thomasson 2003, 2007; see also Searle 1995). Assuming the artefactuality of creativity, an analysis of our linguistic and conceptual practices will provide reliable insight into the nature of creativity since the latter is partly constituted by the former.

A second and more general reason for analyzing creativity by analyzing practices is this. One plausible, and ontologically neutral, way to think about properties is in terms of the role/s

⁶ The present analysis thus targets human creativity. But if non-human animals, computers, and other organisms or systems can be creative, then some of the claims here may be overstated (e.g. computer creativity as such will not depend upon *human* intention). According to the analysis that follows, if such systems and their behaviours are to count as creative, they will have to be agents in some minimal way. And insofar as agency involves intentional action, a creative animal or computer will have to act with intention.

that we assign them in our theoretical and ordinary practices. Plausibly, properties are used to identify resemblances between distinct particulars. This is true as much in circumstances of categorizing and sorting the objects and events we encounter in everyday life, as it is in metaphysical discussions about the *problem of universals* or *One over Many*. As David Lewis suggests, a “version of the property role ties the properties more closely to the meanings of their standard names, and to the meanings of the predicates whereby they may be ascribed to things” (Lewis 1986: 56). On this basic conception of properties, one analyzes a property by analyzing the ascriptions we make using the relevant predicate and the conceptual constraints on those ascriptions.

Both recognition of the artefactuality of creativity and the general property-as-role conception motivate the analytical approach taken here: theorize creativity by consulting conceptual practices. Both motivations are independently plausible but since either one is sufficient only one need be granted. Assume, then, that an analysis of creative properties is an analysis of our concept of creativity. The most minimal of conceptual analyses of creativity will reveal that creativity is a compound property. Attribution of creativity implies attribution of more basic properties. (And of course, these attributions may themselves imply attributions of even more basic properties). The following analysis identifies and clarifies two of these more basic properties: namely, being the product of agency and being novel.

Agency

Human creative processes involve thought and action. Some action theories analyze thoughts as actions. Other theories make thought a necessary condition on actions, since action

involves intention. A theory of creativity can avoid much of the detail of this debate, committing only to the claim that creativity necessarily involves an agent who thinks and acts intentionally. Depending upon one's theory of action, this claim might just be one about thoughts (if thoughts are the essential component of actions) or one about actions (if thoughts are actions). What is crucial is that the following intuition is accommodated: we do not attribute creativity in the absence of a responsible creator. An attribution of creativity implies an attribution of agency. Or more precisely, an attribution of creativity implies attribution of the property 'being a product of agency.'⁷

Consider the following:

- (1) That thunderstorm is creative.
- (2) The Special Theory of Relativity is creative, but no one is responsible for it.
- (3) The design of *La Sagrada Familia* is creative, but no one is responsible for it.

Each of these locutions is silly, but diagnosing their silliness is instructive. The silliness consists in misunderstandings of 'creativity'. Creative processes, and the objects and events they produce, involve intentional agency. A thunderstorm may be breathtaking or ominous, but never creative. The theory of relativity is an artefact of scientific effort. And designing *La Sagrada*

⁷ One may want to weaken the agency condition (as it is developed here). If, for example, one thinks that non-human animal or computer creativity is possible, then an agency condition might be construed more broadly. For example, one might take agency just to require autonomous action which involves, minimally, behaviour mediated by internal mechanisms of a system and some degree of input/output flexibility. So an organism or system is an agent so long as elements or mechanisms internal to the system can produce varying outputs given any particular input (see XXXXa, XXXX). As already suggested, the analysandum for this paper is human creativity, which calls for a richer notion of agency. This will be assumed for the remainder of the discussion.

Familia involved the thoughts and actions of an architect. (1) through (3) misuse ‘creativity’ by failing to acknowledge the dependence of creativity on agency. (1) attributes creativity where there is no agent; (2) and (3) attribute creativity but deny agency simultaneously.

One condition, then, on creativity is this:

agency: an x is creative only if x is the product of agency.

