MINIMAL CREATIVITY: A COGNITIVE MODEL

by

DUSTIN R. STOKES

B.A., Missouri State University, 1998

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

(Philosophy)

THE UNIVERSITY OF BRITISH COLUMBIA

July 2006

Abstract

Creativity is remarkable in at least two ways. It is remarkably important to all manner of human practice, culture, and theory. It is, however, remarkably neglected by philosophers and psychologists. This neglect is symptomatic of the complexity of the phenomenon, of a number of conceptual confusions, and of a tradition of unconstrained theoretical analyses. But it can be avoided: an informative model of creativity that accommodates its importance and complexity is possible.

After diagnosing the confusion and misanalysis, we endorse an alternative explanatory strategy: begin at the fundamental level of creativity and model creativity as a psychological phenomenon. This encourages a descriptive analysis: all evaluative concerns are to be initially set aside. Moreover, concerns specific to this domain or that—say to physics or painting—are to be set aside. The explanandum is thus narrowed to the psychology of *minimal creativity*. The explanans is cognitivist and naturalistic: all features of the target phenomenon are explained in terms of cognitive processes and from within the bounds of natural science.

Chapter 1 offers a conceptual analysis of minimal creativity: the minimum conditions independently necessary and conjointly sufficient for attributions of 'creativity.' The definition that results is adopted for the remainder of the thesis. Chapter 2 identifies a number of extreme cognitivist treatments of creativity. There is much to be learned from these views, but a balance between them is both preferable and possible. Chapter 3 provides a functional analysis of problem solving and then identifies the functional similarities between problem solving and creative thought: both comprise context, problem, solution, and heuristic roles. Chapter 4 is an analysis of the different

modes of imagination; it concludes with an inference that imagination uniquely enjoys cognitive freedom. Chapter 5 identifies the role of imagination in creative cognition. Imagination, as a cognitive state that is both workful and playful, is ideally suited to service creative thought. Chapter 6 provides a naturalistic model of incubated cognition and its role in creative cognition, focusing on neural plasticity and the effects of attention on the functional networking of the brain.

Table of Contents

Preface

Creativity is important for obvious reasons. It is, or so tradition and intuition would have it at least, essential to artistic creation and scientific theorizing. Perhaps less obviously, it is needed for problem solving in a vast array of circumstances: social, academic, professional, political, sporting, gaming, among others. It is thus, for these reasons and others, a remarkably complex and intriguing phenomenon. It should invite and sustain the interest of theorists of all stripes.

All of this said, creativity is a phenomenon largely neglected by philosophers and psychologists. Philosophical analyses have been offered here and there, and psychological studies done from time to time, but there has been no sustained discourse about the topic. Other topics in the arts, say the importance of artistic intention or definitions of art, and topics in mind and psychology, say folk psychological theorizing or the status of artificial intelligence, have been anything but neglected. These topics, although no doubt important, lack the *scope* of importance to human life possessed by creativity. Yet it is the former and not the latter that have enjoyed volumes of research. Some diagnosis of these asymmetric symptoms, for this neglect, is called for.

Diagnosis

Spooks, *confusion*, and *mud*. These are at least three of the causes for this neglect, all of them similar and causally connected in ways that prevent any clean or determinate etiological story. But such a story is not needed. Instead, we merely need clarify the three identified causes and then turn to what remedies are available for the neglected phenomenon.

Spooks: Creativity has, since the days of ancient, been tagged as spooky and mysterious in a number of ways. Here are some of the more common spooks. *Novelty*: creative thoughts or acts are very special kinds of things. They have not happened before and do not obviously connect with what happened before. This, one might think, is in fact a point of definition: novelty is partly or wholly what makes something suitable for attributions of 'creative.' But if a novel event or thought is not and does not depend on previous events, then it emerges *ex nihilo*. Things that come from nowhere are mysterious if anything is. *Flash phenomenology*: creative ideas famously come to us in flashes of insight. We learn this through reading the introspective reports of geniuses. And we see this firsthand by introspecting our own creative insights, assuming we've had one or two. Attention may be right there, elsewhere, or nowhere, and a burst of brilliance comes to us as if inspired by a supernatural God or muse. Given both the novelty of the idea and the lack of immediate bidding that seemed to beget it, this phenomenological feature is certainly one that needs explaining. *Incubation*: not only do ideas come to us in a flash, but they sometimes come to us when our minds are elsewhere altogether. The reports of geniuses brings this point home all the more: shouts of *Eureka!*, images of self-enveloping snakes, fire-side reveries, and opium-filled dreams are but a few of the romanticized examples that populate biographies of renowned artists and scientists. The concept of incubation is introduced to accommodate this feature of creativity. Creative thought requires a period of incubation, where the unconscious self or unconscious processing does the work that the conscious mind failed to do. When conscious attention is returned, mysteriously, one has the answer that was not gotten, and perhaps could not be gotten, before.

Confusion: Creativity, by the looks of it, is not a well-behaved concept. Or at least this much is true: it is traditionally treated in so many different ways that it isn't clear that the treatments are all treatments of the same concept. Much of the biographical literature is in case-study format: identify a genius and study all features of his or her character, circumstances, and life. Literature in the arts has focused on products—artifacts like novels, paintings, and symphonies—and, sometimes, on geniuses as well. Psychological studies, what few there have been, primarily target creative processes or thought. Philosophers have focused on an amalgam of these, sometimes distinguishing them, sometimes not. Authors and researchers taking these various approaches occasionally clarify how 'creativity' is to be understood. All too often, they do nothing of the sort, proceeding instead as if it were a well understood concept and now the job is just to say how this person, product, or process has it.

Mud: Creativity is present in perhaps all domains of human life. This fact has made for some very muddy waters: addressing too many competing issues at once makes for a messy explanandum. Here are just two sets of issues that likely overwhelm the theorist of creativity. *Creativity-in-D*: creative acts and thoughts do not occur in a vacuum; they occur in certain contexts or domains. This motivates many a theorist to isolate her concerns to, say, creativity in art or creativity in science or, more narrowly, creativity in plastic arts or creativity in physics. There may be good reasons to do this, but it is problematic as a start. Theorists of creativity in art, often get stuck in more traditional problems of art and aesthetic theory or wind up biasing their framing of and approach to the target problems in favor of the other problems, or both. Theorists of creativity in science can be described analogously, as can any theorist who starts with

creativity-in-D. *Evaluative concerns*: this point is partly related to the last. Creativity, some argue, is an ineliminably evaluative concept. Ascribing creativity to some object, act, or person involves evaluation or judgment. Or stronger, objects or events in the world are not determinately or objectively creative, but rather depend for such status on the evaluations of some individual or group of individuals. If these points are right, then creativity is a value concept, and its analysis invites many of the problems and issues that attach to value theory—relativity/objectivity of value, instrumentality/intrinsicality of value, kinds of value, to name just three. The choice for an evaluative over a descriptive treatment exacerbates an already challenging task.

Spooks, confusion, and mud are wonderful things for funhouses, ambushes, and psychedelic festivals, but they do a real number on theoretical analysis. In addition to causing their own distinct problems, they also overlap and combine, resulting in two extreme views. Take enough mystery, conceptual confusion and complexity, and you often end up with an unrigorous or problematically unconstrained analysis or explanation. This is certainly true of creativity. Theories of creativity often embrace the spookiness of the features discussed, offering supernatural and mystical explanations thereof. They shrug off any obligation to clarify the property (or properties) that so-called geniuses share with one another, forcing their readers to guess what a definition might be. They fail to separate and set aside some of the complex and intertwined issues that surround creativity, embracing an overly complex explanandum that results in an overly complex, and often unintelligible, explanans.

This view has a rippling effect, further encouraging an opposite view, itself already encouraged by the spooks, confusion, and mud. Naturalists observe these

features of creativity and the theories that so often embrace and exploit them, and turn tail. Phenomena with mysterious features like those, encouraged by so much conceptual unclarity and complexity, are not deemed scientifically tractable by the naturalistic philosopher or psychologist. And so they have been, and remain, largely neglected by naturalism.

Prognosis

Assuming that the naturalist remains compelled by the theoretical importance of creativity, how might she respond to this diagnosis? How is she to treat the spooks, confusion, mud, and the theories they engender?

Spooks: The naturalist can treat the spooks the way naturalism should. Stick to one's naturalism. Deny their spookiness. If they are observable phenomena—and since they have been consistently observed we should grant that they are—then they can be explained by our best natural theories. Any spooky features are reduced to merely phenomenal features, and that makes them no different from northern lights, after-images, or static electricity.

Confusion: The naturalist avoids confusion by starting with the basics: offer a careful and thorough conceptual clarification, a definition if possible, of 'creativity.' Then narrow the scope of one's explanation by focusing on just one creative kind—one kind of attributee of the concept. It is a trivial point to say that there are similarities between creative persons, products, and processes. But it is not trivial to say that there are differences as well. Isolating just one creative kind as the explanandum avoids confusion.

Mud: The naturalist avoids getting stuck in the mud in much the same way she avoids confusion: identify and narrow the explanandum. Avoid, if possible, the evaluative issues. Start with description and see how far that gets the explanation. This of course limits the explanatory scope of the approach. A descriptive approach may not run very far, but it is nevertheless the first step and should always precede the part(s) of the explanation that is value-laden. This can be done in a way that remains consistent with—and in fact provides a solid theoretical foundation for—a variety of value-rich theories of creativity. At the same time, generalize. Do not focus on creativity in this or that domain. Rather, attempt to give a model of creativity as it may be found in any domain. This tack becomes more feasible once the evaluative issues are dropped, since many of them derive from some particular domain or other.

Application

This thesis offers a naturalistic model of creativity, following the prognosis on offer. It acknowledges the spooky features traditionally identified as part of the explanandum while refusing to exit the realm of naturalism to explain them. The reasoning here is straightforward. The common mention of features like flash phenomenology and incubation in both folk reports and theoretical analyses betrays their conceptual significance to creativity. But this significance and the ostensible mystery of these features does not require or justify the unconstrained speculation they have received. Rather, we might instead think that any phenomenon sufficient to generate *so much* speculation, is ripe for naturalistic analysis. The history of science teaches this lesson time and time again: compare an Aristotelian model of vision with one today, an

early modern explanation of phantom limb sensation versus one from contemporary neuroscience, the shift from religious to clinical treatment of mental disorders like schizophrenia.

The thesis begins with a conceptual analysis of creativity, isolating the minimal psychological requirements for creative thought. The target is thus narrowed to creative cognition. This analysis, and the model that emerges from it, are purely descriptive. The suggestion to set aside evaluative concerns is taken seriously. Such concerns are important to understanding creativity but have, as argued above, barred much theory from getting off the ground. The model is, moreover, intended to be general. One will find differences between creativity in art versus creativity in science, or between creativity in mathematical problem solving versus some other domain. It is not obvious, however, that these are differences *in the creativity* from the two compared domains; rather, they may well be differences in the domains themselves. We find out by proceeding with the general, descriptive model, and seeing how far it takes us. Moreover, creative instances in each domain share something, or they would not warrant creativity ascriptions. The most likely candidate for that point of commonality is the cognitive process involved. Despite their differences, creative accomplishment in bronze sculpture, downtempo electronica, organic chemistry, plate tectonic theory, classical linguistics, and planar geometry all involve similar processes of thought. To identify those similarities is to identify creative cognitive processing. And doing that puts us one step closer to offering a comprehensive model of creativity.

