

## MINIMALLY CREATIVE THOUGHT

DUSTIN STOKES

---

**Abstract:** Creativity has received, and continues to receive, comparatively little analysis in philosophy and the brain and behavioural sciences. This is in spite of the importance of creative thought and action, and the many and varied resources of theories of mind. Here an alternative approach to analyzing creativity is suggested: start from the bottom up with minimally creative thought. Minimally creative thought depends non-accidentally upon agency, is novel relative to the acting agent, and could not have been tokened before the time it is in fact tokened, relative to the agent in question. Thoughts that meet these three conditions—*agency*, *psychological novelty*, and *modal*—are what may be called *cognitive breakthroughs*. Even if such breakthroughs are not necessary to or definitive of richer creativity, they are indeed central to much of creativity. The minimal analysis provides a more workable explanandum for theories of creativity of varied motivation and method.

Keywords: cognition, creativity, thought.

---

Creativity is a broadly important phenomenon in artistic, scientific, theoretical, and practical contexts. Creative achievements, processes, and persons involve or have creative thoughts. So creative thought is broadly important. In spite of this importance, creativity and creative thought are relatively neglected research topics in philosophy and cognitive science.<sup>1</sup>

The philosophical relevance of creativity is not exclusive to—though certainly is important for—practices of art and art-making. In philosophy of science, an interest in theory construction and change implies a concern with creativity at group and individual levels. In philosophy of language and linguistics, issues concerning language acquisition and generativity may partly concern creative behaviour and thought. In philosophy of

<sup>1</sup> The neglect is most notable in analytic philosophy, though there are exceptions. More work has been done in the cognitive sciences, but little relative to other topics. Recent examples of philosophical and cognitive scientific research are discussed below. See Barsalou and Prinz 1997; Prinz and Barsalou 2002; Boden 1999, 2004; Carruthers 2002, 2007, forthcoming; Finke, Ward, and Smith 1992; Gaut 2003, 2010; Gaut and Livingston 2003; Kronfeldner 2009; Martindale 1999; Novitz 1999, 2003; Simonton 1999; Smith, Ward, and Finke 1995; Sternberg 1999; Stokes 2007, 2008; Stokes and Bird 2008.

mind and cognitive science, an interest in concepts, problem solving, or action, among other issues, will benefit from an analysis of creative thought. Creativity thus should be 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—consciousness, action, perception, to name three—it receives almost no philosophical attention. Creativity does get more attention in cognitive science. But here again, it is neglected relative to other comparably important phenomena—for example, concepts, imagination, and folk mind-reading ability.

This situation might be remedied by a new approach to the phenomenon. Geniuses, masterworks, and theoretical revolutions are exciting and may be an ultimate explanatory goal. However, if one wants to give a philosophical or empirical theory of creativity, geniuses and the like are not the best initial choice of explanandum. Alternatively, start from the bottom up: identify and clarify some fundamental, even if insufficient, features of creativity. The bottom-up approach suggested in this article focuses on *minimally creative thought*. It concludes with the following analysis:

*MC*: Some thought  $x$  is minimally creative if, for some agent  $A$ ,  $x$  is the non-accidental result of agency;  $x$  is psychologically novel; and  $x$  could not have been tokened by  $A$  before the time  $t_i$  when it actually was tokened by  $A$ .

It is convenient to speak of these conditions as individually necessary for minimal creativity. However, the suggestion here is less committal: the first two conditions, agency and novelty, are independently necessary and when conjoined with a third condition, the modal condition, are sufficient for minimal creativity. Thus the modal condition, while *non*-necessary for reasons considered below, is one way to complete an analysis of minimal creativity.

Minimally creative thought, as analyzed, is a more tractable explanandum for both purely philosophical theories and scientifically oriented theories of creativity. The analysis is deliberately sparse, avoiding many of the problems that plague value-rich and domain-specific theories, as well as explanations of high-level creativity. That said, it is just a first conceptual step, and is thus flexible to supplementation as varied theoretical goals and disciplines dictate.

## 1. Creativity and Responsibility

### 1.1. Agency Condition

An attribution of creativity implies an attribution of agency. First, consider the following comparison. We may attribute beauty or other

aesthetic properties, but we do not (properly) attribute creativity to an unusual array of cracks in a rock wall or to the image of a mythical creature in the clouds. If, however, we come upon an abandoned artefact of some sort, say a painting, we might attribute all of the same properties *plus* creativity. Withholding an attribution of creativity in the first case and allowing for it in the second depends upon the same criterion. We see the cracks and clouds as lacking in any marks of agency; while the painting betrays the fact that it depends upon agency. We are willing to call the second, but not the first, creative. Perhaps this is too fast: paintings and other artefacts do not provide fail-proof evidence for agency: one can mistake agents for non-agents and non-agents for agents. This invites a second point. A debate about artificial intelligence and creativity centres, in part, on the question of agency. Whether a computer or its products are creative (*actually* as opposed to just *apparently*) depends upon a more fundamental question, namely, whether the computer is an *agent* with certain cognitive or behavioural capacities. Is the computer autonomous and responsible for its computations and products or is it, as we say, “just running its program”? (Boden 1999, 2004; Cope 1991, 2001; Dartnall 1994; Hofstadter 1994, 2002; Hofstadter and FARG 1995.) Whatever epistemic difficulties there may be, one does not properly attribute creativity to the computer until one properly attributes agency to that computer. Both of these considerations motivate the same point: creativity requires agency.

Agency involves action. Creative things, whatever else is true of them, are things we *do* as agents. Thus:

Agency: an *x* is minimally creative only if *x* is the non-accidental result of agency.

This condition is intended to pick out a (perhaps artificial) class of actions and consequences that non-accidentally result from agents. One way to individuate this class is to consider actions for which one may be praised or blamed. Kicking someone in the shins, stealing your little sister’s lunch money, or cheating on an exam are all blameworthy. Praiseworthy acts and products also depend on agents with intentions. We praise a person for performing well, making well, doing well. Although we may appreciate any benefits or interest we derive from the results, we do not reasonably praise persons who accidentally do something well. We might be thankful or interested or surprised, but we do not praise a person who haphazardly trips the purse snatcher or whose unawareness leaves an aesthetically pleasing trail of mud. The withholding of praise here derives from the lack of non-accidental responsibility on the part of the agents and not from a lack of valuable consequences. Thus the class of actions of interest to a theory of creativity are what we might call *candidates* for praise or blame. The agency condition goes no further than that: it individuates the

non-accidental responsibility component of praise or blameworthiness while remaining neutral on the value component.<sup>2</sup>

Agency, on most accounts, involves 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 an isosceles triangle drawing, must Bob intend to draw an isosceles triangle? The answer to this question is no. As Donald Davidson recognized, attributions of intention are opaque: Hamlet intends to kill the man behind the arras, but he does not intend to kill Polonius, who 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. Intention is only 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). Just as Hamlet is the agent of Polonius’s killing because he intended to kill the man behind the arras, Bob is the agent of the isosceles triangle drawing in so far as he intended to draw, say, a triangle.<sup>3</sup> This addresses one concern: agency may require intention, but not in so strict a way that Hamlet is not the killer of Polonius and Bob not the drawer of the isosceles triangle.