‘Product of agency’ is a technical term, meant to pick out a class of actions and consequences that require responsible agents, that is, those that result *non-accidentally* from agents. Creative things are things that we praise. We praise or blame a person’s actions only when those actions result in a non-trivial way from that person’s agency; we praise persons for performing well, making well, doing well. We do not reasonably praise persons who accidentally do something well, even if we may appreciate benefits that result from the accident. One can think of many scenes from Charlie Chaplin films that exploit this conceptual feature of praise (or blame) for dramatic effect. The Tramp accidentally saves the day, is praised by certain ignorant members of the situation, and then bashfully takes credit for the save. The irony consists both in the fact that we know what the praisers don’t know and that we know that the Tramp knows that he shouldn’t be praised for desirable consequences of a mere accident. ‘Non-accidental’ is a placeholder for a desideratum to be met by an appropriate theory of action. The desideratum can be understood in terms of the following two worries about *agency*.

Agency, on most accounts, requires intention. Does this mean that to be an agent of an event *e* one must intend to bring about, specifically, *e*? To be the agent responsible for the death of Polonius, must Hamlet intend to kill Polonius? As Davidson recognized, attributions of intention are opaque: it is true that ‘Hamlet intended to kill the man behind the arras’, but it is

false that ‘Hamlet intended to kill Polonius’, where Polonius is in fact the man behind the arras. Although we would deny that Hamlet intended to kill Polonius, under that description, we cannot deny that Hamlet did kill Polonius. So intention is a semantic *criterion* for agency, while the *expression* of agency is extensional. “[A] person is the agent of an event if and only if there is a description of what he did that makes true a sentence that says he did it intentionally” (Davidson 1980a: 46). Hamlet is the agent of Polonius’s killing since it is true that he intended to kill the man behind the arras.⁸ This addresses one concern: agency may require intention, but not in so strict a way that Hamlet is not the killer of Polonius.

An opposite worry regards the weakness of an agency condition. Given a few common though not uncontroversial assumptions, agency comes easy. I don’t intend to splatter bits of paint in an aesthetically interesting array of colour, but I do intend to move five heavy, uncovered cans of paint from one side of my studio to the other in one trip. The Davidson/Anscombe thesis of action identification says that ‘if a person *F*s by *G*ing, then her act of *F*ing = her act of *G*ing.’ Coupling this thesis with Davidson’s semantic criterion for action implies that I am the agent of the aesthetically interesting splattering.⁹ In spite of my better intentions, the splattering of paint is my action. Davidson takes this to be the appropriate result: accidents and mistakes, even if we don’t anticipate or much like them, are still our actions. But I am responsible for the splattering in a relevantly trivial way: the splattering was merely an accident resulting from my clumsiness. An observer of the series of events would not reasonably praise me for the splattering since the event, or at least the relevant results, are only loosely

⁸ In alternative terminology, the killing of Polonius is an act of Hamlet’s (Wilson 1989:89).

⁹ See Anscombe 1959, Davidson 1980b; see also Wilson 1989.

connected with my intentions. So, an x may depend upon the agency of some A , but this alone does not make x the non-accidental product of A .

The concern about the strength of the agency condition makes salient the importance of intention. The concern about the weakness of the agency condition makes salient the importance of how actions are (or are not) to be distinguished from their effects. It further reveals that agency of an event does not guarantee praise or blame for the event or its effects. Here as always, a job for the action theorist is to determine the role of intention in the assignment of agency and to provide principles for individuating actions. An agency condition for creativity aims to pick out just those thoughts, actions, and results that are the non-accidental products of agency. ‘Products of agency’ is thus silent on the general ontological issue of action identification: whether Coltrane’s playing of a series of notes is identical with certain movements of his body is orthogonal, so long as the playing is, non-accidentally, an *act of* Coltrane’s. What is important is that the connection between the agency and the event (or result) in question is rich enough to warrant praise or blame. No interest in creativity is an interest in accidental action.

‘Non-accidental’ thus placeholds the work to be done by a general theory of action. The challenge is to say what additional feature an event must have to be one that is non-accidentally dependent upon the acting agent. The challenge may be met in a number of ways but needn’t be met here.¹⁰ The two scope worries considered above frame the desideratum: an agency condition

¹⁰ One might make the event or results counterfactually dependent upon an agent and her intended action or plan. One might make appeal to an ideal observer, such that if the observer would withhold praise or blame from the thought or action then it at most trivially depends upon agency. One might stipulate that the event or results cannot be the result of any (or too much) luck. One might make the exercise of some relevant skill or knowledge a condition on non-accidental action. For some relevant theories of action, see Ginet 1990: 72-89; Harman 1976; Wilson 1989: 88-167; and the collection Mele 1997.

on creativity must strike a balance between requiring of creativity that an agent intends to get precisely the results that she does get (which is overly exclusive) and allowing for accidental performances of creative action (which is overly inclusive).