Chapter 1
Minimal Creativity: A Conceptual Analysis

What do we mean when we use the term 'creativity'? We ascribe creativity in multifarious ways: to the radical artist, the brilliant scientist, and the ingenious problem solver; to a brush stroke method or prose style, to the procedures of a chemist or biologist, to a mode of problem solving; to a work of art, to a scientific discovery, to a solution to a tricky puzzle. What conditions underpin our application of the concept of creativity to creative persons, their processes, and the products that result from them? There is confusion in folk ascriptions and philosophical theorizing of 'creativity.' As we will see, much of this confusion stems from the evaluative issues that attach to many uses of the term. In what follows, a minimal analysis of creativity is offered in the form of three descriptive conceptual conditions. The analysis takes the following form:

> *MC:* Some x is minimally creative if and only if, for some agent A, x is the product of the agency of A; x is psychologically novel; and x could not have been tokened by agent A before the time t_i when it actually was tokened by A.

It is convenient to talk of these conditions as individually necessary and jointly sufficient for attributions of *minimal creativity*. At the very least these conditions are importantly symptomatic of—and likely constitutive of—our concept of creativity, and construing them as criteria for creativity is progress. Moreover, this analysis enables a clearer start to a study of creativity, since it resists commitment to (but remains consistent with much of) the complex and often confusing evaluative treatments in the literature.

1.1 Creativity and responsibility

1.1.1 Agency Condition

Creative acts and products are things we *do* and *make*, things for which we are commended and admired. Thus some entity x is minimally creative only if x is the product of agency. Products of agency require responsible agents—they are actions we perform and consequences that non-trivially result from such performances.[1] One way to individuate these kinds of products is to consider actions for which we may be praised or blamed. Kicking someone in the shins, stealing your little sister's lunch money, or cheating on an exam are all products of agency. Note that these actions are worthy of blame not praise. We praise a person for performing well, making well, doing well. Praiseworthy acts and products, like blameworthy ones, are artifacts of intentional action and praiseworthy persons are persons who perform such acts, all of which presupposes agency. 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 who with eyes closed chooses the correct answer out of 100 choices, or whose clumsiness causes an aesthetically pleasing puddle of paint. The withholding of praise here derives from the lack of responsibility on the part of the agents and not from a lack of valuable consequences. We thus might think of products of agency as *candidates* for praise or blame. However, the agency condition takes us no further than that: it individuates the responsibility component of praise/blameworthiness while remaining neutral on the value component.

---

[1] I take for these purposes 'actions' to include mental or cognitive actions, that is, thinking.

The concept of creativity—and thus any ascription of creativity—is bound up with agency. No matter if we attribute creativity to a person, an artifact, or some process, our attribution implies a responsible agent, since agency is built into the property attributed. Here are two reasons for thinking this is so. First, consider the following comparison. We may attribute beauty or other aesthetic properties, but we do not attribute creativity to an unusual array of cracks in a rock wall or to the image of our favorite mythical creature in the clouds. If, however, we come upon an abandoned artifact 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 of any marks of agency; the painting betrays the fact that it depends upon agency. We are thus willing to call the second, but not the first, creative. Perhaps this is too fast: paintings and other artifacts do not provide fail proof evidence of agency. This brings us to the second point. The debates about artificial intelligence and creativity center, in part, on the question of agency. Whether we want to call a computer or its products creative (*actually* as opposed to just *apparently*) depends upon a more fundamental question, namely, whether or not the computer possesses the kind of cognitive capacity requisite for agency. We must ask: is the computer responsible for the product 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). Both of these considerations motivate the same point: creativity requires agency.

The agency condition does not require, however, that a person always intends to get precisely the results that she does in fact get. It merely requires that she is attempting

to complete some task where the resulting thought or artifact counterfactually depends upon this attempt. Had *A* not attempted to complete some task *t*, the resulting thought or act *c* would not have been had or performed by *A*. This implies that at least some of the features of *c* counterfactually depend upon *A* and her attempts but does not require that *A* intended to have or perform *c* in particular. *A* is thus non-trivially responsible for *c*.

Problem: this exposition of non-trivial responsibility is trivially satisfied. Had my friend not forced herself out of bed this morning, she would not have dominoed a book into a full glass of red wine on to my computer by spinning in her chair. Had I not attempted to cross the street to buy a slice of pizza by jaywalking between 1st and 2nd Avenues, I would not have forced a tiny Smart Car to wrap itself around a telephone pole in a surprisingly interesting arrangement of metal, plastic, and paint. Had Bob not gone for that sixth donut, he would not have arrived to class with a maple-leaf shaped stain of jelly on his white dress shirt. These counterfactuals all satisfy the above schema of non-trivial responsibility and thus the agency condition. However, none of them involve the kind of intentions on the parts of their agents requisite for creative action or thought. So as it stands, we have not given a workable schema for *non-trivial* responsibility. Here is the task: we need to strike a balance between requiring of creativity that an agent intends to get precisely the results that she does get (which is too exclusive) and allowing for accidental performances of creative action (which is too inclusive). This is no doubt a tricky task. But let us assume that it can be done: where the line is to be drawn is unclear but it is safe to proceed on the assumption that there is a distinction there. Our focus then is on the side of the distinction where agents are, in some sense, non-trivially responsible

for their cognitive processes and the results they afford. Only when such a condition is met do we (justifiably) call the person or act creative.

1.1.2 Agency and inspiration

The agency condition already puts the present analysis at odds with many traditional views. Going back to Plato, creativity has often been modeled on notions of supernatural inspiration. Call the motley crew of such views *inspirationalism*. According to Plato, poets are 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 charioteering, and so clearly was not responsible for his descriptions thereof (Plato 1989). Schopenhauer places greater emphasis on madness or irrationality. For him, the genius differs from the insane only insofar as the former manages to channel his irrationality or worldlessness into the production of art. This accomplishment, however, is out of the control of the genius (Schopenhauer 1958). 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 "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 (Kivy 2001: 11).

Kivy qualifies inspirationalism in at least two ways. First, the Platonic model is necessary to accommodate many and perhaps all examples of radical creativity, but it is generally not sufficient: it must be conjoined with a Longinian model which 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.*

Commitment to the agency condition is inconsistent with as-if inspirationalism. Here are a few additional reasons against as-if inspirationalism. Kivy's first qualification does little to assuage worries regarding appeals to inspirationalism. The Platonic model, even when complemented by the Longinian model, does not sufficiently *explain* geniuses in particular or creativity in general: it merely puts a finger on the sometimes unprecedented nature of radical creativity. The as-if inspirationalist may respond by saying that this is just the point: we cannot explain genius but only conceptually mark it. So this disagreement—regarding whether explanation is possible, and what suffices for explanation—may be intractable. Sufficiency aside, it hardly seems plausible that inspirationalism is *necessary* to accommodate all instances of genius. More to our purposes, but perhaps not Kivy's, inspirationalism of whatever stripe does not seem necessary to explain general (less radical) creativity. (More on non-radical creativity in the next section.)

Perhaps this is too fast. We should consider the motivation for inspirationalism, of both the literal and as-if varieties. There is an obvious sense in which creative ideas 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 reports:

> 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 (quoted in Boden 2004: 26).

Creative ideas are often described in such ways, as ones that "just happen" or "just come to us" unwilled in flashes or bursts of insight. This is part of the phenomenology of creativity; call it the *flash phenomenology*. This phenomenological feature might motivate appeals to the inspirationalist model: one might infer from this feature that the ideas in question are out of the control of their subjects. Flash phenomenology entails lack of responsibility. And lack of responsibility entails inspirationalism. Grant the second entailment. Our question then regards the first entailment: does it hold?

The following fact implies a negative answer: the inspirationalist explanation of flash phenomenology is not the only one available. Consider some more mundane mental acts. Beliefs are not the kinds of cognitive states that we deliberately form or change *just like that*. Pascal knew this, thus his suggested first step was to make the wager that God exists given the stakes, but this commitment is not sufficient to get you 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, you might acquire the belief that accompanies the wager. This is true of beliefs in general, we lack immediate control over their formation. 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 of our 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 play. Here again, the phenomenology of the state's formation possesses features of abruptness and involuntariness (Millgram 1997). Finally, consider memories evoked by some sensation or other: I may recall my mother's kitchen when walking by a cozy bakery or my little league days when smelling fresh cut grass. Such memories come to me in a flash, without my intending them much less predicting them. Thus many thoughts just come to us: they involve something like *flash phenomenology*. Granting this feature, are we forced into positing inspirationalism for beliefs, desires, and memories? No, we are not. Most contemporary philosophy and psychology does not endorse such a position and lacks no explanatory power as a result. In spite of the fact that they sometimes just come to their bearers, it is reasonable to model beliefs, desires, and memories as states for which the agent is non-trivially responsible. We can maintain that such cognitive states *feel* this way, and even acknowledge that we lack immediate control over them, without exiting the realm of naturalistic explanation.[2] With respect to belief, for example, William Alston argues that we at least have *indirect voluntary influence*: we have direct control over belief forming habits and practices which encourage and prevent the formation of particular beliefs (Alston 1989). The same explanatory options are available for creativity. Creative ideas may *feel* like uncontrolled flashes of insight, but surely we should hesitate to infer from that that they *are* out of our control. So we block the

[2] Michael Polanyi discusses a similar phenomenon with respect to bodily movement (Polanyi 1981). The muscle contractions requisite for movement often *feel* out of our control. I recognize that I have the thought to clench my fist and intend to clench my fist, but the clenching *feels as if* it just happens to me once I've formed (often sub-consciously) the relevant intention. This phenomenology does not force inspirationalism. There are standard neurophysiological explanations for bodily movement, and none of them appeal to supernatural movers.

inference from flash phenomenology to lack of responsibility, thus blocking the inference to inspirationalism. Flash phenomenology is not sufficient to motivate inspirationalism.

The anti-inspirationalist can in fact grant that *some* such states are out of the control of their possessors. However, inspirationalism requires more than this: it 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 (ostensibly 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, desires, or memories.)

Contrary to the Platonic, Schopenhauerian, and Longinian models, then, the present analysis is anti-inspirationalist, anti-irrationalist, and anti-nativist, respectively. We should and do require more responsibility of agents and acts that we call creative. The agency condition is intended to accommodate this fact.