An opposite worry regards the weakness of an agency condition. Given a few common—though not uncontroversial—assumptions, agency comes easy. My friend did not intend to spill her glass of red wine onto my laptop computer, but she did intend to spin in her chair with her arms flailing (where, as a matter of fact, the tipping over of the wine was done by the chair-spinning and arm-flailing). The Davidson/Anscombe thesis of action identification says that “if a person *Fs* by *Ging*, then her act of *Fing* = her act of *Ging*.” Coupling this thesis with Davidson’s semantic criterion for action implies that my friend was the agent of the wine spill, since she intended to spin and flail and spilled wine by spinning and flailing.<sup>4</sup>

<sup>2</sup> Or at least neutral on the *relevant* value component: one might hold a theory of action that makes mere agency or responsibility an evaluative notion. The point here is just that satisfaction of the agency condition is silent with respect to evaluating an action as blameworthy or praiseworthy. See also Ginet 1990, 80.

<sup>3</sup> In alternative terminology, the killing of Polonius is an act of Hamlet’s, and the drawing of the isosceles triangle is an act of Bob’s (Wilson 1989, 89).

<sup>4</sup> The Davidson/Anscombe thesis, as it is called in the philosophy of action literature, can be found in Anscombe 1959 and Davidson 1980a; see Wilson 2009 for discussion. There is much to debate here. *Identificationists* like Davidson identify what some would distinguish as two (or more) actions or an action and an outcome (or outcomes) of an action: Booth’s pulling the trigger just is, for Davidson, Booth’s shooting Lincoln. Davidson is motivated by both the intuition that the “doing by” relation is not a causal one and by a parsimonious ontology of events (Davidson 1980b). *Anti-identificationists* resist this identification, sometimes by appeal to distinct spatiotemporal properties of the identified events. For example, trigger pullings and shootings occupy distinct spatiotemporal regions and are thus, one might think, distinct actions (Thomson 1971). See Pietroski 1998 for discussion and an analysis that attempts to satisfy both identificationist and anti-identificationist intuitions.

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. 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. Here again, in spite of my better intentions, the splattering of paint is my action. By the same criterion as above, I am the agent of the aesthetically interesting splattering (the splattering is an act of mine), just as my friend is the agent of the spilling. But I am responsible for the splattering in a (relevantly) trivial way: the splattering and its array of colour and form were merely an accident resulting from my clumsiness. An observer of the series of events would not reasonably praise me for the aesthetic properties of the splattering, since the event, or at least the relevant results, is only loosely connected with my intentions. So agency, which must be attributed in both cases, is not sufficient for candidacy for praise or blame. An  $x$  may depend upon the agency of some  $A$ , but this alone does not make  $x$  the non-accidental result 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. This is standard fare for theories of action. The action theorist must determine the role of intention in the assignment of agency, and she must provide principles for individuating actions. An interest in creativity is an interest both in creative events—thoughts and action—and in results of those events: say in Coltrane's modal jazz improvisation and Einstein's theory of special relativity.<sup>5</sup> The agency condition aims to pick out just those thoughts, actions, and results that non-accidentally result from agency. The condition 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) is rich enough to warrant praise or blame. No interest in creativity is an interest in accidental action.

"Non-accidental" is thus a placeholder, to be filled by a plausible 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. What is, as some have put it, the "right kind of causation for intentional action" (Wilson 2009)? The challenge might be met in a

<sup>5</sup> And this point, to be clear, is compatible with a variety of commitments, e.g., both a theory that says that products are the locus of creativity and a theory that says processes are the locus. These two theorists, in spite of their differences, will not deny that *both* processes and products are relevant to a theory of creativity.

number of ways. One might make the event or results counterfactually dependent upon an agent and his 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 (see Ginet 1990, 72–89; Harman 1976; Mele 1997; Wilson 1989, 88–167). This challenge needn't be met here. Instead, the two scope worries considered in the last few paragraphs frame a desideratum: an agency condition on creativity must strike a balance between requiring of creativity that an agent intends to get precisely the results that he does get (which is too exclusive) and allowing for accidental performances of creative action (which is too inclusive). Where the line is to be drawn is unclear, but it is safe to proceed on the assumption that there is a distinction here. Creative thought and action are on the side of the distinction where agents are non-accidentally responsible for their actions and results.<sup>6</sup>

### 1.2. *Agency and Inspiration*

Commitment to an agency condition already puts the present analysis at odds with many traditional views. Going back to Plato, creative thought has often been modelled on notions of supernatural inspiration. Plato took poets to be mere media for their muses, conduits for divine inspiration without any real knowledge or understanding of the contents of the lines they compose. Homer knew nothing of war or chariotteering, and so was clearly not responsible for his descriptions thereof (Plato 1997). Schopenhauer placed greater emphasis on madness or irrationality. For him, the genius differs from the insane person only in so far as the former manages to channel her irrationality or worldlessness into the production of art. This accomplishment, however, is out of the control of the genius (Schopenhauer 1958). Call any such view, where there is an absence of a responsible agent, *inspirationalism*.

Inspirationalism is not just one for the ancients or moderns. In a recent book on musical genius, Peter Kivy argues that something like Plato's model is necessary to account for masterworks and masterminds: "Bright ideas are not generated by acts of will through application of some

<sup>6</sup> If one thinks that animal or computer creativity is possible, then an agency condition might be construed more broadly. For example, one might take agency to only require autonomous action that 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 Stokes and Bird 2008). The analysandum for this article is human creativity, which calls for a richer notion of agency. This will be assumed for the remainder of the discussion.

'method.' Bright ideas just 'happen' to people. People who get them are patients, not agents. That was Plato's (or Socrates') discovery. Insight is a kind of 'infectious disease' that one succumbs to. One might well call it the 'passive' notion of genius" (2001, 11). Kivy qualifies inspirationalism in at least two ways. First, the Platonic model is necessary to accommodate many and perhaps all examples of radically creative thought, but it is generally not sufficient: it must be conjoined with a Longinian model that appeals to innate abilities and dispositions. So divine inspiration plus innate creative dispositions explain geniuses like Handel and Beethoven. Second, Kivy recognizes the Platonic model as a kind of myth not to be taken literally. Rather, geniuses require us to treat them "as if" they have been divinely inspired. So the inspirationalist model provides a kind of conceptual marker, tagging phenomena that cannot be fully explained on naturalistic grounds. Call this *as-if inspirationalism*.

Consider the motivation for inspirationalism, of both the literal and the as-if variety. Creative ideas may come to their bearers unbidden like bumps on the head. Speaking of his discovery of the ring structure of the benzene molecule, Friedrich von Kekulé famously reported: "I turned my chair to the fire and dozed. Again the atoms were gambolling before my eyes. This time the smaller groups kept modestly in the background. My mental eye, rendered more acute by repeated visions of this kind, could now distinguish larger structures, of manifold conformation; long rows, sometimes more closely fitted together; all twining and twisting in snakelike motion. But look! What was that? One of the snakes had seized hold of its own tail, and the form whirled mockingly before my eyes. As if by a flash of lightning I awoke" (qtd. Boden 2004, 26). Creative ideas are often described in this way, as ones that "just happen" or "just come to us" unwilling in flashes of insight. This *flash phenomenology* is part of the phenomenology of (some of) creativity. It might motivate the inspirationalist model: one might infer from it that the ideas in question are out of the control of their subjects. The argument would go as follows. Flash phenomenology entails lack of responsibility. Lack of responsibility entails inspirationalism. Grant the second entailment. Does the first entailment hold?