Novelty

Creativity involves novelty. A creative x is new or different in some relevant way. Maggie's justification of her creativity attribution appealed to the ways in which elements of Pollock's generative process—his thoughts, actions, and the results they produced—were different from some relevant (features of a) context. This is necessarily true of how and to what we attribute creativity. Creative objects and events are novel objects and events:

novelty: an x is creative only if x is novel.

Plausibly, novelty is a relational property. An x is novel only relative to some comparison class C . C might be some class of human culture, some system, population, conceptual space, or social context, or an individual agent's behaviour. Novelty *simpliciter* may be relative novelty where the comparison class comprises all of history.

Human creativity involves novelty of two kinds: historical and behavioural. *Historical novelty* specifies some (part of) history as the comparison class: some x is historically novel if and only if x is new with regard to the history of some population. The harmonic complexity of Dizzy Gillespie's playing is novel relative to the class of trumpeters before him. A heliocentric cosmology is novel relative to the class of pre-Copernican cosmologies. *Behavioural novelty* relativizes to a particular agent. Some action or thought b is behaviourally novel for some particular agent A if and only if b is new with regard to the previous behaviour of A . My making

a certain dance move is novel relative to my previous behaviour. Your solution to a logic problem may be novel relative to your previous behaviour. These two concepts of novelty—historical and behavioural—are conditions on corresponding concepts of creativity.¹¹

Perhaps creativity consists in just these two properties: resulting non-accidentally from agency and possessing the appropriate relative novelty. If so, then my dance move and your logic proof are creative in only the loosest sense, novel relative only to our respective behavioural histories. They are at most *behaviourally* creative. Gillespie's trumpet playing and Copernicus's cosmological theorizing, however, are creative in a much richer sense; they are *historically* creative. The difference is in the scope of comparison class. Both Gillespie and Copernicus thought and acted in ways novel relative to a broad comparison class. Thus perhaps the richness of creativity corresponds to the breadth of the comparison class relative to which the act or thought is novel.

There are, however, good reasons to think that agency and novelty are insufficient for creativity in any interesting sense. I can construct a simple but unique word search process where I input random words to an internet search engine, then choose words from the hits that correspond to my birth date numbers, and generate a result like the following: 'Sweet faces typically manipulate satellite amenities.' Some of the thoughts and actions that compose this process, and the result itself, are richly historically novel and depend non-accidentally on my agency. However, it is a stretch to attribute creativity in this circumstance, and in spite of the

¹¹ Margaret Boden distinguishes *historical* novelty (and thus creativity) from *psychological* novelty (Boden 2004). *Behavioural* novelty as understood above subsumes psychological novelty, making room for a broader variety of novel phenomena than merely those that are cognitive.

satisfaction of *agency* and *novelty*. So there may be additional conditions on creativity.¹² In any case, the task here is not ultimately one of definition or conceptual analysis. Rather, the conceptual analysis has been adopted as a means for identifying some features of creativity as it is attributed, in order to identify enough constituent properties to coarsely type creativity as a property. This much has been achieved: creativity is type individuated, perhaps only minimally, by the constituent properties ‘being the product of agency’ and ‘being novel.’ *Agency* and *novelty* that is to say, are necessary conditions on an *x* possessing the property creativity.

Creativity as an artefactual, extrinsic, characterizing property

A definition has not been provided. Being the product of agency and being novel relative to some comparison class are not presented as the necessary and sufficient conditions for creativity. They have been presented as necessary, but whether (and what) additional conditions are needed for sufficiency has been left open. This should be no criticism of the present analysis, and for two reasons. First, it isn’t clear that creativity, like so many of our ordinary concepts, will admit of a real definition. We may get close, but there is little reason to think that ‘creativity’ will somehow be more amenable to definition than concepts like ‘art’, ‘equality’, or ‘knowledge’. And we at least have good inductive evidence that these latter concepts do not admit of definition. Second, the questions of primary concern in this paper are ontological ones. An ontology of *F*s is not obligated to providing a definition of ‘*F*’. One might, for example, reasonably give an ontology of artworks without providing a definition of ‘art’—by fitting