1.2 Creativity and novelty

1.2.1 Psychological versus historical novelty

To count as creative, a thought or act must be novel. But novel for whom? It is a mistake to think that our concept of creativity requires that a thought be novel *simpliciter*, or even relative to all cognizers or agents. Imagine Carl, a 10 year old whiz kid who is working solely from a rather antiquated periodic table. If Carl manages to hypothesize all of the elements missing from the table, we would want to call his actions creative. 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 our account is mistaken if it dictates

otherwise.[3] Margaret Boden distinguishes *psychological creativity* (P-creativity) from *historical creativity* (H-creativity). Carl's actions are not creative in the latter sense, since they fail to be novel with respect to the whole of human history. His actions are, however, P-creative, since they involve or are motivated by ideas that are fundamentally novel with respect to *his* mind (Boden 2004: 43). P-creative ideas and acts may occur multiple times in history; H-creative ideas occur only once.[4]

1.2.2 Novelty and minds: Abandoning generative rules

Boden characterizes P-creative ideas as ones which the agent *could not* have had before now. The 'could' here is a computational one, relative to generative rules. An idea *could* have been produced before, according to Boden, if that idea is describable/producible by an existing set of generative rules. An idea *could not* have been produced if it is impossible with respect to such rules. There is some ambiguity about just 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" (Boden 2004: 51). "A creative mathematician explores a given generative system, or set of rules, to see what it can and cannot do" (Boden 2004: 57).

---

[3] This discussion and much of what follows begs the question against those who drive a wedge between artistic creation and scientific discovery, presupposing instead that both achievements are *creative* ones. No argument is offered for this presupposition, but the central claims made do not depend upon its concession: if one prefers to dichotomize scientific discovery and artistic creation, then the present discussion speaks to the latter and not to the former.

[4] Lawrence Barsalou and Jesse Prinz make a similar distinction. They distinguish *mundane creativity* from *exceptional creativity*, emphasizing the importance of the former to general intelligence. They claim that concept acquisition involves mundane creative activity. Only a relative few enjoy exceptional creativity, while all humans enjoy mundane creativity. It is an open question whether such activity is creative, since the symbols that structure concepts, and thus the process of acquiring concepts, is inaccessible to consciousness and so is non-deliberate (Barsalou and Prinz 1997, 2002). See 3.3.3 for more on Prinz and Barsalou, and 1.3.1 and 5.4 for more on mundane creativity.

Elsewhere Boden describes the generative rules in more subjective terms, where the relevant subjects are the receivers of or audience for creative acts and ideas. "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"* (Boden 2004: 41-2). 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* (Boden 2004: 52).

But the second proposition does not follow from the first. The impossibility of an act or 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 epistemological circumstances.[5]

So Boden seems faced with the following dilemma: either the generative rules relative to which P-creative ideas are impossible are subjective—just the ones that an audience or group of people have in mind—or they are objective. As just indicated, opting for the first horn spells trouble: a descriptive story of creativity will have little explanatory purchase if relativized to some believer or other. Moreover, there will likely be inconsistencies between such perspectives: what you and I consider the relevant generative rules to be may well be different, and so whether or not something *is* creative will depend upon who you ask. This implies a relativism we should want to avoid. So

---

[5] Unless, of course, the relevant modality is *epistemic*, such that the event in question is/was impossible *for all we know*. Given most of Boden's discussion however, it seems she wants something stronger—some non-epistemic modality—for her impossibility claims.

perhaps the set of generative rules is objective. Opting for this horn also spells trouble. The thrust of distinguishing P-creativity from H-creativity was to weaken the conceptual requirements so that a thought or act may qualify as creative even if it had been tokened (by someone else) before. If, however, P-creative ideas could not have been tokened before relative to some set of generative rules then the reference point is not in fact the *individual mind* but an abstract class of computational devices. The only ideas that will qualify will be ones that involve breaking (or perhaps bending) such rules. This comes very close to dissolving the distinction between P-creativity and H-creativity: an idea or act which is indescribable in terms of agent-independent generative rules would seem to better capture what we want to call historical creativity. So either way, relativizing creativity to generative rules results in an overly narrow definition of creativity.

The fix is to adopt the notion of p-novelty while rejecting the relativization to generative rules, objective or subjective. The novelty condition thus states that 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.

One final qualification. Boden chooses to use the term 'creativity' in making the psychological/historical distinction. But this distinction is, at bottom, one with respect to novelty. And while novelty seems to be an essential part of an analysis of creativity, it is not the entire story: novelty, psychological or historical, is not enough. For example, I can now imagine an orange dwarf in a giant cookie jar eating purple crayons. The content of 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 the kind of thought that we want to call 'creative'? Perhaps so, perhaps not. The point is that novelty is, at most, a necessary condition for application of 'creative', not a sufficient one. So we should hesitate to identify creativity with novelty.[6] For this reason, 'P-novelty' and 'H-novelty' are the more appropriate terms.

1.3 Creativity and possibility

1.3.1 Modal condition

Boden is right: creative ideas (in the psychological sense) are ones which *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 a group of people believe to constrain the relevant domain. Instead, we should stick to the spirit of the psychological/historical distinction as it is initially presented: "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 ones that *could not* have been tokened before by the mind in question, ones that were impossible relative to *that* agent's cognitive position. The modal condition thus says: x is minimally creative only if x could not have been tokened by A before t_i. The relevant modality is nomological or, more specifically, psychological.

Consider an agent A who at time t_i tokens a creative idea c. Before t_i, A lacked some knowledge, beliefs, concepts, imaginative states, perhaps even a grasp of some set

[6] The insufficiency of novelty is what, at least in part, motivates Boden to appeal to impossibility relative to generative rules. As discussed, this move suffers various problems, both in terms of internal coherence and theoretical consequences.

of generative rules, which are requisite for the tokening of *c*. So before t_i, *A* lacked some of the prerequisite cognitive goods, *G*. Lacking *G*, *A* thus could not have tokened *c* prior to t_i. Note two things: this speaks neither to the possibility allowed by an objective set of generative rules nor to what is impossible for some *other* mind or mind(s) given the set of rules they grasp or believe.

There will be more and less interesting satisfactions of this condition (and thus, when the other two conditions are satisfied, more and less interesting instances of minimal creativity). Sometimes an agent could not have tokened a certain thought for the simple reason that she was not thinking about the right domain or set of contents. For example, when learning French, I will have a number of thoughts regarding French verb conjugation that I could not have had before learning about particular features of the language (e.g. a significant quantity of verb vocabulary, plus the various pronouns, tenses, etc.). More interestingly, Kekulé could not have tokened the thought that the benzene molecule is ring-like in structure prior to the time he did, since he did not allow himself the relevant conceptual machinery—organic chemistry at the time deemed such molecular structures impossible, and Kekulé presumably followed that tradition until his creative insight. He imagined something unorthodox, and this imagining was requisite for the tokening of his creative thought(s).[7]

Many satisfiers of MC's modal condition, plus its agency and novelty conditions, will be rather mundane instances of cognition: the thoughts about French verb conjugation mentioned above provide an example. *MC* thus bears the consequence that thoughts we might not ordinarily call creative are at least minimally creative. Is this

[7] More on conceptual innovations like Kekulé's in 3.2.2, 3.3.1, and 3.3.2. More on the role of imagination in chapter 5.

consequence—call it the *mundane creativity consequence*—a problem? If we keep a finger on Boden's distinction between psychological and historical creativity, with an eye towards describing minimally creative cognitive processes without invoking richer evaluative senses of 'creative', then the answer is no. In calling some instances of concept, belief, and knowledge acquisition, among other kinds of cognitive states, 'minimally creative', we are not committing to the historical importance, radical novelty, or other rich evaluations of such states. We are only recognizing that such states bear marks that are fundamental to creative thought: they bear a certain kind of psychological novelty, agency, and modal status. At bottom, this is just to say that thinking often requires a bit of creativity on the part of their agents. Learning French or quantum mechanics are tasks that require *of me* novel cognitive processing, much of which was not previously available *to me*. Mundane creativity is thus not a problem, but a welcomed consequence.[8]

1.3.2 The condition profligacy charge

Another worry goes as follows. If an idea *c could not* have been had by *A* prior to t_i then, quite trivially, *c would not* have been had by *A* prior to t_i. So an idea or act that meets the modal condition will meet the novelty condition. Thus satisfaction of the

[8]As mentioned above (see footnote 5) Barsalou and Prinz argue that concept acquisition is mundanely creative (1997; 2002). 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*. No matter: the Barsalou/Prinz view is instructive, since it compells us to acknowledge *how* mundane cognition might be called creative. Granting that we form concepts (either in the way the perceptual symbol theory proposes, or in the way a more traditional amodal concept theory proposes), we are obligated to say that doing so is a psychologically novel act. This point generalizes to a variety of cognitive states. Belief formation, skill acquisition, and a variety of other kinds of learning and thought require the same novelty of their agents. This should at the very least make the attribution of 'creativity' to such cognitive states palatable.

modal condition entails P-novelty. Thus the second subsumes the first and we can get by with two conditions not three. Call this the *condition profligacy charge*.

This would be a reasonable charge only if the classes that meet the two conditions really were co-extensive, but they are not. Some ideas will meet the novelty condition while failing to meet the modal condition: I may form a novel thought but given my cognitive position I could have formed it before. So even if we grant the entailment from satisfaction of the modal condition to satisfaction of the novelty condition, we need not grant the entailment in the opposite direction: not all P-novel ideas are ones that *could not* have been had before by that agent. The two conditions are thus distinct in their extension.

One might counter by claiming that lack of co-extensionality is irrelevant. If the entailment runs one way—namely from satisfaction of the modal condition to satisfaction of the novelty condition—then that is all that matters. If any thought that satisfies the modal condition also satisfies the novelty condition, then satisfaction of the modal and agency conditions are enough and we needn't talk about novelty.

Our response at this point is to deny the entailment from satisfaction of the modal condition to satisfaction of the novelty condition. 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 which she has not before tokened. A satisfaction sentence for this condition takes this form: $(\exists x)(\exists y)(\exists t_i)(\forall t)[(t<t_i \supset \sim Hxyt)\ \&\ Hxyt_i]$, where 'x' is a cognitive agent, 'y' an idea, 't_i' some particular time, and 'H' is the relation to have or token. The modal condition is just the opposite, it is silent with respect to *actual cognitive profiles* concerning, instead, modal cognitive profiles. A satisfaction sentence

for this condition takes this form : $(\exists x)(\exists y)(\exists t_i)(\forall t)[t<t_i \supset \sim\Diamond Hxyt]$. Closing either of these sentences truly does not entail a true closure of the other. Disarming the condition profligacy charge is thus instructive, since it reveals the fundamental distinction between the two conditions. The novelty and modal conditions come apart with respect to *actual* thoughts and *possible* thoughts. The idea is of course that both conditions must be satisfied for an idea to be creative. And if the two conditions come apart as suggested, then we should keep them explicit in our analysis.