The inspirationalist explanation of flash phenomenology is not the only one available. Consider some common mental act types. One does not deliberately form or change beliefs *just like that*. Pascal knew this, thus his suggested first step to religious belief was to make the wager that God exists given the stakes. But this commitment is not sufficient for the relevant belief. You either have the belief or you don't, and if you want it, you must go through the motions of religion and then, maybe, acquire the belief that accompanies the wager. This is true of beliefs in general: belief formation is not under immediate control. Moreover, beliefs often just come to us. These two features of belief motivate *doxastic involuntarism* (Alston 1989; Bennett 1984, 1990; Williams 1973). The same point can be made for desire. Many desires just come to us: I may suddenly have a

craving for a beer, or some ice cream, or to finish work for the day and go have some fun. Here again, the phenomenology possesses features of abruptness and involuntariness (Millgram 1997). States like beliefs and desires can come in a flash, without one's immediately willing them, much less predicting them.

Does granting flash phenomenology imply inspirationalism for beliefs, desires, and like propositional attitudes? No, it does not. In spite of the fact that they sometimes just come to their bearers, it is reasonable to understand beliefs and desires as states for which an agent is non-accidentally responsible. One can maintain that such cognitive states *feel* this way, and even acknowledge that we lack immediate control over them, without denying that they are states for which we are responsible. With respect to belief, for example, we arguably have indirect voluntary influence: we have control over belief-forming habits and practices that encourage and prevent the formation and maintenance of particular beliefs. Thus praise or blame for a believer may ultimately be traced to his epistemic responsibility (Alston 1989). The same explanatory options from philosophy of mind and psychology are available for creative thought. Creative ideas may (sometimes) *feel* like uncontrolled flashes of insight, but this alone is insufficient to motivate the claim that they *are* (entirely) out of our control. A flash of insight—just like the formation of a belief—depends upon previous, deliberate cognition, on the acquisition of information, application of concepts, imagination and hypothesis generation, the execution of skills, and so on. These considerations block the inference from flash phenomenology to lack of responsibility, and thus block the inference to inspirationalism.<sup>7</sup> Flash phenomenology is not sufficient to motivate inspirationalism. Attributions of creativity imply responsible, non-trivial agency. The agency condition is designed to accommodate this fact.

## 2. Creativity and Novelty

### 2.1. *Relative Novelty*

Creativity implies novelty. Creative thought and action is new or different in some relevant way. One might think that novelty is just newness *simpliciter*. Some  $x$  is creative only if  $x$  has never occurred before. This may characterize instances of radical creativity or genius. But there is little reason to think that it characterizes all instances of creativity. And, what is

<sup>7</sup> The anti-inspirationalist can in fact grant that *some* such states are out of the control of their possessors. However, the (traditional) inspirationalist requires more than this: his position requires that all or at least most creative thoughts are out of the control of their possessors. And this is precisely the inference he cannot have. Even if we grant the (apparently dubious) assumption that all creative thoughts bear flash phenomenology, this does not entail a lack of responsibility for all or most of those thoughts (anymore than it would entail the analogous proposition with respect to beliefs or desires.)



perhaps more important, thinking about novelty in relative terms is more tenable. 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, social context, or an individual agent's behaviour. Novelty *simpliciter* may be relative novelty where the comparison class comprises all of history. Whether a particular Beethoven symphony or Picasso cubist painting meets this criterion is a difficult question, but not one that a general account of creativity must answer.

An interest in human creativity is an interest in 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  $a$  is behaviourally novel (for some particular agent) if and only if  $a$  is new with regard to the previous behaviour of that agent. "Behaviour" here can be understood broadly so as to include thought; bodily action and thought are both things an agent does, and both may be novel relative to the agent that does them. (But note that construing behaviour as including thought is not intended to imply that thought is or reduces to behaviour.) Thus my doing 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 notions of novelty—historical and behavioural—are conditions on corresponding notions of creativity.

Perhaps creativity consists just in the right kind of agency and the appropriate relative novelty. So my dance move and your logic proof are creative in only the loosest sense, novel relative only to our individual behavioural histories. Gillespie's trumpet playing and Copernicus's cosmological theorizing, however, are creative in the richest of senses. 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. An analysis of creativity might stop here, with the following novelty condition (plus agency):

Novelty: an  $x$  is minimally creative only if  $x$  is novel.

Even some mundane actions and thoughts are novel if the comparison class is an individual behavioural history. If the theoretical interest is in something richer, then the comparison class for novelty is broadened. The novelty condition leaves open how the novelty is relativized and perhaps enriched, and thus how the condition is applied in particular cases.

## 2.2. *Psychological Versus Historical Novelty*

An interest in creativity from the perspective of philosophy of mind or cognitive science does better to begin at the lowest level, with something like the behavioural novelty just discussed. Even if the ultimate explanatory goal is broadly novel thought and action, a sensible starting place is the most mundane of creative acts. Moreover, historical comparisons are not the business of cognitive science; thought and behaviour are. So, what must occur for an agent to think or act in ways novel relative to her own behavioural and cognitive history? The remainder of the analysis focuses on this question.

Imagine Carl, a ten-year-old whiz kid who is working solely from a rather antiquated periodic table. Suppose Carl manages to hypothesize all of the chemical elements missing from the table. In spite of the fact that these elements have already been discovered, there is an obvious sense in which Carl's actions are creative, and an account that dictates otherwise is mistaken. Margaret Boden distinguishes *psychological novelty* (p-novelty) from *historical novelty* (h-novelty).<sup>8</sup> Carl's actions are not novel in the latter sense, since they fail to be new with respect to the whole of human history. His actions are, however, p-novel, since they involve or are motivated by thoughts that are novel with respect to *his* mind (Boden 2004, 43). P-novel thoughts may occur multiple times in history; h-novel thoughts occur only once.