¹² There are a number of candidates on offer. For example, some suggest a value condition on creativity (Gaut and Livingston 2003; Gaut 2003; Martindale 1999); some a utility condition (Novitz 1999); and others a modal condition, such that creative thoughts or actions are ones that could not have been tokened before they in fact were (Boden 2004; XXXXb).

artworks into some basic category or categories of existents and providing some principles for individuating artworks one from another. Now of course, having some conceptual clarification of ‘*F*’ enables and indeed may be necessary to answering ontological questions about *F*s. And this is what the two identified conditions accomplish: they help to deliver and support answers to some of the ontological questions on the table.

If creativity is at least partly constituted by the identified properties of agency and novelty then it is not, as was already suggested, a natural kind. This may be an altogether obvious point, since creativity would seem to be an artefactual kind *par excellence*. Natural kinds, if they do exist, are mind-independent, serving in causal explanation and prediction and playing essential roles in laws of nature. On the realist picture at least, they are things that we discover. Creativity, by contrast, depends on human minds and in at least two ways: creativity is the product of intentional agency and depends for its novelty on its relations with classes of human activity, thought, and culture. Creative things are things we make and our concepts and terms for creativity pick out those products of agency. This is not to imply, however, that creativity is an entirely unnatural kind. Some metaphysicians maintain that naturalness of kinds admits of degree, and that a compound property is natural to the degree that its constituents are natural.¹³ Creativity then, may be natural to the degree that its constituent properties are natural. On the face of it, neither of the identified properties are natural—being the product of agency and being novel relative to some class *C* would fail to meet any of the usual criteria for naturalness. But perhaps agency itself, or some other third constituent property (if there is one, or more) is natural. Leave this as an open question.

¹³ See Koslicki 2008: 200-33; Lewis 1986: 60-1.

Assuming, then, that creativity is not a (purely) natural kind, what does this imply for a theory of creativity? Although it may not be a surprising metaphysical insight, it should further relieve some of the worry about the incomplete conceptual analysis offered above. On the one hand, as an artefactual kind, we should have special epistemic access to creativity as a property and should achieve that access through analysis of our practices of attribution. In other words, as the creators of creative things, we are (at least more) immune to error regarding the nature of such things, as contrasted with the possibilities for massive error that we face in attempting to delineate natural kinds.¹⁴ On the other hand, creativity as a non-natural property may just be a somewhat gerrymandered and heterogeneous property. It may be well characterized by certain, more simple properties—the above analysis identifies two—but may ultimately lack an essence. We should then expect that a theory of creativity will only achieve a limited amount of precision.

Some properties are intrinsic, others extrinsic. There is debate about how the intrinsic/extrinsic distinction should be made. Here are two common marks for the distinction; by both marks, creativity comes out on the side of the extrinsic.¹⁵ ‘Intrinsic’ is perhaps most commonly taken to oppose ‘relational.’ A property F of x is extrinsic just in case it is a relational property of x . And relational properties are properties with more than one argument place, such that they cannot be attributed to an object solely by reference to that object (and its internal features or structure). Properties with just one argument place, specifiable just by reference to the object that possesses them, are by this criterion intrinsic. If the intrinsic/extrinsic distinction is made this way, creativity is clearly an extrinsic property. Whatever additional features

¹⁴ See Thomasson 2003, 2007 for a general defense of our privileged epistemic access to artefactual and social kinds as contrasted with natural kinds.

¹⁵ See Humberstone 1996; Weatherson 2006.

contribute to type-individuating creativity, an x is not a product of agency or novel in a non-relational way. One might alternatively distinguish intrinsic/extrinsic in line with a qualitative/non-qualitative distinction (Lewis 1983; Sider 1993). On this conception, qualitative properties are just those properties always shared by perfect duplicates: such that if y is the perfect (actual or possible) duplicate of x , the properties that never differ between x and y are the qualitative properties of x (and y). Intrinsic properties are qualitative properties, and extrinsic properties are non-qualitative—those that can differ between duplicates. By this criterion of intrinsicness, creativity is extrinsic. For example, for any creative x in the actual world W_a (say a painting, a musical performance, a philosophical argument) we can imagine its duplicate y in a possible world W_l where the contextual factors in W_l are relevantly different from those in W_a such that y is not novel relative to W_l (or some comparison class in W_l). Thus x is (minimally) creative while y , the qualitative duplicate of x , is not. By either criterion, then, creativity is an extrinsic property.¹⁶