1.3.3 A worry about the strength of the modal condition

Here's another worry regarding the modal condition. "That's silly: surely Beethoven *could have* composed his 8th Symphony a day or two before he did, or even an hour or two before he did. He just didn't. Perhaps he was too tired or preoccupied. Perhaps he was shopping or traveling or chasing women. Perhaps his piano was being tuned. 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." There is a sense in which this is true: the world could have turned out slightly different such that Beethoven composed his 8th a few hours or days earlier. There is a close possible world where Beethoven composed the piece an hour earlier, another close possible world where he composed it two hours earlier, and so on.[9] It is important, however, to emphasize that the modal condition concerns a narrower modality than the broad metaphysical one just invoked. The modal condition invokes a nomological modality, targeting possibilities

[9] Although these worlds may not be as close as one might initially think. Even slight qualitative differences in a cognitive perspective may require significant changes to a world, such that the proximity of the relevant possible world, if there is one, is not so close after all. For alternative glosses on "nonactual epistemic possibility," see Simchen 2004.

relative to an actual cognitive perspective. What an agent *can* think depends upon what she *has* thought, *how* she has thought about it, and *how* she is able to think generally. A modal cognitive profile thus depends upon an actual cognitive profile: broadly, the intentional *contents* of cognitive states, the *characters* of states (e.g. beliefs or desires or auditory perceptions or visual images, and so on), plus general cognitive *capacities* (technical skills, reasoning abilities, and the like).[10]

Cognitive profiles may be individuated with greater or lesser fineness of grain. We might individuate a cognitive profile simply at the level of mental state tokens. This, however, would be too fine-grained for our purposes: some existing mental tokens will be relevant to the possibility of tokening some thought *c*, while others clearly will not. Thinking about environmental selection pressures may be relevant to the *possibility* of a breakthrough thesis in evolutionary theory, but beliefs or desires about the hockey game surely are not. Thoughts about the hockey game might of course trigger a string of thoughts that lead to the relevant thoughts about evolution, but the point is that other thoughts, perhaps about boxing or chocolate or whatever, *could have* initiated the same string. So although *some* tokens and their contents and characters will determine the modal status of some thought for an agent, equally important will be 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 broader method of individuating (parts of) cognitive profiles.[11] A creative thought *c* is causally contingent upon a chain of cognitive states. Some of those states are necessary for the tokening of *c*,

---

[10] The distinction between mental state *content* and mental state *character* is thoroughly discussed in chapter 4.

[11] See Currie for an application of this conceptual machinery to the ontology of artworks (Currie 1989: 46-84).

others could be omitted from the chain with no causal upshot for *c*. As a first sketch at least, the former set of thoughts constitutes the heuristic path to *c*, and it this set whose members must be either tokened *or* possible for an agent (possible in the sense that all of the enabling knowledge, beliefs, skills and other states and capacities are in place) in order for *c* to be a possible thought *for that agent*. A heuristic path will of course have its own enabling conditions—for example, the heuristic path to the discovery in evolutionary theory depends upon, at least, certain empirical and theoretical knowledge, skills of theoretical and statistical interpretation, and certain motivations—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 will be individuated at a broader level of description, one that describes certain organizations of states and capacities—something like a heuristic path or process. (Further analysis and application of the notion of heuristic paths or methods in creative thought is offered in 3.3.2, 5.1.1, and 5.2.3.)

So while it may be true that Beethoven might have composed his 8th at some earlier time, it is not true, given that Beethoven had a certain cognitive profile up to t_i, that Beethoven could have composed his 8th before t_i, where the 'could' is nomological. Given his actual cognitive profile, Beethoven could not have tokened the relevant cognitive states prior to t_i. The tokening of creative thoughts, like any thoughts perhaps, requires particular cognitive circumstances. Sometimes the only component missing is simply sustained attention to some task or other (maybe that is all Beethoven lacked and perhaps for the mundane reasons listed above); more often what's missing is some combination of cognitive states—beliefs, motivations, imaginings, concepts (maybe this

is what Beethoven lacked, namely, one or more component of or enabling condition for the causal process that led to his 8th); still other times some special training or knowledge must yet be acquired. Whatever the case, not until the right combination of any or all of these kinds of elements is in place, can an agent token a psychologically novel thought.

1.3.4 Modality, ex nihilism, and emergence

The modal condition goes some way towards addressing the various worries that fall under the rubric *creation ex nihilo*, which any account of creativity is under some obligation to address. These worries—which in spite of their antiquated flavor still have some grip on contemporary discussions of creativity, sometimes under the guise of the *paradox of novelty*—can be divided into strong and weak forms. The strong form says that creative ideas, bearing properties not had by their makers, must come out of nothing. Call this *strong ex nihilism*. This position can be understood by analogy with the traditional claim that the world contains properties not had by God—namely material ones—and thus God must have created the world *ex nihilo*. By this criterion, a creative thought could not be tokened before it actually was, since there is nothing in existence to produce it: creative thoughts, quite literally, come out of nowhere.

MC is inconsistent with strong ex nihilism. While the strong ex nihilist says that no existing process could produce genuine novelty, the modal condition requires of novelty certain circumstances of production. Not only are existing knowledge, beliefs, concepts, and other cognitive states instrumental in the production of creative thoughts, they are essential. The creative thought may involve some novel and praisable recombination or re-cognition of such materials, and it may involve some newly acquired

information or learning. Whatever the case, the idea did not come out of nowhere: it emerged from a particular cognitive background and necessarily so. This is a desirable consequence, since strong ex nihilism invites a dilemma, forcing us either into a supernatural explanation of creativity or, if one is naturalistically inclined, to an abandonment of the project of *explaining* creativity. *MC* and its modal condition reveal the hole in this dilemma.

But *MC* is consistent with *weak ex nihilism*, which holds that creative ideas and acts possess properties not had by the processes which led to their making and thus, in some sense, come out of nothing. The weak ex nihilist thus maintains a fairly robust notion of novelty, while allowing for a platform of production. The modal condition makes this platform necessary—creative thoughts require particular cognitive profiles and processes—but does not imply that the creative thought produced be reducible to the underlying processes and states. So some residue of *creation ex nihilo* remains, but any paradox is dissolved. A creative thought *c* is novel in the sense that it *emerges from* a certain cognitive organization O, and *c* cannot be identified with or reduced to O. But this is just a garden variety relation of emergence which, though it invites its share of problems, does not invite any strong form of ex nihilism. Philosophers of mind worry about downward causation, explanatory exclusion, and property dualism when addressing emergence but they do not worry that the emergent properties emerge from nothing!

1.4 Further issues

1.4.1 Constructive condition

Another condition, rooted in the thoughts of Henri Poincaré (1913), advanced by David Novitz (1999; 2003) and apparently endorsed by Berys Gaut (2003; see also Gaut and Livingston 2003: 10-11) and a number of psychologists (Martindale 1999; Sternberg and Lubart 1991, 1992, 1999), requires that creative acts or thoughts be of "real value" where this consists in the possession of "properties that are of actual or potential benefit to sentient beings: that either do or can increase enjoyment of life, enhance security, health, prosperity, and so on...The mad scientist who creates nothing but harm is ingeniously destructive but his actions are not properly described as creative" (Novitz 1999: 78). Thus x is creative only if x is constructive; destructive thoughts, acts, or artifacts should not or would not be described as creative. Whichever claim, normative or descriptive, the *constructive condition* is problematic.

First, it implies that the inventions of nuclear, chemical, and biological weaponry, among other weaponry technologies, are not creative or the products of creative thought. It is also at odds with many theories of art, namely any theory that requires that artistic or aesthetic value be, at least in part, intrinsic (Kant 1987; Beardsley 1958; see also Lamarque and Olsen 1994). These results are counterintuitive. Such inventions may strike us as dangerous or menacing to society, but these features do not obviously block attributions of creativity. And although the mentioned theories are not without their problems, we should not dismiss theories of intrinsic artistic or aesthetic value out of hand. But we can do better than just appeal to our intuitions: we can vindicate them by recognizing the following distinction. An act can be creative in the *productive* sense,

resulting in some useful artifact or progress. This sense of creativity derives straightforwardly from the verb 'to create.' Call this *creativity*$_p$. Alternatively, an act can be creative in the *novel* sense: some actions, events, and objects are judged (or just are) new relative to some standard of comparison. Call this *creativity*$_n$.

Only creativity$_p$ is beholden to the constructive condition. The invention of the atom bomb is not creative$_p$, since it is generally destructive of progress (though this too is contentious). Despite its destructive qualities, however, this invention may still be creative$_n$, this sense not inconsistent with being destructive. The invention of such an artifact involved surprisingly novel thoughts and actions (both in the psychological and historical senses of 'novelty'). To motivate the constructive condition, then, one must take one of the following three lines.

First, one could argue that creativity$_p$ is a condition on creativity$_n$. Perhaps to have something creative in the novel sense, that thing must be productive. If this is a fact we would expect it to surface in ordinary language, but it doesn't. Instead, linguistic practice betrays that we use these two distinct senses in distinct contexts, and neither use implies the other. Productive ideas can be old news; new ideas can be useless wastes of time. There may be considerable overlap in extension, but if the terms come apart at all then creativity$_p$ is not a condition on creativity$_n$.

The second option is to focus just on creativity$_p$. This is not viable, for the simple reason that we are concerned with, in making and analyzing creativity ascriptions, creativity$_n$. It is this sense of the term that we typically ascribe to Einstein, to *Guernica*, to the Beatles' *White Album*; it this sense that is the target of the present analysis (and the target of Novitz's and Gaut's analyses as well).

Maintaining the constructive condition is left to just one final option. One can motivate the condition and deny creativity to destructive acts only by shifting between creativity_n and creativity_p. This is clearly not a viable option. Equivocation is equivocation: we must use a term univocally throughout our analysis. The constructive condition should thus be rejected for the present purposes.[12]

1.4.2 Persons, processes, and products

The above discussion sometimes concerns agents, sometimes processes, and sometimes thoughts, actions, or things. In some ways, this is as it should be: we commonly ascribe creativity to all three P's—persons, processes, and products. There are undoubtedly important distinctions to be made between creativity ascriptions to these three sorts of things. The present analysis is concerned with the psychological phenomenon of creativity, thus primarily targeting thoughts, with some emphasis on the cognitive processes that enable certain kinds of thought. We might take creative thoughts to be the *products* of such *processes*. But perhaps we wish to reserve 'product' for whatever tangible artifact results from the thought, in which case we may need to tell a story for products different from the one given for thoughts. So maybe we should think of creative thoughts as instances of (cognitive) *process* or *processing*. No matter: this is a

[12] Commitment to conditions like the constructive condition (and some other evaluative conditions) may be symptomatic of the following mistake. When analyzing some term or concept Φ, it is easy to confuse the meaning of Φ with its use (or its use in particular contexts). That is, some property F may appear constitutive of Φ given that in some or many circumstances, attributions of Φ imply attributions of F. This, however, is a mistake if the goal is to identify literal meaning, as distinct from utterance meaning on the one hand and utterer's meaning on the other. So in some contexts of utterance, it may seem inappropriate to attribute creativity—say to atom bombs in discussions of world peace or global community—since it lacks a property of utility, namely, of furthering world progress and production. And indeed it is inappropriate to attribute creativity in such circumstances. But this only informs the way we *use* the term, in certain contexts, and does not entail anything about the literal *meaning* of the term. Thank you to Patrick Rysiew for emphasizing this distinction.

decision point we need not address. What's important is that we have with *MC* given an intuitive conceptual analysis of the psychology of creative cognition, identifying the minimal conditions that must be in place for creative thought.