Boden characterizes p-novel thoughts as ones that the agent *could not* have had before now. The "could" here is a computational one, relative to generative rules. A thought *could* have been produced before, according to Boden, if that thought is describable/producing by an existing set of generative rules. A thought *could not* have been produced if it is impossible with respect to such rules. This is ambiguous: which generative rules are relevant? Boden sometimes speaks of an objective set of rules. "A merely novel idea is one which can be described and/or produced by the same set of generative rules as are other, familiar, ideas. A radically original, or creative, idea is one which cannot" (2004, 51). "A creative mathematician explores a given generative system, or set of rules, to see what it can and cannot do" (2004, 57). Elsewhere Boden describes the generative rules in more subjective terms, where the relevant

<sup>8</sup> Boden's distinction is in fact between *p-creativity* and *h-creativity*. However, what renders an instance of creativity psychological rather than historical (or vice versa) is the novelty involved. Moreover, while novelty is an essential part of an analysis of creative thought, it is not the entire story: novelty, psychological or historical, is not enough. For example, I can now imagine an orange monkey in a giant cookie jar eating purple crayons. This thought is novel with respect to my mind and (so far as I know) novel with respect to the history of ideas. But is it a thought that we want to call "creative"? Perhaps so, perhaps not. The point is that novelty is a necessary but insufficient condition for creative thought. So we should hesitate to identify creativity with novelty. For these reasons, "p-novelty" and "h-novelty" are the more appropriate terms.

subjects are the receivers of or audience for creative acts. “Fundamentally creative” ideas require, according to Boden, that “our surprise at the creative idea recognizes that the world has turned out differently not just from the way we thought it would, but even from the way we thought it *could*” (2004, 41–42). And again: “[T]he surprise that we feel on encountering a creative idea often springs not merely from an unfamiliar combination, but from our recognition that the novel idea simply *could not* have arisen from the generative rules (implicit or explicit) which we have in mind. With respect to the usual mental processing in the relevant domain (chemistry, poetry, music . . . ), it is not just improbable, but *impossible*” (2004, 52). But the second proposition does not follow from the first. The impossibility of a thought does not depend upon the generative rules that are *believed*, by some particular person, to constrain that domain; such a broad metaphysical conclusion cannot be inferred from such narrow epistemic circumstances.<sup>9</sup>

So Boden is faced with the following dilemma: either the generative rules relative to which p-novel thoughts are impossible are subjective—just the ones that an audience or group of people have in mind—or they are objective. Opting for the first horn spells obvious trouble: an account of creative thought will have little explanatory purchase if the relevant properties are novel only relative to some believer or other. Moreover, there will likely be inconsistencies between such perspectives: what you and I recognize as the relevant generative rules may well be different, and so whether or not something *is* creative will depend upon whom you ask. This is a relativism to avoid. Perhaps the set of generative rules is objective. Opting for this horn also spells trouble. The thrust of distinguishing p-novelty from h-novelty is to weaken the conceptual requirements so that a thought may qualify as novel even if it has been tokened (by someone else) before. If, however, p-novel thoughts could not have been tokened before relative to some set of generative rules, then the comparison class is not in fact the *individual mind* but an abstract class of computational structures. The only thoughts that will qualify will be ones that involve breaking (or perhaps bending) such rules. This dissolves the relevant distinction between p-novelty and h-novelty: a thought that is indescribable in terms of agent-independent generative rules would seem to better capture what we want to call historical novelty.

The fix is to adopt the notion of p-novelty while rejecting the relativization to generative rules, objective or subjective. An individual-relative novelty condition is:

<sup>9</sup> Unless, of course, the relevant modality is *epistemic*, such that the event in question is/was impossible *for all we know*. However, as discussed below, this epistemic reading would result in an undesirable relativism. And, given most of Boden’s discussion, it seems she wants something stronger—some non-epistemic modality—for her claims about impossibility.

P-novelty:  $x$  is minimally creative only if  $x$  is psychologically novel.

The novelty here is relative to the psychological agent in question: a thought is p-novel for some agent just in case the agent has never tokened the thought before. (Note that this condition is silent with respect to the modal status of creative thoughts.) This kind of novelty is an important explanandum for any explanation of creative thought, in particular, for any psychological, naturalistic, or cognitive scientific theory of creativity.

### 3. Creativity and Possibility

#### 3.1. Modality and Cognitive Change

As suggested above, given the right kind of novelty (that is, novelty relative to a sufficiently broad comparison class), agency and novelty may suffice for rich creativity. If, however, one is interested in creative thought, both mundane and radical, this sparse analysis is unsatisfactory. First, if the interest is not just in radical creativity but also in more mundane instances of creativity, then agency and novelty are not enough. One can wilfully token novel thoughts, one after the other. Only some of these thoughts are creative in any interesting sense. And second, even if the ultimate explanatory goal is radical creativity, a good start is to understand minimal instances of creative thought and what features distinguish it from *merely* novel thought. What, in addition to agency and p-novelty, suffices for minimally creative thought?

Boden is right: creative ideas (involving psychological novelty) are ones that *could not* have been had before. Her mistake is to relativize the modality to either an agent-independent set of generative rules or to the rules that an individual or a group of people believe to constrain the relevant domain. Instead, the spirit of the psychological/historical distinction as it is initially presented should be maintained. “The psychological sense concerns ideas (whether in science, needlework, music, painting, literature . . .) that are surprising, or perhaps even fundamentally novel, with respect to *the individual mind* which had the idea” (Boden 2004, 43). Creative ideas are often ones that *could not* have been tokened before by the mind in question, ones that were impossible relative to *that* agent’s cognitive position. Thus:

Modal:  $x$  is minimally creative only if  $x$  could not have been tokened by  $A$  before  $t_i$  when it actually was tokened.

The relevant modality is nomological or, more specifically, psychological. (The modal condition is here construed as a necessary condition. However, as discussed below, there are difficult questions about the necessity of this condition, ones the answering of which seems to outstrip

conceptual intuitions. The condition will thus finally be suggested as non-necessary but jointly sufficient for minimal creativity.)

Modal is motivated by an acknowledgement that creative thoughts are novel and novelty involves change of some kind. What distinguishes some merely p-novel thoughts from minimally creative thoughts is that the latter are thoughts that the agent in question could not have had before. If a thought  $c$  is psychologically novel for some agent  $A$ , then  $A$  made some cognitive change in order to enable that novelty. Before this change,  $A$  *could not* have tokened the novel thought. Cognitive changes can be made in many ways,  $A$  may acquire new skills, information, beliefs, concepts, perform some kind of mental operation, use imagination or hypothetical reasoning, or some combination thereof. Before the time of tokening  $c$ , the suggestion is,  $A$  lacked some of these cognitive elements and activities, without which  $c$  was not psychologically possible.

There is an important worry here. It is not clear that novelty, even of the mere psychological kind, requires significant cognitive change. Put in the terms of modal, it is not clear that all creative thoughts, even minimally creative thoughts, are ones that the agent *could not* have had before the time that he in fact did. Reflection upon some masterminds drives this worry: Picasso produced so many great works, it may be implausible to think that each such production was not possible until the time it was actually produced. This worry is considered in section 3.2 below.

The p-novelty and modal conditions are importantly related but are not equivalent. The two conditions come apart with respect to actual thoughts versus possible thoughts. Some ideas meet p-novelty while failing to meet modal: I may form a novel thought, but given my cognitive position I could have formed it before. Not all p-novel thoughts are ones that *could not* have been had before by the agent thinking them. The entailment does not run the other way either: satisfaction of modal does not imply satisfaction of p-novelty. So, the p-novelty condition (and indeed the novelty condition) is silent with respect to *modal cognitive profiles*, requiring for its satisfaction that an agent, as a matter of fact, tokens some thought that she has not before tokened. A satisfaction sentence for this condition takes this form:  $(\exists x)(\exists y)(\exists t_i)(\forall t)[(t < t_i \supset \sim Hxyt) \ \& \ Hxyt]$ , where “ $x$ ” is a cognitive agent, “ $y$ ” an idea, “ $t_i$ ” some particular time, and “ $H$ ” the relation “to token.” The modal condition is silent with respect to *actual cognitive profiles* and concerns, instead, modal cognitive profiles. A satisfaction sentence for this condition takes this form:  $(\exists x)(\exists y)(\exists t_i)(\forall t)[t < t_i \supset \sim \Diamond Hxyt]$ . Truly closing either of these sentences does not entail closing the other truly.