There is one final feature of creativity as a property worth mentioning, and one that will connect the ‘type of property’ question from this section with the ‘property of’ question from *I.1* above. Creativity is a *characterizing* property: it doesn’t provide counting principles in the way that *particularizing* properties (Armstrong 1978; Swoyer 2008) or *sortal* properties (Strawson

¹⁶ There are many problems with distinguishing intrinsic and extrinsic properties. In addition to those already mentioned above, see Franciscotti 1999 and Yablo 1999, among others. These particular applications are thus only as good as these particular ways of making the distinction. Classifying creativity as extrinsic depends ultimately on whether there is a robust distinction here, and some philosophers worry whether there is one. Here as above, however, metaphysics can be tempered with pragmatism and, as some have suggested, the intrinsic/extrinsic distinction can be made in ways that do important philosophical work (Humberstone 1996; Weatherston 2006). For the purposes at hand, application of the distinction is instructive, making more salient that creativity is a compound property and one *of* a process rather than merely of a single product.

1959) do. Thus in a way that is analogous to quantifying over red things, one can only ask how many creative things are present (at a time and place) by piggybacking on the quantity *of things* (at that time and place); we don't count creativity even if we can count creative things.

This may provide additional motivation for denying the claim that there is a locus for creativity. The product view forced us to interpret attributions of creativity as attributing a property *of* the perceptible product before us. On a pure product view, this is tantamount to saying, of the product and its intrinsic properties only, that '*that* is the creativity'. This result is at odds with our practices. The refined product view treats creativity attributions as saying that '*that* is the creativity' where '*that*' refers to the product understood as possessing a package of relational or extrinsic properties. This, the analysis above suggested, may be ontologically acceptable, but should be avoided provided a more parsimonious alternative. The process thesis provides just such an alternative. Recognition that creativity is a characterizing rather than particularizing property provides additional leverage for the process thesis over the alternative product theses. Both the pure and refined product theses implicitly suppose that creativity is countable, that identifying creativity in any instance is to identify the locus of creativity, in the same way that one appeals to principles of individuation for 'chairs' or 'tables' and then counts the chairs and tables in the room. This is mistaken. Although we can count paintings or theories, which may be creative, we cannot count creativity. Creativity enables characterization of things, but it does not enable partitioning of things into a specific quantity. It could do this only if it were a sortal or particularizing property.

This brings the features of the proposed ontology together. Creativity, the process thesis maintains, is fundamentally a property *of* processes. As such, there is never a singular locus for

creativity. Instead, creativity is an attribute of processes, where processes take place across some segment of time and generally, but not necessarily, proceed towards a terminus. One may appreciate the creative process through its terminating product, its acting agent, or by observing some other parts of the process. Creativity, as a property, is type individuated by more basic properties, two of them identified here: being the product of agency and being novel relative to some comparison class. As compounded out of at least these two properties, creativity is an artefactual, extrinsic, characterizing property. All of this is significant progress towards answering the most basic question: what is creativity? On this ontology, creativity *is* a property of processes, namely those involving, minimally, intentional agency and novelty.

§II. Semantics and epistemology

An ontology of *Fs* will have semantic and epistemological implications. If *Fs* are supposed to be understood as being of this or that category, what does this imply about *F*-talk? And what or how do we come to know or identify *Fs*? A brief answer to each of these questions is offered in turn. Answers to both questions reveal an inherent reflexivity in the overall mode of analysis. Creativity was theorized by consideration of attributive practices, and ontological claims were motivated. Now the questions concern those practices themselves (rather than as a means for analysis), where the answers will be informed by the ontological conclusions derived. For some philosophers, this kind of circularity is never vicious, since all we ever have is our practices and theoretical frameworks. Of those for whom it seems threatening, bear in mind that creativity is not a natural kind (and this seems clear even without the analysis offered in *§I*

above). An analysis of creativity thus seems bound to follow this order: Draw some inferences about the phenomenon by consulting linguistic and conceptual practices regarding that phenomenon. Then, based on the theses inferred, derive semantic theses about our talk of that phenomenon and some epistemic constraints on how we should go about identifying instances of the phenomenon. Ideally, much of our ordinary rational practice is maintained in the output of this analysis, but the semantic and epistemological theses will often be normative in character. And when revision of the existing practice is implied, some kind of error theory is required. This is the general approach taken here, and is analogous to the approach taken by contemporary analyses of, for example, knowledge, art, or value.