Taking this analysis as basic, we can move on to other analysanda from here, to products like artworks and scientific theories, to the production of creative artifacts and to persons, perhaps geniuses. We might, for example, sketch an analysis of creative persons as follows. Here Boden's psychological/historical distinction proves especially important. P-creative persons are, presumably, just persons who token psychologically novel thoughts which they could not, given their cognitive perspective, have tokened before. (This might not be enough: we generally require more of creative persons, even in the psychological sense. So perhaps P-creative persons are ones who have P-creative ideas with some frequency.) H-creative persons are far more rare: they are persons who token P-novel thoughts which they could not have tokened given their cognitive perspective and, as a matter of fact, which have never been tokened by *anyone*. (This may not be enough: perhaps a person is H-creative only if he has H-creative ideas with considerable frequency. Thus people like Einstein, Newton, Mozart and Picasso would all be H-creative, while one-hit wonders would be relegated to, at best, P-creativity.)

1.5 Fundamentals: Descriptive or Evaluative Creativity

An important decision point for any analysis of 'creativity' is the choice between a descriptive or evaluative account of the concept. Can we describe creative persons, processes, and products in value-free ways? Most recent philosophical analyses have answered 'no' (Boden 2004; Gaut 2003; Gaut and Livingston 2003; Novitz 1999, 2003);

while some psychological analyses have answered 'yes' (Finke et al 1992, 1995). Consider our analysis *MC*:

> *MC:* Some x is minimally creative if and only if, for some agent A, x is the product of the agency of A; x is psychologically novel; and x could not have been tokened by agent A before the time t_i when it actually was tokened by A.

Does *MC* force us into an evaluative model of creativity? One might think the answer is affirmative here given the appeal to concepts like novelty.

However, this last inference would be too quick. Both the agency and novelty conditions can be understood in purely descriptive terms. The agency condition is intended to capture only the responsibility for the act or process: x is a product of agency only if x is something for which an agent is non-trivially responsible. This feature of thought and action can be assessed in purely descriptive terms. Borrowing a point from Boden, the novelty condition is narrowed to capture *psychological novelty*. P-novelty consists in novelty to some individual mind: an idea c tokened by an agent A is P-novel only if c has not been tokened by A before. This is a strictly descriptive matter: an agent either had an idea in the past or she didn't. Roughly the same point can be made about the modal condition: it is a descriptive matter whether or not an agent *could have* tokened an idea given her cognitive perspective. Thus none of the three conceptual conditions commit the present analysis to an evaluative account.

Grant that the account is descriptive. We then must say why a descriptive account is useful. Remaining at the descriptive level of analysis allows us to identify the fundamental psychological conditions required of creativity, while remaining consistent with but non-committal to narrower or higher level analyses of creativity. Why think that

this analysis is fundamental? Recall that *MC* says that creativity is ascribed to thoughts that are, at minimum, products of agency, psychologically novel, and which could not have occurred relative to the cognitive perspective and time in question. These conditions are fundamental to ascriptions of creativity given the kind of property that creativity is. If one is attributing creativity to an artifact, an idea, or a person, one is attributing, in part, some kind of cognitive agency. (Which is not to imply that we attribute agency to artifacts or ideas: if the attributee is an artifact or idea, then agency is attributed to the producer or possessor.) One is also attributing a kind of psychological novelty, comprising both an actual and a modal component. These properties are constitutive of creativity. Put another way, creativity is conceptually bound to agency and novelty. Creativity is thus a higher order, emergent property, and *MC* is a fundamental analysis since it identifies the properties fundamental to the higher order one. So if I attribute creativity to *x*, I attribute, at least, a satisfaction of *MC*.

Minimal creativity is neither a technical notion, nor one reserved for scientific experts. The folk certainly use creativity in this sense. I may ascribe creativity to a child learning and combining new concepts or to a school project that she might bring home. You might ascribe creativity to a cab driver who takes an unexpected shortcut to avoid the Chicago traffic. We call television, radio, and periodical advertisements 'creative' without second thought. Spectators might ascribe creativity to a chess-player who uses a piece in a slightly unorthodox way. Some of these ascriptions may be more or less evaluative. But they can all be described in the terms of *MC*: in each case, the property ascribed comprises both agency and a particular kind of novelty. These ascriptions no doubt require of the folk some recognition of the psychology of the attributee (or of the

agent responsible for the attributee, i.e. the product), but this requires no more than folk psychological ability. The controversial *nature* of mindreading and metarepresentational capacity notwithstanding, one thing we do know is that we *do have* such abilities.

Finally, a descriptive analysis of the fundamentals of creativity is a powerful and flexible research strategy since it allows us to say something useful about the psychological phenomenon of creativity, without getting stuck in certain debates in art and aesthetic theory, general value theory, and philosophy of science. With respect to creativity in art, for example, we avoid debates about intrinsic versus instrumental value, the significance of art-historical context, or artistic intention. One may choose to give analyses of more specialized, evaluative, or idiosyncratic senses (and uses) of 'creativity.' Assuming that they are senses of the same phenomenon, *MC* will be consistent with such analyses and, in fact, provide them with a useful starting point. Once the psychological conditions are in hand, one can build in evaluative and contextual concerns, sociological-historical factors, features of the theoretical or scientific domain, and more particular features of the creative agent's psychology and environment. 'Creativity' may often be used in these other ways, many of them value-laden, but minimal creativity as defined is fundamental to such uses.

Chapter 2
Cognitivist extremes

Our analysis of creativity is cognitivist in motivation: it centers around thinking and processes of thought. It is also naturalistic: all components of the analysis are consistent with, if not supported by, the natural sciences. These choices motivate the following theoretical constraints.

A cognitivist account is constrained to cognitive processes in its analysans. Call this the *cognition constraint*. Appeal to this constraint immediately strikes inspirationalist views of creativity from consideration. There will be no room in a cognitivist model for divine inspiration or magical insight. Thinking creatively is just that: *thinking* creatively. (Recall from 1.1.2 that we have independent reasons for abandoning inspirationalism.)

A general constraint on metaphysical theories is that one is parsimonious in postulating entities. This constraint carries over to the metaphysics of mind. Call this, quite simply, the *parsimony constraint*. The parsimony constraint may be coupled with the *naturalistic constraint*, which derives from the empirical evidence offered by the natural sciences. This coupling casts considerable doubt upon the existence of unique creative mechanisms, modules, or capacities. If creative processes are cognitive processes then they, like other better understood processes such as belief, desire, and decision making, are likely embedded and intertwined with the various components of the cognitive and perceptual system. This sort of holism, despite its deficiencies, remains an attractive strategy for modeling the mind. So the likely scenario is that creative thought involves a number of mechanisms also used for other cognitive processes.

2.1 Cognitivist extremes

A cognitivist theory assumes a very simple explanatory strategy: whatever is to be explained is to be explained in terms of cognitive processing, in terms of thought processes that typically lead to or underwrite knowledge and awareness of the world. A cognitivist model of creativity thus takes creative processes to be explainable in terms of thinking.

Of the few cognitivist theories of creativity on offer, there are two extremes. On one end, we have theories that posit either special creativity-specific or creativity-needed mechanism(s). The idea here is simple. In its strong form, such a model proposes that creative thoughts are enabled only by a special, modularized mechanism whose only function is the realization of creative processing. The weaker form says that creative thoughts require *some* special, fairly modularized mechanism, one of the functions of which is the enabling of such thought. Call any such model, strong or weak, a *special mechanism model.* The motivations for such models include providing a comprehensive and fairly domain-specific theory of mind, accommodating the wide disparity in creative capacity between human beings and all other species, and accommodating the ostensibly wide disparity in creative productivity among humans. A special mechanism approach meets all such desiderata: a creativity mechanism processes input that requires novel (kinds of) processing; humans have this mechanism while other species either do not or have a much less sophisticated mechanism; some humans have a more developed or sophisticated mechanism than others. Some additional empirical motivations for the special mechanism approach are considered below.

On the other extreme are models that deny any special mechanism, and moreover deny any especially creative capacity or ability. Creative abilities are, at most, just more sophisticated, enhanced, or developed cognitive abilities. Call any such model a *nothin' special model*. Such models share the motivations that the special mechanism models possess, with the subtraction of the domain-specificity motivation and the addition of greater emphasis on parsimony. Nothin' special models accommodate such motivations equally well: they explain creative capacities in largely domain-general ways; make fewer, or at least less specific, commitments by way of brain architecture; and explain human-to-other species and human-to-human differences in terms of more and less sophisticated cognitive capacities for problem solving and learning.

Problems arise for these models both where they match and mismatch in motivation. But most of the problems lie at the extremities of their respective treatments. To see this, however, we must first consider each extreme a bit more closely.

2.2 Special mechanism models

The Darwinian model of psychologist Donald T. Campbell, recently advanced by Dean Simonton, provides an example of a strong special mechanism model of creativity (Campbell 1960, 1965; Simonton 1999). The model comprises three basic conditions. First, there must be what we might call a *variation-generator* in the cognitive system: a mechanism that, by analogy with Darwinian evolution and genetic mutation and variation, produces variations on ideas and concepts. Second, there must be a *selection-mechanism*. This mechanism, by analogy with the role of natural or sexual selection in biological evolution, selects ideas generated by the variation-generator. Finally, there

must be a *retention mechanism*: by analogy with genetic inheritance, creative ideas are the ones retained and stored by memory, and passed on further through communication. The second and third mechanisms can be accommodated easily enough by other cognitive mechanisms—some kind of reflection, metacognition, or introspection for the second and memory for the third. But the first mechanism, the variation-generator, seems to be a creativity specific mechanism.

Psychologist Colin Martindale argues for what he calls the *Cortical Arousal Theory*, which centers around the nature of focuses of attention (Martindale 1977, 1981, 1995, 1999; Martindale and Armstrong 1974; Martindale and Hines 1975). He proposes a multi-stage model of problem solving, which if the right mechanism is possessed, leads to creative thought. In the initial stages, information is gathered, various approaches are taken to the problem, and there is a high level of cortical arousal with a narrow focus of attention. As information increases and the problem remains unsolved, two kinds of responses may occur. The first kind of response is to keep attempting the same solutions to the problem such that the arousal and attention focus stay high and narrow, respectively. Alternatively, some persons experience a decrease in cortical arousal coupled with a wider range of attention focus. Information then enters what Martindale calls *primary processing*: a kind of subconscious cognition not under the complete control of the agent. It is this kind of processing, and the arousal mechanisms that enable it, that distinguish creative people from non-creative. The first kind of response typically results in frustration and failure, while the second often results in creative insight. (Further discussion and criticism of the cortical arousal theory is offered in chapter 6).