Nonetheless, the conceptual interaction between novelty and modal is important. The two conditions are motivated by the observation that thought is generally systematic. What thoughts one has tokened depend on one’s broader cognitive profile. What one *can* think, how one can think about it, and how one is able to think and reason, generally depend upon

this profile and the skills that one possesses (plus other situational and environmental circumstances). A modal cognitive profile—the cognitive changes that one can make—thus depends importantly upon an actual cognitive profile. The next question concerns just how tight this link is.

### 3.2. *The Non-necessity of the Modal Condition*

One might worry about modal as follows. “That’s silly: surely Beethoven *could have* composed his Eighth Symphony an hour or two before he did, or even a day or week before he did. He just didn’t. Perhaps he was too tired or preoccupied. Perhaps he was shopping or travelling or chasing women. Whatever the case, he could have had the relevant thoughts before he in fact did and thus could have composed the piece before he in fact did.” A first response to this worry is to recall that the modal condition concerns a nomological modality, targeting possibilities relative to an actual cognitive perspective. So while it may be true that there are close possible worlds where Beethoven composed his Eighth a few hours or few days earlier, in the actual world Beethoven composed his Eighth just when he did. Given its novelty, the thoughts involved in this composition required cognitive change, and thus before those changes the relevant thoughts were (nomologically) impossible. The changes may have been slight, involving mental imagery or some conceptual tweaking rather than acquisition of new skills or knowledge.

However, this response may miss the deeper worry about any modal condition on having creative thoughts. Working artists, for example, continue to use the same skills and knowledge to create many artworks, even if each work is novel. Likewise, a physicist uses the same theoretical knowledge to form new hypotheses, an engineer the same mathematical understanding to solve a structural challenge. Thus, although the artist or scientist (or whoever) may think in new ways, why should such thinking be impossible before it actually occurred?

This general worry is a complex and important one. It is underwritten by at least two concerns. First, we have limited epistemic access to the mind of anyone other than ourselves, and are thus in a less than ideal position to identify the necessary cognitive steps to Beethoven’s creation, or to anyone’s creative thought for that matter. What’s more, our introspective abilities are fallible in this regard. This makes answering the above questions about modal, for any one particular case, practically impossible. Second, the analysis to this point has proceeded as if creativity is a natural or psychological kind, but it may not be. It may depend upon or reduce to such kinds. For example, perhaps agency is a psychological kind, and creativity necessarily involves agency. But there may be no one natural, psychological kind that answers to “creativity.” It may instead be a social or artefactual kind (see Searle 1995; Thomasson 2003, 2007); or it may be a variegated or gerrymandered kind. In any case, conceptual

intuitions about creativity vary. Some will find modal intuitively implausible. Others will find it plausible. And others will be indifferent. It is not clear what should adjudicate such intuitions if indeed there is no well-individuated kind to discover or if our attributive practices for creativity just diverge.

This forces a dilemma. On the one hand, conceptual intuitions and resources underdetermine how an analysis of creativity should be completed. But a conceptual analysis that consists of only two necessary conditions is disappointing. So, on the other hand, one may insist on a complete analysis of creativity. The trouble here is that any completed analysis of creativity will to some degree be stipulated if intuitions are in fact as varied as they seem. The middle ground is to accept the ostensible fact that creativity is not a well-individuated kind and that the concept may be indeterminate or ambiguous. The role of conceptual analysis then becomes—beyond accommodating any uncontroversial conceptual conditions—one of enriching that analysis in ways that are informed but oriented toward some theoretical goals or context. The question, after agency and novelty are in hand, becomes: What is the theoretical context of interest?

The present theoretical interest is in creative thought, and thus in laying a conceptual groundwork for a cognitive architecture of creative thought. The analysis should aim to reveal or make theoretically tractable the cognitive events, processes, and changes involved in creativity. In this light, modal remains an attractive condition. The necessity of the condition is controversial. But the analysis can remain neutral on the necessity of modal while maintaining that when conjoined with agency and *p*-novelty, it is sufficient for minimal creativity. A completed analysis of minimally creative thought is this:

*MC*: Some thought *x* is minimally creative if, for some agent *A*, *x* is the non-accidental result of agency; *x* is psychologically novel; and *x* could not have been tokened by *A* before the time *t<sub>i</sub>* when it actually was tokened by *A*.

Reasons were given above for doubting the necessity of modal. Beethoven's Eighth may be creative in virtue of non-accidental dependence upon Beethoven's agency and the novelty of the thoughts and actions involved in its composition—novel relative to a broad and important comparison class, say, pre-nineteenth-century Western music. So maybe agency and novelty are sufficient. It is hard to say just what cognitive changes were necessary for this composition, such that Beethoven could not have composed it before he in fact did. And so it is hard to motivate modal as any kind of condition on this creative accomplishment. But leave Beethoven for the moment and consider more mundane creativity. As suggested above, modal, although distinct, is importantly connected to

p-novelty. It provides a distinguishing mark for merely p-novel thoughts versus minimally creative thoughts. Some thoughts are novel relative to an agent, while others are novel and could not have been tokened (by that agent) before they were. Thoughts in the second class meet both p-novelty and modal. This reveals an important question for any theory of creative thought: What kinds of cognitive efforts and changes enable thoughts that meet both p-novelty and modal?

Modal cognitive profiles depend upon actual cognitive profiles. For some agent *A*, individuating *A*'s actual cognitive profile informs us of the cognitive possibilities for *A*. A cognitive profile may be individuated with greater or lesser fineness of grain. Most finely, a cognitive profile might be individuated at the level of mental-state tokens. However, this would be too fine-grained: some mental tokens are relevant to the possibility of *A*'s tokening some thought *c*, while others are clearly not. Thinking about environmental selection pressures may be relevant to the possibility of a breakthrough thesis in evolutionary theory, but desires about the hockey game surely are not. Thoughts about the hockey game might, as a matter of fact, trigger a string of thoughts that result in the breakthrough, but the point is that other thoughts, perhaps about boxing or chocolate or whatever, *could have* initiated the same string. So although *some* tokens (partly) determine the modal status of tokening some thought for an agent, equally important are the relations between certain thoughts and the general abilities and skills of the agent.

The notion of a heuristic path is useful, since it provides a coarser method of individuating cognitive profiles vis-à-vis their modal properties.<sup>10</sup> A creative thought *c* is causally contingent upon a cognitive process. Some of the states that compose that process are necessary for the tokening of *c*, others could be omitted or replaced with no causal upshot for *c*. As a first sketch at least, the former set of (necessary) thoughts and their relations constitute the heuristic path to *c*. It is this set of thoughts, understood at the level of organization rather than mere individual tokens, that makes *c* a possible thought *for that agent*. A heuristic path has its own enabling conditions—for example, the heuristic path to the breakthrough in evolutionary theory depends upon, at least, certain empirical and theoretical knowledge, skills of theoretical and statistical interpretation, certain motivations, and novel conceptual combinations—but it will also be causally open to a variety of initiating thoughts, say, thoughts about hockey or chocolate or television game shows. Roughly, then, an actual cognitive profile can be individuated at a coarser level of description, one that describes a certain organization of states and capacities and how they relate to environmental circumstances. This organizational structure is the heuristic path to a creative thought.