Semantics

If creativity is typed partly by the properties ‘being the product of agency’ and ‘being novel’, then a statement about creativity depends for its truth on the presence or instantiation of (at least) these latter two properties. This is perfectly consistent with our ordinary and critical linguistic practice. Part of what we mean when we say ‘creative’ is that the thing spoken of depends in a non-accidental way on an agent and is novel in some important way, and we rationally withhold or challenge an attribution if either of these conditions is not satisfied. This semantic fact is further revealed by our justifications in response to challenges or queries: one will almost certainly mention the agent of an x and the novelty of x if one is explaining the creativity of x .

The process thesis maintains that creativity is fundamentally a property of generative processes. This forces some revision of the semantics of certain creativity statements, namely those that ostensibly involve attribution to a person or product rather than a process. Thus a statement of the form

(4) *o* is creative

where ‘*o*’ is a particular object or event is elliptical for a statement of the form

(5) *o* is the result of some creative process *p*.

And a statement of the form

(6) *A* is creative

where ‘*A*’ is some particular person is elliptical for a statement of the form

(7) *A* is the agent of some creative process *p*.¹⁷

The elliptical treatment of these attribution types accommodates the diversity of our attributive practices, while maintaining the process thesis. Moreover, it seems to do so in a way that requires no major revision: it maintains the surface form of all ordinary statements *and* the ways that we explain the relevant statements, namely, by appeal to a generative process. The proposed ontology thus implies an intuitive even if incomplete semantics for creativity statements. It provides, if one likes, a (partial) sense for ‘creative’: being the product of agency and being novel. And it provides a general category of referent, namely, processes.

¹⁷ Depending upon context, (7) may need to be qualified with a frequency clause (i.e. Pollock acted creatively more often than I do and so this should be revealed by analysis even if it doesn’t show up in the surface language).

Epistemology

If creativity is a property of processes, as the process thesis has it, then identifying creativity is to identify a generative process. Maggie's attribution of creativity was to an object, Pollock's *White Light*. In justifying that attribution, Maggie invoked features of Pollock's creative process: his bodily actions, techniques used, intentions, desires, and so on. She also invoked the novelty of features of this process relative to art history. This example betrays the nature of the knowledge required for identifying creativity as such. Identifying creativity requires background knowledge of context; mere perceptible properties of an object or event are insufficient. And in the case of an inanimate (non-*processing*) object or event *o*, one identifies that *o* resulted from a creative process. This fact is not perceivable in *o* alone: one must have knowledge either of the generation of *O*s—of objects or events of that general category—or of the generation of that particular *o*. Maggie had the former sort of knowledge—of Pollock's action paintings as a category and of the generative process behind them—but one can imagine circumstances where one has knowledge, for example, of the particular history of the work *White Light*.

If creativity is a property of the type suggested above, additional epistemological consequences follow. Consider the two constituent properties of creativity, understood in terms of the conditions *agency* and *novelty*. One might think that artefacts are easily identified as results of agency. However, it is not clear that agency can always or often be identified (wholly) perceptually. Without the right background knowledge, one might take faces in the clouds or rock formations as resultant from agency (which is of course consistent with most of us having

this background knowledge). Running the other direction, one might interpret facsimiles of nature or human-manipulated structures in nature as *not* resultant from agents. In any case, supposing that agency is perceivable in some x alone, this does not suffice for identifying x as the product of agency. The latter requires recognition that x was the non-accidental result of the agent's thoughts and actions. This is not information one can acquire *just* by perceiving x . One cannot identify the aesthetically pleasing paint splattering as the non-accidental result of agency in the absence of some information about the agent of the splattering (e.g. whether the agent was earnestly imitating Pollock or just clumsily moving several cans of paint across her studio). This identification requires background knowledge.