Both examples of special mechanism model are vulnerable to criticisms. Both fail to meet, or at least do not meet well, the parsimony constraint on cognitivist theories of mind. The Darwinian model posits at least one special creativity mechanism: the variation generator; the Cortical arousal theory posits a special arousal mechanism which, at least under some interpretations, is a creativity-specific mechanism.[13] Moreover, both models should be considered with respect to the naturalistic constraint: one might question the empirical plausibility of the evolution or development of a creativity-specific mechanism. These criticisms are not decisive, but they are counts against the respective theories.

The Darwinian model in particular is open to the following criticism. If creative ideas are initiated by a particular cognitive mechanism, the variation generator, then there is an important sense in which responsibility is stripped from the thinker. That is, if we stick with the analogy with genetic mutation, variations are generated blindly, and thus are not to be credited to the thinker. These cognitive circumstances fail to meet the agency condition of *MC*: randomly generated ideas are not ones that result from agency (see 1.1).

The cortical arousal theory posits a special mechanism, and it proposes that this mechanism marks the distinction between creative persons and all other persons. So we really have here a theory of genius or creative persons, and not a theory of more mundane, psychological creativity. This invites problems. First, the theory demarcates creative acts by less creative persons, or what we might call one-shot creative acts, as being of a kind different from the acts of consistently creative persons. This is

[13] This is most clear in the way that Martindale distinguishes creative from noncreative persons. See below and 6.2.4.

problematic for conceptual reasons: we use 'creative' to describe both kinds of act, irrespective of the frequency with which their producers typically produce such acts. It is also problematic for empirical reasons: the theory implies that creative persons develop or inherit a cognitive structure that most of us simply lack. This invokes a heavy explanatory burden. So the benefits of the cortical arousal theory come at a high theoretical price.

2.3 Nothin' special models

At the other extreme are the nothin' special models. Perhaps the best example is Robert Weisberg's incremental problem solving model (Weisberg 1986). Weisberg is especially critical of special mechanism views, dispelling what he takes to be a number of myths regarding creativity.[14] Creativity does not involve the use of special creative abilities or mechanisms; there is no special incubation or tacit cognition stage in the creative process; so-called geniuses are not different in kind from you and me. Instead, Weisberg proposes that creative thinking consists in normal, persistent problem solving, requiring intelligence plus skills and knowledge specific to the problem space in question. Thus much of what we say and what is theorized about creativity is symptomatic of folk myths and hasty science and is, strictly speaking, false.

Nothin' special models like Weisberg's are also vulnerable to a number of criticisms. While the special mechanism models have difficulty explaining humdrum or minimal psychological creativity, the nothin' special models are impoverished in the other direction: they say little to distinguish radically creative persons and processes from

[14] For a useful comparison of the Cortical Arousal Theory and Weisberg's theory, see Partridge and Rowe (2002).

less creative ones. That is, their explanation is too coarse grained to accommodate historical-creativity and the wide variance in psychological creativity. Second, *reductions* of creativity to basic problem solving are easily undercut by counterexamples. Examples where we have one and not the other are not difficult to come by. Think of the brilliant artist who is embarrassingly poor at solving simple puzzles; inversely, many people are quite good at solving such problems, say in math or logic, but seem entirely devoid of creative abilities. Theories like Weisberg's might of course be finessed around such counterexamples, but perhaps not without either loss of explanatory power or inheriting somethin' special in their explanation.

Weisberg's theory also fails to accommodate an important feature of creativity: flash phenomenology (what he calls the "A-ha myth"). His move is as follows. If creative ideas sometimes come in flashes of insight, then they must be disconnected from previous knowledge or memory. Creative ideas are not disconnected from previous knowledge or memory but instead require them. Therefore, creative ideas do not come in flashes of insight and can instead be explained in terms of basic problem solving. This inference is both too fast and bears consequences we should want to avoid. We can and should accept Weisberg's claim about the necessary connections between knowledge and creative thoughts, but we can do that without accepting his dismissal of flash phenomenology. We block his modus tollens by denying the conditional on which it operates. One can admit flash phenomenology into the explanandum—since it is a point about the phenomenology and not the metaphysics, or the cognitive architecture—while maintaining that creative ideas emerge from and are thus logically bound to a particular

epistemic perspective. We thus meet the naturalistic constraint that cognitivism imposes upon us, without dismissing a pre-theoretical datum of creativity.

2.4 Lessons learned

These discussions suffice to show that neither cognitivist extreme is to be preferred. However, each of the models considered has virtues, answering questions and filling holes left open by the alternatives. The Darwinian model identifies a role for introspection or metacognition, rendering creative processes higher-order cognitive processes. The cortical arousal theory addresses questions regarding flash phenomenology, which Weisberg resolutely dismisses. By the same token, Weisberg emphasizes the importance of background and domain knowledge, and thus locates the responsibility of creative acts in the agent. Weisberg also rightly identifies the similarities between creative cognition and problem solving; his error is to reduce the first to the second. These considerations inform our modeling strategy: an analysis can learn from the extremes without itself becoming excessive. A hybrid approach is preferable: we need opt neither for ultra-specialization nor ultra-non-specialization to explain creative processes. Just such a model is developed in the chapters that follow, splitting the difference between creativity-specific mechanisms and nothin' special skepticism.

Chapter 3
Minimal Creativity and Problem Solving: A Functional Analysis

3.1 Problems, tasks, and spaces

A strategy for analyzing creative processes is to model such processes on more standard problem solving tasks. Weisberg takes this position to its extreme, but a number of cognitive scientists working on artificial intelligence have used similar models, while making weaker theoretical claims (Buchanan et al. 1976; Hofstadter and FARG 1995; Langley et al 1987; Lindsay et al 1980; see also Boden 2004: 199-232). This kind of proposal is often met with skepticism. Here is one such response. Problem solving involves a determinate problem and a determinate solution. Creative persons, be they painters or chemists, do not have in mind the particular solution or result they end up with when they set out to do their work. In fact, such persons often do not even have a determinate problem and are thus logically precluded from having a determinate solution. Therefore, creative processes should not be modeled on problem solving processes. Call this the *determinacy worry*.

3.1.1 A functional analysis

The intuition behind this bit of reasoning is a good one: there is more to creative processes than mere problem solving. But the reasoning is flawed. The flaw consists in an overly narrow construal of problem solving frameworks. The assumption is that problem solving always takes the following form: an agent has a problem and the problem entails a solution or a set of solutions. To solve the problem, the agent thus must locate or find the solution or one of the elements of the set of solutions. Problem solving

is thus discovery, not creation. But surely problem solving should not be so narrowly understood. A scientist *may have* a problem that takes the above schema, in which case she just sets to work to discover the solution. This process may be more or less creative. It is often the case though, in art, science, and other domains, that the problem is not determinate and so does not determine the solution. Instead, a scientist may have some phenomenon—itself an analysandum in need of clarification—and wish to analyze and explain that phenomenon in the terms of her particular science. An artist may have a certain emotion he wishes to express, without any terribly lucid notion of how to express it or even what emotion it is. A marketing strategy team may have a set of products and a set of markets with the simple goal of selling the first to the second. Here the problems and solutions vary in determinacy. One might respond that these are *tasks* and not *problems*, and the second is a different kind of thing from the first. However, it is better to think of tasks and problems as same in kind, differing at most only in matter of determinacy. We see this by thinking of problem solving in functional terms.

Consider the following very simple theory of problem solving.

> (T) Given some agent *x* in some **context**, the **context** presents *x* with a **problem**, which requires for its completion some **solution set** achieved via some **heuristic set**.

Some preliminary notes about our theory. First, it is a minimal account of problem solving: it only identifies, at most, features fundamental to problem solving, without treating any subtle nuances or offering any higher level analysis. No matter: we are only after a basic structure for the environments that force the solving of problems upon their occupants. Moreover, this theory is founded on a commonsense understanding of problem solving: it is not a scientific theory. Second, this theory may not tell us much

but it does imply that the various elements are causally related in important ways. For example, problems are not the sole determinant of solutions: problems are bound to contexts and agents, where all three elements constrain the range of possible solutions and heuristics. The theory thus allows for problems that, depending upon how the other variables are filled, may be solved in only one way, and problems that may be solved in a variety of ways.[15] Finally, one will note that some of the terms used (those in **bold**) are in fact the very terms that one would hope a theory of problem solving to illuminate. So at least as it stands, this theory is circular and of little significant value.

Using the Ramsey-Lewis method (Ramsey 1929; Lewis 1970, 1972), we can translate (T) into its Ramsey sentence by replacing each theoretical term with an open variable and existentially generalizing.

> (T_R) There exists some P_1, P_2, P_3, P_4, and some x, such that x is in P_1, and P_1 presents x with P_2, which requires for its completion P_3 achieved via P_4.

Or, if one prefers, we can formalize (T_R) even further, following the method of Ned Block (Block 1980a, 1980b).

$$(\exists P_1)(\exists P_2)(\exists P_3)(\exists P_4)(\exists x)\ [T(P_1, P_2, P_3, P_4)\ \&\ x \text{ is in } P_1]$$

Block's level of formalization is not, however, necessary for our purposes; we will stick with (T_R).

The advantage of Ramseifying—translating (T) to (T_R)—is that we eliminate all theoretical terms (T-talk) that might themselves be in need of further explication. We replace the theoretical terms with open variables while keeping the non-theoretical

---

[15] So note that in using the terms 'solution set' and 'heuristic set,' we are not requiring that problems be multiply soluble nor soluble via a variety of heuristics: the sets in a given context may contain many members, contain only one member, or even be empty (in the case the problem is in fact insoluble for the agent in question). Using 'set' allows us to keep all such options open.

terminology constant. We thus maintain the causal relations between the constants and the variables. The theoretical purchase of this method is that we identify the functional role of each of the theoretical elements of (T).[16] 'Context' is replaced with the functional role P_1; 'problem' with P_2; 'solution set' with P_3; 'heuristic set' with P_4. This exercise in functional definition affords many lessons.

First, we find that these functional roles are not uniquely realizable. Consider functional role P_2. We should immediately see that if we occupy P_2 with 'task' or 'challenge', while filling the other roles as in (T), (T_R) remains intact. In other words, we can close the formula in any of these ways without affecting the truth of the theory: if (T_R) is true for problems, it is true for tasks or challenges. Reverting back to the talk of (T), things that answer to 'tasks' and 'challenges' relate, according to their contexts, to solutions and heuristics in the same ways as problems. The flexibility, as it were, of this role generalizes to the other three. And we reveal this simply by identifying the functional role of 'problem' and the other theoretical terms.