<sup>10</sup> See Currie (1989, 46–84) for an application of this conceptual machinery to the ontology of artworks.



The suggestion, finally, is that modal, when appended to agency and p-novelty, is sufficient for minimally creative thought. One way to understand a minimally creative thought is as follows. An agent *A* has a cognitive profile up to  $t_i$  such that prior to  $t_i$ , *A* could not have some thought (or set of thoughts) *c*. Leading up to  $t_i$ , *A* makes cognitive changes that enable the heuristic path to *c*.<sup>11</sup> Prior to this time, *A* lacked elements of that heuristic path; perhaps he lacked certain skills, had to connect or apply some concepts in new ways, perform some hypothetical reasoning, or imagine things to be some way other than they are.<sup>12</sup> Once *A* makes the relevant changes to enable the heuristic path to *c*, *A* tokens *c*; *c* is the (non-accidental) product of *A*'s agency, is p-novel (for *A*), and could not have been tokened (by *A*) prior to  $t_i$ . According to *MC*, *c* is minimally creative; *c* could be a thought with just about any content: a new way of slicing vegetables, an insightful but common description of a philosophical thesis, a clever shortcut for one's walk home; it all depends upon *A*'s cognitive profile. What makes the satisfaction of these three conditions sufficient for minimally creative thought is that they are only satisfied by cognitive breakthroughs. This kind of breakthrough, which the relation between p-novelty and modal makes salient, is at the heart of an interest in creative thought. Breakthroughs are breakthroughs precisely because some change had to occur to make them *possible*. A new thought that was possible before, requiring no significant cognitive change, is not a breakthrough, not creative, but merely novel.

#### 4. Value and Creativity

*MC* might not fully capture creativity like Beethoven's. Perhaps Beethoven could have composed his Eighth before the time that he in fact did. And even if he couldn't have, one might worry that satisfaction of *MC* does not suffice for Beethoven's creativity. Instead, the richness of Beethoven's creativity, and other geniuses like him, might be better captured by agency and novelty relative to a specific comparison class, thus involving broad historical novelty rather than mere psychological novelty. Perhaps all of this is so. No matter. As stated at the outset, this analysis does not aim at explaining radical creativity or genius but instead aims at some of the cognitive fundamentals of more mundane creativity. In this respect, *MC* is a successful analysis, since it reveals what is valuable about

<sup>11</sup> Nothing in the phenomenon or the analysis requires a precise account of timing here. The changes simply had to occur before the novel thought(s), and most likely involve a process across time, progressing towards the culminating thought(s).

<sup>12</sup> Here the strictly philosophical analysis can be supplemented with some of the work in cognitive science, for example, on conceptual combination and exploration (Finke, Ward, and Smith 1992; Finke 1995; Ward, Smith, and Finke 1999), assembly and activation of action schemata (Carruthers 2007), and unconscious, *incubated* cognition and neural plasticity (Smith and Blankenship 1989, 1991; Stokes 2007).

more mundane creativity. And this is progress that may ultimately contribute to an analysis of genius.

*MC* individuates a class of thought and action that one might call, as intimated above, *breakthroughs*. A thought that meets the three conditions of *MC* is, for the agent in question, a cognitive breakthrough. Any such thought depends non-accidentally upon the intentional action of the agent, is novel relative to that agent, and requires some significant change before which it was impossible for that agent. Likewise for a bodily action that meets these conditions. It requires little argument to show that these breakthroughs, even if common in everyday human life, are valuable. A breakthrough might be needed to comprehend and use new concepts, to apply new skills, solve simple problems in one's environment, perform multiple tasks more efficiently, and so on. On a group level, such breakthroughs are important for adaptive changes in the face of challenges in the environment. Cognitive changes are simply valuable because they enable novel opportunity for thought and action.

Minimal creativity is also a valuable concept, since it describes a more tractable explanandum for a theory of creative thought. Even granting that satisfaction of modal, and thus *MC*, is unnecessary for richly creative behaviour like that of Bach or Beethoven, it remains plausible that much of richly creative behaviour does involve the kind of cognitive breakthrough individuated by the minimal analysis. If this is the case, then studying the cognitive architecture of minimally creative thought is a promising angle on the richer phenomena. One might, for example, attempt to give an account of the role of imagination in enabling previously impossible thought (see Gaut 2003). Presumably an explanation of imagination in minimally creative thought would provide some insight on the analogous role for imagination in the richer creativity of masterminds. And one could tell similar stories about other cognitive capacities and operations and their role in minimally creative thought, with an ultimate explanatory goal of genius or rich creativity.

Rather than baldly asserting that creativity is valuable, the minimal analysis thus provides *reasons* for thinking that creativity, even when minimal, is valuable. This is an advance beyond recent creativity literature in philosophy and cognitive science. Many theorists assume that creativity is a value-laden concept and make value a condition for creativity. "There is a broad consensus that creative products and acts must exhibit originality and *be valuable*" (Gaut and Livingston 2003, 10; emphasis added). An *x* is creative only if *x* is valued or to be valued. This is intuitive but uninformative. Endorsing a value criterion does nothing to enhance the explanatory power of a theory of creativity that, say, makes agency and novelty the criteria for creativity. A bald value criterion is silent on the important value question, namely, *why* creative thought and behaviour is valued. Consider an analogy. If one asks about the nature of, say, a carburettor and is told, "A carburettor is a very useful part of an internal

combustion engine,” one comes away with no insight on what a carburettor does or how it does it. At best, one has a reason via testimony for thinking that a carburettor is valuable, but no idea *why* it might be valuable. Moreover, an *x* may be valuable in virtue of some property *F* of *x*, but this does not imply that *value* is a property of *x*. A carburettor may be valued *in virtue of* its individuating characteristics or functional properties, but this does not imply that value is one of those properties. The same is plausibly true of creativity. And so a better strategy for analyzing creativity is to grant that creative things are valuable, and then attempt to identify reasons for thinking they are valuable—conditions on creative thought and behaviour. *MC* does just this, even if the explanation is not yet complete.

### 5. Minimal Creativity: Virtues and Concerns

Many mundane thoughts are, according to *MC*, minimally creative. Comprehending complex concepts of quantum mechanics or successfully working through formal logic proofs might, for some agent, satisfy the three conjointly sufficient conditions. Such thoughts, for some persons, are minimally creative. This qualification makes no commitment to the historical importance, radical novelty, or other rich evaluations of the relevant cognitive events. It simply recognizes that these events possess (some) features fundamental to creative thought. They are cognitive breakthroughs for the relevant agents. They are distinct from merely novel thoughts, being both non-accidentally dependent upon agency and previously impossible for that agent.