An x is novel only relative to a comparison class C . Thus novelty is not seen, heard, or otherwise perceived in x alone. It is instead identified as a relation between x and C . Identifying novelty thus requires knowledge of C . In the most minimal case, this is a matter of identifying the novelty relative to the behavioural (or perhaps psychological) history of the agent in question. Thus to identify creativity in a child's solution, s , to an algebra problem, one must identify the novelty of s not relative to the history of mathematics or even to the other members of the child's peer group, but rather relative to *that* child's previous behaviour.

Interest in the creativity of Pollock or Bach or Einstein is interest in a broader, historical novelty. To identify the creativity of a Pollock work, one must know, for a start, that Pollock performed actions that were novel relative to the art culture that preceded and surrounded him when he began dripping, throwing, spilling, splashing, and flinging paint on a giant canvass. Knowing facts like these implies knowing some of the context of art, namely, just those events, features, and behaviours relative to which Pollock's techniques and behaviours were novel.

Identifying historical novelty thus depends upon knowledge of (some of) the relevant history, in this case, the history of art.¹⁸

One general epistemological consequence, then, is this: one does not perceive creativity minus the relevant background knowledge. Does this imply that one can never just see or hear creativity? How one answers this question depends upon what one means in asking it, and this depends upon certain theoretical commitments regarding the contents of perceptual experience. On sparse theories of perceptual content, only low-level properties like shape or colour are represented in experience. And so the answer to the question is a trivial ‘no’, since that is the answer given by such theories with regard to the perceptual representation of any high-level property. Kinds, emotions, or semantic properties are not represented in experience.¹⁹ For such theories, the relevant epistemological consequence of the ontology on offer is this: one only judges or forms perceptual beliefs about creativity of an *x* given the appropriate background knowledge about *x*. Alternative theories of perceptual content argue that higher order properties are represented in experience—properties like ‘being a cat’ or ‘being angry’ can be part of the

¹⁸ The epistemic consequences considered are silent on issues of epistemic justification. The claims here are just about identification, where identification is understood veridically. One can thus identify *x* as creative only if one has the relevant background knowledge. (And, as suggested below, one can have an experience of *x* as creative, only given that background knowledge.) It may be appropriate to think about identifications of creativity as being a kind of propositional knowledge. If so, what that knowledge involves depends partly on one’s theory of epistemic justification in a way that seems no different from any other knowledge that *p*. Part of that story will presumably involve mention of the epistemic status of the supporting background knowledge (how it was acquired, what its content is, how much of it one has) and how that knowledge serves as a reason for (or contributes to a reliable means of forming) the belief that ‘*x* is creative’. So while the materials to tell the justificatory story seem readily available to the present theory, no claims are made about the epistemic justification of the identifications or experience.

¹⁹ See Dretske 1995; Tye 1995.

content of experience.²⁰ Perhaps for these theories creativity—as at least partly constituted by the properties ‘being the product of agency’ and ‘being novel’—can be seen or heard. The ontology on offer is consistent with this possibility, requiring simply that to have experiences with these contents, one must have the relevant background knowledge. One does not have an experience that represents x as angry minus a concept ANGER, and one does not have an experience of x as creative minus some background knowledge about x : knowledge about, at least, the agent of which x is a non-accidental product and x ’s relations with some comparison class C .

And so one may or may not just see creativity. It depends upon which theorist of perception you ask. And the answer depends ultimately upon what kinds of things, more fundamentally, we *just see*. Only on some of the relevant views is creativity, as a high-level property, something we may experience. So in this sense one *may* just see creativity. What is clear, no matter one’s commitments regarding the contents of experience, is that one does not see creativity without the relevant knowledge of context. And furthermore, one identifies creativity only with that background knowledge. Identification of creativity then, is much like identifying art among Arthur Danto’s famous indiscernibles. Identifying creativity, like identifying some of modern art, may involve experience. But experience is insufficient, because the perceptible properties of the objects or events are insufficient. One needs knowledge of the relevant context to have the relevant experience or make the right judgement. So, in this sense, one never *just sees* creativity.

²⁰ See Peacocke 1992; Siegel 2006. See Siegel 2005 for general discussion.