We also learn that we can close the formula in more or less determinate ways without affecting the integrity of the structure. For example, we might occupy P_2 in a very determinate way and, presumably, affect how P_3 and P_4 can be occupied.[17]

---

[16] 'Function' is a term used and misused in a variety of ways. To be clear, the present usage is non-teleological and non-commital with respect to design. So to attribute a function is just to attribute a certain causal role to an element situated in a system. This is the sense of 'function' posited by Robert Cummins (1975), who claims that functions are dispositions or systematic relations that an element bears with regard to other elements in some larger system. We appeal to Cummins-functions in order to explain *how* the larger system works in the ways that it does, but not *why* it works in such ways.

[17] Two terminological notes. One, we must distinguish the metaphysical and epistemic senses of 'determinacy.' In some problem scenario, the problem and solution might, as a matter of fact, be entirely (metaphysically) determinate. But for present concerns, the metaphysics of problem determinacy is less relevant than the epistemology of problem determinacy. What matters is whether a problem is determinate in the sense of being clearly formulated for some cognitive agent, and likewise for a solution. Thus problem and solution determinacy, in this epistemic sense, admit of degree and are agent-relative. Two, terms like 'filling' and 'occupying' are used to indicate that an open variable—which in the case of our

However, nothing about the functional relations between these roles requires that the determinacy of one will increase the determinacy of the other (and the same can be said for lessening the determinacy of one). All the analysis tells us is that the determinacy of the roles depends upon the determinacy of the others, but it does not specify the ratio or mapping of such relations. Reverting back to our theoretical terms, the point is just that a more or less determinate problem will affect the determinacy of the solution set and the heuristic set, but could do so in a nearly infinite number of ways. This serves to dispel the determinacy worry, which proposed that problem solving requires wholly determinate problems and thus equally determinate solutions, and creativity involves neither. Our functional analysis reveals that there is no such conceptual requirement on problem solving. Less determinate problems, tasks or challenges relate to the other elements in (T_R) in the proposed ways, thus fulfilling the role P_2. Problem and solution determinacy are not necessary marks of basic problem solving frameworks.

We must be very clear here: (T) and (T_R) do not, to be sure, constitute a theory of creativity or even of creative problem solving. One can truthfully close the formula (T_R) in a variety of ways that do not amount to creative achievement: think of chickens faced with the challenge of crossing the road. So we are not about to limit our model to a functionalist theory of creativity, since it is unlikely to be fruitful. However, we have, with our little theory, revealed the basic functional structure of problem solving. What we have found is that this structure is fundamentally similar to that of creative thinking. We have thus dispelled some of the skeptical concerns about modeling creativity on problem solving. (T) and its functional equivalent (T_R) provide a framework against

---

(T_R) is a functional role—has been named. The term 'closed' is used to indicate that each open variable of the formula (T_R) has been named.

which we can analyze the various components of creative cognition, supplementing where distinctions from problem solving are needed. We will continue by looking more closely at P_2 and its relation to the other elements in (T_R).

3.1.2 Creativity without problems

Consider a gloss on the determinacy worry that goes like this. Consider situations where there seems to be *no* determinate problem or solution whatsoever, yet some resulting novelty. This is after all, one might think, what we have in the case of creating art. A painter or artist is not addressing some more or less determinate problem. Rather, she is simply *doing* her art: using her particular medium to create some artifact. The result is just that, but it is no solution since there is no problem being attended. (T_R) is surely not true of these kinds of cases and thus does not accommodate much of what we call creative. Much of this criticism simply misses the mark: artists often do think in very problem-oriented terms (Baxandall 1986). "How do I express this thought in two lines?" "How can I photograph that scene so as to accentuate the shadowed portions?" "What can I use to get a rusty, nostalgic effect in the third verse of the song?" "What color best contrasts with this part of the canvass?" We could list thousands of like questions, all of them posing problems to their askers. But even if we set these aside for the moment, we can still see how the proposed "problem-free" cases will fit (T_R).

Boden provides an analogy that is especially helpful here. When a task or problem—the occupant of P_2 in (T_R)—is very general, or perhaps even indeterminate, we should think of the action as one of exploration of conceptual space analogous to undestined geographical exploration. "Some explorers of planet Earth seek something

specific: Eldorado, or the source of the Nile. But many simply aim to find out 'what's there': how far does the plain extend, and what happens to this river when it gets there?; is this an island?; what lies beyond the mountain-range? Likewise, the artist or scientist may explore a certain style of thinking so as to uncover its potential and identify its limits" (Boden 2004: 59). So here one does not have a clearly defined problem to address but, just like the undestined explorer, a certain *space* to explore. Boden refers to this space as the conceptual space; we might call it, less specifically, a problem space. In any case, here the occupant of P_2 is just the exploration of this space. One is challenged with the task of exploring, mapping, and comparing the terrain of the relevant problem space. We can thus reasonably say that (T_R) is satisfied even in situations involving no determinate problem.

In cases like these, where the occupant of role P_2 is very loosely determined—or to use Boden's terms: where the task or problem just is conceptual exploration—we should note how important the occupant of P_1 becomes. The context in which the agent is situated—or better, a sub-space of that context—is the space to be explored. Thus if the task is to explore *that* space, then that space will importantly constrain what might count as a solution/s and the heuristics available for reaching it/them. If P_1 and P_2 are filled as we have been discussing, then in virtue of situation in some context (the occupant of P_1), *x* is presented with the task (the occupant of P_2) of exploring that context or some proper part of it. What can occupy P_3 and P_4, is thus constrained by the occupant of P_1, while the occupant of P_2 derives (almost entirely) any constraints it places on P_3 and P_4 from

P_1.[18] This is true (to varying degrees) when P_2 is more determinate (put in T-talk, when we have a clearly defined problem). But consideration of cases where the occupant of P_2 is less/non-determinate makes more salient the importance of P_1. In the following section, we will consider the general constraints that P_1 places on P_3 and P_4.

3.2 Context and epistemic constraints

So we are now considering an instance of (T_R) where the occupant of P_2 just is the exploration of a problem space—some subset of the relevant context (the occupant of P_1). We might ask, for some agent *x* situated in this context, what solutions and heuristic methods might be appropriate. In other words, in such cases what constraints does the occupant of P_1 place on *x* and how the roles P_3 and P_4 may be filled? Let's consider Boden's undestined exploration analogy once more.

3.2.1 Undestined conceptual exploration

In undestined exploration, explorers often possess a map of the territory to be explored. Other times they must make one as they go, either not having a pre-made map to start with, or not having one desirable or suitable given other considerations. (Destined exploration may involve either as well: sometimes we know where we are going, but don't have a map to get there. Other times we have both the knowledge and the map.)

As Boden indicates, having a pre-made map allows us to revisit places we've already been. It also tells us where we can go and, roughly, how to get there. Finally, it

[18] In discussions of particular satisfactions of (T_R), the open variables 'P_1', etc. (identifying functional roles) are sometimes used as names for constants rather than the more appropriate 'occupant of P_1.' This is done merely for convenience, and will be avoided in instances where greater precision is required.

contains information regarding the ways of exploration that are possible for the mapped territory as it is mapped—different routes, segments, timings, kinds of terrain, and so on. It is in this sense that the map constrains the exploration. Boden will unpack this analogy with respect to conceptual spaces in terms of generative rules. "These maps of the mind, which are themselves *in* the mind, are generative systems that guide thought and action into some paths but not others" (Boden 2004: 59).

We must be careful in how we read Boden here: the conceptual maps are *in the mind*, but typically only insofar as a person has knowledge or understanding of some objective conceptual map of some objective conceptual space. In the paragraphs that follow this last quotation, she tells us that theoretical maps, in art and science, help their users to explore the mapped domains. These maps are typically not the product of one single mind, but instead emerge from some particular theoretical domain, be it chemistry, biology, or realist painting. So we want to distinguish the objective constraint on exploring or thinking in some domain—provided by the broadly agent-independent map—and the subjective constraint on such thinking—provided by the relevant agent's knowledge of the map and other cognitive factors specific to that agent. As argued in 1.2.2, Boden goes wrong to relativize P-creativity to the first, to some agent-independent set of generative rules. But we can nonetheless acknowledge the general importance of such rules—namely, the constraints they place on one's epistemic perspective—for any instances of (T_R), creative or not.[19] The relevant question to ask is this: when there is a

[19] Here and elsewhere, I understand 'epistemic perspective' and 'epistemic constraint' broadly. Epistemic agents—that is, organisms who possess knowledge as well as other cognitive and perceptual states—are situated evidentially somewhere in the world, and possess certain states representative of the world, which are thus importantly constrained by the world. So epistemic perspectives are not limited to sets of knowledge and epistemic constraints are not exclusively constraints on knowledge, but rather are perspectives had by and constraints on knowledge-possessing agents.

conceptual map of some domain available, what kinds of constraints will it place upon some agent situated in that domain? A distinct but closely related question is: what kinds of constraints does the space or domain itself place upon the agent?[20] This recapitulates the question we began this section with: for some agent *x* situated in P_1 which presents x with P_2—where here P_2 is occupied by the task of exploring some subset of P_1—what constraints does P_1 place on *x* and how P_3 and P_4 may be filled?

Most simply, such a problem solving circumstance is one where the occupant of P_1 significantly constrains what *can* occupy roles P_3 and P_4. That is, the conceptual map and domain will constrain the kinds of thoughts that an agent can have *about that domain* (insofar as she is working within that domain). So if the task is one of exploring some space, the map will contain a finite amount of information and some guidelines on how the map represents what is mapped and how the map is to be used, and the job of the map-wielding explorer is thus to acquire some quantity of this information and use it to explore the space. So exploring will involve, for a start, both the acquisition and use of propositional knowledge, as well as the acquisition and use of some procedural knowledge. Minus the exploration-talk, thinking within some conceptual domain, say neurobiology or modal jazz, is importantly constrained by that domain, both in terms of what one can know about the domain as it is currently conceptualized and what skills can be learned from and used within the domain. (It will also be constrained by the explorer.

[20] To be clear, conceptual maps and conceptual spaces (or domains) can be distinguished thus. A conceptual map is the information made more or less explicit about some domain. The relevant conceptual space will include this, but will also include logical (or perhaps natural, as 'logical' may not always be the appropriate term here) consequences of the explicit content of the conceptual map. There will typically always be some features of the space that have not yet been made explicit, that is, not yet mapped. So maps only represent some proper subset of the conceptual space.

More on that below.) This should not force us, however, to conclude that domain-bound thinking is precluded from being creative, as we will soon see.

3.2.2 Map-using, map-changing, and map-making

With respect to some conceptual space, an agent may work on the map, off the map, or even off the space. Consider the following examples.