Relative to all of the other theories discussed, the minimal analysis takes a new approach to the phenomenon. Although the eventual explanatory goal is the same—creativity—*MC* identifies an important, even if minimal, explanandum.<sup>13</sup> A minimal analysis of creativity has a number of theoretical benefits. Here are three, each briefly clarified in the paragraphs that follow. First, inverting the common approach and working instead from the bottom up, the analysis provides conceptual clarification for any theory of creativity. Second, the analysis provides some answer to the value question about creativity, and makes a start for theories that take

<sup>13</sup> Similarly, Lawrence Barsalou and Jesse Prinz (1997, 2002) distinguish *mundane creativity* from *exceptional creativity*. Only a relative few enjoy exceptional creativity while, they argue, all humans enjoy mundane creativity. Barsalou and Prinz locate mundane creativity in concept acquisition. They claim that concept acquisition requires that an agent abstract from new perceptual experiences (often representing unfamiliar stimuli), memory, and existing concepts to form concepts that are novel with respect to the mind of that agent. It is an open question whether we form concepts in this way or in non-perceptual ways, and whether such concept formation is something *we do*. In any case, the Barsalou/Prinz view is instructive since it offers another compelling example of *how* mundane cognition might reasonably be called creative.

creativity to be value-laden without committing to those theories. Finally, the analysis makes some advance on certain purportedly spooky features of creativity, or at least suggests how they should not trouble a philosophical or scientific theory of creativity.

Although some of the discussed theories of creativity offer definitions of the concept “creativity,” few if any offer an extended analysis of the conditions they argue for or, in some cases, just assume. This is one culprit for the relative neglect of the topic. Conceptual clarification should precede any theory of  $x$ , whether the theory be strictly empirical or richly philosophical. Instead, much of the research on creativity to date has chosen domain-specificity of creativity over general conceptual analysis of creativity. This choice is motivated by the fact that creativity does not occur in a vacuum; it occurs in specific domains or contexts. Many theorists isolate their concerns to, say, creativity in art or creativity in science, or, more narrowly, creativity in plastic arts or creativity in physics. Theorists of creativity in a domain  $D$  often get stuck in more traditional problems of  $D$  and theories of  $D$ , or end up biasing their framing of and approach to the target problems (of creativity) in favour of the domain-specific problems, or both. The result—an analysis of  $x$  in  $D$  without any clarification of  $x$  *simpliciter*—is not inviting to the newcomer. *MC* is an improvement in this regard. Worries about modal notwithstanding, *MC* provides motivation for and clarification of both agency and novelty. Any theory of creativity should benefit from this clarification. And with the addition of modal, the minimal analysis targets cognitive breakthroughs, which plausibly figure into richer instances of creativity. This conceptual clarification is offered in a domain-neutral way.

The minimal analysis offers some reason for thinking that creative thought and action, even when minimal, are valuable. And it does so without making any explicit commitment to a value condition for creativity. Many theorists, however, are committed to a value condition. *MC* is consistent with this commitment. In fact, the minimal analysis provides a good first step even for a theory that insists that creativity is an essentially value-laden concept. The first step in an account of value is to identify what one can at a descriptive level. Consider some analogies. If we want to ask whether someone’s actions are praiseworthy or blameworthy, we first get straight on the descriptive facts of her action. If a court is deciding whether a defendant is guilty of some crime, it first debates and identifies the physical evidence. If you and I are arguing over the aesthetic value of some artwork, we had better be sure that we agree what the artwork is and what its physical properties are. Individuating these descriptive facts precedes consideration of the evaluative ones in the order of analysis. These cases are disanalogous from our target, since they concern individual instances of evaluative judgment, not general accounts of value or inherently evaluative phenomena. Nevertheless, the progression from descriptive to evaluative issues is common to both, and this is

instructive for theories of creativity. *MC* provides the first step, identifying some general descriptive features of creativity. Once these conditions are in hand, one can build in evaluative features befitting one's theoretical commitments and goals.

Perhaps the most obvious culprits for the philosophical and cognitive scientific neglect of creativity are the many spooks and mysteries that accompany attributions of creativity. Two of these spooks—novelty and flash phenomenology—have been addressed above. The novelty of creative ideas encourages some to infer that creative ideas emerge *ex nihilo*. And the flash phenomenology of creative ideas implies *inspirationalism*. If these inferences are correct, then creativity is not an attractive topic for such fields of research as analytic and naturalistic philosophy and cognitive science. The minimal analysis shows why both inferences are mistaken. Novelty is a relational property, and even in cases of richly historical novelty, the thoughts or actions in question are novel relative to some definite comparison class: some social or historical category of behaviour and or artefacts. And the combination of the p-novelty and modal conditions shows how at least some cognitive breakthroughs emerge, if emergence is the appropriate concept, from a particular cognitive profile. No creation *ex nihilo* here. The discussion of agency shows that although one can admit that many creative thoughts possess flash phenomenology, this alone does not imply *inspirationalism*. A theorist of creative thought has available to him the same resources as the philosopher of mind who wants to explain ordinary mental states in a way that is not agency-independent. At the very least, *inspirationalism* is no more a forced consequence for creative thought than it is for belief or desire.

All of this should be encouraging for a naturalistic or analytic philosopher, a psychologist, or a cognitive scientist, among others. Creativity is a widely important phenomenon, and theories of mind have a wealth of varied resources. The present analysis shows that, when approached minimally, creativity is a theoretically tractable phenomenon.

*Department of Philosophy*  
*University of Toronto*  
*170 St. George Street*  
*Toronto, ON M5R 2M8*  
*Canada*  
*d.stokes@utoronto.ca*

### **Acknowledgments**

I have benefitted from conversations on this topic with many people over the past several years. Thanks to Lucy Allais, Jon Beatty, Vince Bergeron, Maggie Boden, Jon Bird, Peter Carruthers, Ron Chrisley, Stephen Davies, Dom Lopes, Matthew Kieran, Amanda Klein, Aaron Meskin, Andy

McGonigal, Michael Morris, Murali Ramachandran, Patrick Rysiew, Elliot Paul, A. David Smith, Kathleen Stock, Chris Thornton, Anami Vakil, Jonathan Weinberg, and any others whom I have unfortunately omitted. This article also benefitted greatly from a work-in-progress workshop in Toronto; special thanks to Stephen Biggs, Imogen Dickie, Benj Hellie, Mohan Matthen, and Jessica Wilson.