Bibliography

- Anscombe, G.E.M. *Intention*, Oxford: Blackwell.
- Aristotle (1984) *The Complete Works of Aristotle*, ed. J. Barnes, Princeton: Princeton University Press.
- Armstrong, D. (1978) *Universals and Scientific Realism, Vol. II. A Theory of Universals*, Cambridge: Cambridge University Press.
- Beardsley, M. (1965) 'On the Creation of Art', *Journal of Aesthetics and Art Criticism* 293: 291-304.
- Boden, M. (2004) *The Creative Mind: Myths and Mechanisms*, 2nd Edn., London: Routledge.
- Currie, G., *An Ontology of Art*, New York, St. Martin's Press, 1989.
- Danto, A. (1964) 'The Artworld,' *Journal of Philosophy* 61: 571-84.
- Davidson, D. (1980a) 'Agency,' in *Essays on Actions and Events*, D. Davidson, Oxford: Oxford University Press.
- _____ (1980b) 'Actions, Reasons, and Causes,' in *Essays on Actions and Events*.
- Davies, D., *Art as Performance*, Oxford: Blackwell, 2004.
- Dretske, F. (1995) *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Francescotti, R. (1999) 'How to Define Intrinsic Properties,' *Nous* 33: 590-609.
- Gaut, B. (2003) 'Creativity and Imagination,' in B. Gaut and P. Livingston (eds) *The Creation of Art*, Cambridge: Cambridge University Press.
- _____ and P. Livingston (2003) 'The Creation of Art: Issues and Perspectives,' in

- B. Gaut and P. Livingston (eds) *The Creation of Art*, Cambridge: Cambridge University Press.
- Ginet, C. (1990) *On Action*, Cambridge: Cambridge University Press.
- Harman, G. (1976) 'Practical Reasoning,' *Review of Metaphysics* 29: 431-63.
- Humberstone, I. L. (1996), 'Intrinsic/Extrinsic,' *Synthese* 108: 205-67.
- Koslicki, K. (2008) *The Structure of Objects*, Oxford: Oxford University Press.
- Lewis, D. (1983) 'Extrinsic properties,' *Philosophical Studies* 44: 197-200.
- _____(1986) *On the Plurality of Worlds*, Oxford: Blackwell.
- Martindale, C. (1999) 'Biological bases of creativity,' in R. Sternberg (ed.) *Handbook of Creativity*, Cambridge: Cambridge University Press.
- Mele, A. (ed.) (1997) *The Philosophy of Action*, Oxford: Oxford University Press.
- Novitz, D. (1999) 'Creativity and Constraint,' *Australasian Journal of Philosophy* 77: 67-82.
- Peacocke, C. (1992) *A Study of Concepts*, Cambridge, MA: MIT Press.
- Ryle, G. (1949) *The Concept of Mind*, London: Hutchinson.
- Searle, J. (1995) *The Construction of Social Reality*, New York: The Free Press.
- Sider, T. (1993) 'Intrinsic Properties,' *Philosophical Studies* 83: 1-27.
- Siegel, S. (2005) 'The Contents of Perception,' *The Stanford Encyclopedia of Philosophy*, (Winter 2008 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/win2008/entries/perception-contents/>
- _____(2006) 'Which Properties are Represented in Perception?' in T. Gendler Szabo and J. Hawthorne (eds) *Perceptual Experience*, Oxford: Oxford University Press.

Strawson, P. F. (1959) *Individuals: An Essay in Descriptive Metaphysics*, Garden City, N.Y.: Doubleday.

Thomasson, A. (2003) 'Realism and Human Kinds,' *Philosophy and Phenomenological Research* 67: 580-609

_____(2007) 'Artifacts and Human Concepts,' in S. Laurence and E. Margolis (eds), *Creations of the Mind: Essays on Artifacts and their Representation*, Oxford: Oxford University Press.

Tye, M., (1995) *Ten Problems of Consciousness*, Cambridge, MA: MIT Press.

Vendler, Z. (1957) 'Verbs and Tenses,' *Philosophical Review* 66: 143-60.

Weatherson, B. (2006) 'Intrinsic vs. Extrinsic Properties,' *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/intrinsic-extrinsic/>

Wilson, G. (1989) *The Intentionality of Human Action*, Rev. ed., Stanford: Stanford University Press.

Yablo, S. (1999) 'Intrinsicness,' *Philosophical Topics* 26: 479-505.

XXXXa [AUTHOR MASKED]

XXXXb [AUTHOR MASKED]

XXXX [AUTHOR MASKED]