## References

- Alston, William. 1989. "The Deontological Conception of Epistemic Justification." In *Epistemic Justification*, 115–52. Ithaca: Cornell University Press.
- Anscombe, G. E. M. 1959. *Intention*. Oxford: Blackwell.
- Barsalou, Lawrence, and Jesse Prinz. 1997. "Mundane Creativity in Perceptual Symbol Systems." In *Creative Thought: An Investigation of Conceptual Structures and Processes*, edited by Thomas Ward, Steven Smith, and Jyotsna Vaid, 267–307. Washington, D.C.: American Psychological Association.
- Bennett, Jonathan. 1984. "Deciding What to Believe." In *A Study of Spinoza's Ethics*, 159–62. Indianapolis: Hackett.
- . 1990. "Why Is Belief Involuntary?" *Analysis* 50: 87–107.
- Boden, Margaret. 1999. "Computer Models of Creativity." In *Handbook of Creativity*, edited by Robert Sternberg, 351–72. Cambridge: Cambridge University Press.
- . 2004. *The Creative Mind*. Second edition. London: Routledge.
- Carruthers, Peter. 2002. "Human Creativity: Its Cognitive Basis, Its Evolution, and Its Connections with Childhood Pretence." *British Journal for the Philosophy of Science* 53: 225–49.
- . 2007. "The Creative Action Theory of Creativity." In *The Innate Mind*, vol. 3: *Foundations and the Future*, edited by Peter Carruthers, Stephen Laurence, and Stephen Stich, 154–71. Oxford: Oxford University Press.
- . Forthcoming. "Creative Action in Mind." *Philosophical Psychology*.
- Cope, David. 1991. *Computers and Musical Style*. Oxford: Oxford University Press.
- . 2001. *Virtual Music: Computer Synthesis of Musical Style*. Cambridge, Mass.: MIT Press.
- Currie, Gregory. 1989. *An Ontology of Art*. London: Macmillan.
- Dartnall, Terry (ed.). 1994. *Artificial Intelligence and Creativity: An Interdisciplinary Approach*. Dordrecht: Kluwer.
- Davidson, Donald. 1980a. "Agency." In *Essays on Actions and Events*, 43–61. Oxford: Oxford University Press.
- . 1980b. "Actions, Reasons, and Causes." In *Essays on Actions and Events*, 3–19. Oxford: Oxford University Press.

- Finke, Ronald. 1995. "Creative Realism." In *The Creative Cognition Approach*, edited by Steven Smith, Thomas Ward, and Ronald Finke, 303–26. Cambridge: Cambridge University Press.
- Finke, Ronald, Thomas Ward, Steven Smith. 1992. *Creative Cognition*. Cambridge, Mass.: MIT Press.
- Gaut, Berys. 2003. "Creativity and Imagination." In *The Creation of Art*, edited by Berys Gaut and Paisley Livingston, 148–73. Cambridge: Cambridge University Press.
- . 2010. "The Philosophy of Creativity." *Philosophy Compass* 5, no. 12: 1034–46.
- Gaut, Berys, and Paisley Livingston. 2003. "The Creation of Art: Issues and Perspectives." In *The Creation of Art*, edited by Berys Gaut and Paisley Livingston, 1–32. Cambridge: Cambridge University Press.
- Ginet, Carl. 1990. *On Action*. Cambridge: Cambridge University Press.
- Harman, Gilbert. 1976. "Practical Reasoning." *Review of Metaphysics* 29: 431–63.
- Hofstadter, Daniel. 1994. "How Could a COPYCAT Ever Be Creative?" In *Artificial Intelligence and Creativity: An Interdisciplinary Approach*, edited by Terry Dartnall, 405–24. Dordrecht: Kluwer.
- . 2002. "Looking Emmy Straight in the Eye—and Doing My Best Not to Flinch." In *Creativity, Cognition, and Knowledge*, edited by Terry Dartnall, 67–100. Westport, Conn.: Praeger.
- Hofstadter, Daniel, and FARG (Fluid Analogies Research Group). 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books.
- Kivy, Peter. 2001. *The Possessor and the Possessed*. New Haven: Yale University Press.
- Kronfeldner, Maria. 2009. "Creativity Naturalized." *Philosophical Quarterly* 59, no. 237: 577–92.
- Martindale, Colin. 1999. "Biological Bases of Creativity." In *Handbook of Creativity*, edited by Robert Sternberg, 137–52. Cambridge: Cambridge University Press.
- Mele, Alfred (ed.). 1997. *The Philosophy of Action*. Oxford: Oxford University Press.
- Millgram, Elijah. 1997. *Practical Induction*. Cambridge, Mass.: Harvard University Press.
- Novitz, David. 1999. "Creativity and Constraint." *Australasian Journal of Philosophy* 77: 67–82.
- . 2003. "Explanations of Creativity." In *The Creation of Art*, edited by Berys Gaut and Paisley Livingston, 174–91. Cambridge: Cambridge University Press.
- Pietroski, Paul. 1998. "Actions, Adjuncts, and Agency." *Mind* 102: 73–112.
- Plato. 1997. *Ion and Republic*. In *Complete Works*, edited by J. M. Cooper. Indianapolis: Hackett.

- Prinz, Jesse, and Lawrence Barsalou. 2002. "Acquisition and Productivity in Perceptual Symbol Systems: An Account of Mundane Creativity." In *Creativity, Cognition, and Knowledge*, edited by Terry Dartnall, 105–27. Westport, Conn.: Praeger.
- Schopenhauer, Arthur. 1958. *The World as Will and Representation*. Translated by E. F. J. Payne. Indian Hills, Colo.: Falcon Wing's Press.
- Searle, John. 1995. *The Construction of Social Reality*. New York: Free Press.
- Simonton, Dean. 1999. *Origins of Genius*, Oxford: Oxford University Press.
- Smith, Steven, Thomas Ward, and Ronald Finke (eds.). 1995. *The Creative Cognition Approach*. Cambridge, Mass.: MIT Press.
- Smith, Steven, and Steven Blankenship. 1989. "Incubation Effects." *Bulletin of the Psychonomic Society* 27: 311–14.
- . 1991. "Incubation and the Persistence of Fixation in Problem Solving." *American Journal of Psychology* 104: 61–87.
- Sternberg, Robert (ed.). 1999. *Handbook of Creativity*. Cambridge: Cambridge University Press.
- Stokes, Dustin. 2007. "Incubated Cognition and Creativity." *Journal of Consciousness Studies* 14: 83–100.
- . 2008. "A Metaphysics of Creativity." In *New Waves in Aesthetics*, edited by Kathleen Stock and Katherine Thomson Jones, 105–24. London: Palgrave Macmillan.
- Stokes, Dustin, and Jon Bird. 2008. "Evolutionary Robotics and Creative Constraints." In *Beyond the Brain: Embodied, Situated, and Distributed Cognition*, edited by Benoit Hardy-Vallée and Nicolas Payette. Newcastle, U.K.: Cambridge Scholars Publishing.
- Thomasson, Amie. 2003. "Realism and Human Kinds." *Philosophy and Phenomenological Research* 67: 580–609.
- . 2007. "Artifacts and Human Concepts." In *Creations of the Mind: Essays on Artifacts and Their Representation*, edited by Stephen Laurence and Eric Margolis, 52–73. Oxford: Oxford University Press.
- Thomson, Judith. 1971. "The Time of a Killing." *Journal of Philosophy* 68: 115–32.
- Ward, Thomas, Steven Smith, and Ronald Finke. 1999. "Creative Cognition." In *Handbook of Creativity*, edited by Robert Sternberg, 189–212. Cambridge: Cambridge University Press.
- Williams, Bernard. 1973. "Deciding to Believe." In *The Problems of the Self*, 136–51. Cambridge: Cambridge University Press.
- Wilson, George. 1989. *The Intentionality of Human Action*. Revised edition. Stanford: Stanford University Press.
- . 2009. "Action." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Available at <http://plato.stanford.edu/archives/fall2009/entries/action/>.