Example 1: I have a small problem. I need to get from 10th and Alma on the westside of Vancouver to 41st and Main on the eastside of Vancouver. Being a rather conservative traveler, I consult a map and find a simple route from my present location to my destination. My method for solving my problem and thus my solution are constrained by the context, Vancouver, and the map of that context. My method is just to read my solution right off of the map. This is a simple case of domain-constrained thinking. The problem (P_2), context (P_1), and my personality dictate that I consult a city map and secure an easy route (P_4) which gets me where I need to go (P_3).

Example 2: Norm has a crush on Eigen. So, naturally, Norm wants to impress Eigen. Impressing Eigen, however, won't be easy: she is a quantum physicist who, so far as Norm can tell, only has eyes for superposition. So what is Norm, a sober-minded sociologist, to do? The obvious answer seems to be to learn some quantum theory. So Norm begins spending his nights sweating over quantum mechanics textbooks and journals, and attending the colloquium series of the physics department on campus. After many sleepless nights and several lost hairs, Norm makes significant progress. He is ready to make his move. Now, never mind whether Norm succeeds in wooing Eigen. The important point is to note the occupants of the functional roles in this instancing of

(T_R). Norm, situated in a complex social and professional context (P_1), is faced with the problem (P_2) of impressing ultra-scientific Eigen. All of these factors contribute to the rather academic method (P_4) employed to secure a solution (P_3). Certain features of this situation constrain what Norm can do to solve his problem—most importantly, Eigen's being a tunnel-visioned member of a very specialized sector of the scientific world. Norm, a vector space virgin, thus must take a rather formulaic approach to tackling his problem.

Example 3: By the early 1960's, epistemologists were quite happy with themselves. Ostensibly, some consensus about an ancient philosophical problem—the nature of knowledge—was cementing. There were still disputes over the details, but most seemed to agree: a bit of knowledge consisted in a belief that was epistemically justified and true. Call this collection of theories *JTB*. Then, in 1963, along came a gentleman named Edmund Gettier who, according to a former professor of mine, was a "veritable counterexample factory." In an article of no more than three pages, Gettier called into question the fundamental approach of JTB, showing that it was at least incomplete if not simply mistaken. His method? Thought experiments and counterexamples: Gettier took the JTB analysis at face value and employed two of the oldest of theoretical methods. He considered two thought experiments, both of which satisfied JTB's purported criteria of knowledge and both of which clearly did not amount to knowledge. These two counterexamples constitute what is today known as *the Gettier Problem* (Gettier 1963). Note that Gettier was strictly constrained by the relevant context (P_1), as he was essentially just exploring a philosophical domain, and a particular theory and its consequences for that domain (P_2) by the basic methods (P_4) considered fair

within that domain. In other words, he challenged JTB from within the theory by playing the game only as the relevant theorists played it, securing a result (P_3) that ultimately changed the game.

Example 4: Consider Bach's famous work, *The Well-Tempered Clavier*, which consists of Book I and Book II, each one comprising a set of one prelude and one fugue devoted to each of the twelve major and minor keys (24 pieces in each Book). Part of Bach's goal or task, as evidenced by the title of the work, was to explore the possibilities of both tempered tuning for keyboard instruments—this tuning method was invented during Bach's time by Andreas Werckmeister—and the newly established tonal scale system.[21] In fact, some theorists hold that he created these works for largely didactic purposes (Tomita 1996, 1998). We thus have a problem (P_2) that is importantly constrained by the problem space (P_1), where in a very literal sense the problem just was the exploration of that space. Bach thus constrained himself both to a certain music-theoretical structure and to certain bits of musical technology. Thus the methods (P_4) and final results (P_3) had to meet such constraints.

These four examples are different in several important ways, but they share the following feature. They are all problem solving circumstances where the occupants of P_1 and P_2 significantly constrain what *can* occupy roles P_3 and P_4. In T-talk, the problem and problem space—where the problem is more or less determined (compare example 1 to example 4)—limit the agent with respect to the heuristic methods she may employ for a solution. What we find is that even when P_1 and P_2 are fixed and P_3 and P_4 are

[21] 'Clavier' during Bach's time just meant keyboard instrument. However, there is evidence that Bach made a more fine-grained distinction, distinguishing organs from smaller-stringed instruments and categorizing only the latter as *claviers*. This category would include the clavichord, fortepiano and harpsichord, which are typically characterized by a somewhat delicate timbre, sharp sound, and small size. See Tomita 1996; 1998.

constrained accordingly, for some *x* situated in P_1, there is often significant freedom with respect to how P_3 and P_4 are filled. In other words, a problem space may broadly constrain the kinds of thoughts that *x* might form in addressing that space, but it leaves open to some degree the ways in which those contents are combined, employed, and manipulated. A variety of methods or heuristics may be used by *x* in achieving her results.

This should strike us as an intuitive consequence, since it is consistent with the following general point. Given some sizeable set of information {I}, different agents will acquire different bits of knowledge about {I}. Moreover, not all thoughts about {I} will amount to knowledge. Some will turn out to be false beliefs, some true but unjustified, others both false and unjustified. Still others will not be doxastic states at all: one might have entertainings, desires, intentions, curiosities, doubts, and imaginings, among other states, with respect to {I}. It is a mistake to assume that all cognitively valuable states regarding some set of information—some conceptual domain—will be knowledge-bound or even truth-bound. Instead, the *knowledge* had with respect to {I} will derive solely from {I}, since to be knowledge of {I} some belief or skill will have to accurately track some element of {I} (either in the sense of being true for propositional knowledge or accurate/applicable for procedural knowledge). However, there are a variety of other *non-truth bound* cognitive states that one might have with respect to {I}. To recognize this is to recognize that there are a number of cognitive states that are considerably less constrained by {I}. These states are crucially important, we will see, for creative thinking.

The Bach example clearly illustrates this point. The *Well-Tempered Clavier* is clearly an exploration of some conceptual space, and is one that stays within the boundaries of such space. Bach thus had to know a lot about the 12 major and minor keys and how to navigate them, about tempered tuning, and about keyboard instruments. However, this knowledge—both propositional and procedural—is surely insufficient for the creation of Bach's masterpiece: the *Well-Tempered Clavier* is not, as it were, a logical consequence of the conceptual space. Bach's heuristic method consisted in something more than mere knowledge of the space and how to use it. An omniscient being could have full knowledge of the relevant music-theoretical space without having knowledge of the structures embodied in the *Well-Tempered Clavier*. Rather, Bach had to cognitively manipulate this space in order to achieve the results that he did, digging for certain melodic structures, "trying on" various combinations of notes, experimenting with pairings of such structures and combinations, and so on. Although he did all of this within the boundaries of a certain conceptual space, the creative results required—in addition to the relevant knowledge—non-truth bound cognitive states. (These issues will be discussed further in 5.1.)

Consider two final examples.

Example 5: One needn't peruse the literature on creativity and scientific discovery for long before coming upon an account of Kekulé's insight into the structure of the benzene molecule (discussed in 1.1.2). This example is popular for at least two reasons. One, Kekulé experienced his insight in a dreamy state while dozing by the fire: his thought process displays flash phenomenology. Two, Kekulé gave first person reports of the creative process involved in this discovery—this one and many others in fact (Findlay

1965; see also Boden 2004: 25-8, 62-71). For present purposes, we are only concerned with the following feature of Kekulé's discovery: his postulation that the benzene molecule was ring-like in structure was deemed impossible by the chemistry of his day. Nonetheless, the structure of benzene could not be adequately modeled given the conceptual framework available. Thus Kekulé's problem (P_2): a novel but scientifically acknowledged phenomenon and no adequate materials for explaining it. The solution? Kekulé broke the rules. He was constrained to some degree by the relevant domain (P_1), organic chemistry circa 1865, but a solution (P_3) required a method (P_4) that simply transcended that domain.[22]

Example 6: e.e. cummings is the kind of poet who elicits both fear and awe in aspiring writers. How could a poet be so boldly cavalier, yet maintain such a marked command of content, form, and feeling? One plausible, albeit oversimplified and incomplete, explanation is that cummings in fact retained many of the features of the tradition he inherited. cummings was rather fond of sonnets and other standard poetic forms. He often used standard techniques of rhyme, and many lines were peppered with alliteration: "Softer be they than slippered sleep/the lean lithe deer/the fleet flown deer." For all that, cummings rejected a great deal of poetic tradition, sometimes in ways simply unthinkable from within that tradition. He famously dropped the use of punctuation and capitalization, refusing to capitalize even his own name or the first person pronoun. More radically, he deliberately distorted syntax, using verbs as nouns and nouns as verbs, and splitting words so that they carry over into the next line. Finally, unlike anyone before him, cummings first realized the potential the typewriter had for making the visual form of a poem aesthetically relevant. It is thus that cummings, in exploring (P_2) the

[22] Boden gives a concise description of this set of circumstances (2004: 62-3).

poetic domain (P_1) maintained many features of that space, while *revealing* (P_3) altogether new ones.

These last two examples make salient a couple of important points. One, sometimes a domain constrains more than the knowledge had about or within that domain. Both Kekulé and cummings did something that the rest of the members in their fields did not *think* of—this, for the very simple reason that the respective conceptual spaces deemed such thoughts impossible: you can't have a ringed molecular structure; nouns work this way, verbs this way. So in exploring the conceptual space, chemists and poets naturally stayed within the boundaries given them. Somehow, Kekulé and cummings managed to transcend that space. However, and this is the second point, both Kekulé and cummings were to some degree constrained by the relevant domain. They did not re-invent the wheel, and so much of their method answered to the dictates of the theoretical context. So even radical innovations are not ones that emerge ex nihilo.

3.3 Conceptual space and creative processes

We can consider these six examples along two dimensions: mapping conceptual space and creative processes. We can ask, first, what kind of conceptual exploration and mapping the examples involve and, second, whether the processes are ones we want to call creative. What we find may be a bit surprising: these qualifications cross-cut one another. Transcending the conceptual map or even changing it are not necessary for creativity. Some instances of creativity follow the map more closely than we might have expected.

3.3.1 Dimension one: Maps and spaces

With respect to conceptual spaces and mapping, we can divide the examples into three types. Examples 1 and 2 are the most constrained and involve the least by way of conceptual change. Here the agents in question, as it were, use the map and stay on the map. Examples 3 and 4, though they involve significant appeal to the concepts of the relevant theoretical domain, involve some conceptual augmentation. Here the agents in question change the map but remain within the space.[23] Finally, examples 5 and 6, though they too involve significant appeal to some relevant theoretical domain, involve conceptual invention. Here the agents make a new map and at least some parts of this map are inconsistent with the standing map and the space originally mapped. Call these problem space situations *Type 1*, *Type 2*, and *Type 3* respectively. Let us briefly consider an example from each type.

*Type 1:*Norm needs a map and he is painfully aware of this fact. He doesn't know the first thing about eigenvectors, entanglement, wave functions, or spin states. (Although, sometimes after hours of crying in his beer over the lovely Eigen, he winds up entangled in waves of nausea and spinning vision. But presumably that doesn't count.) He is thus at the mercy of the information provided by authorities on quantum mechanics: whatever they say about their science, he is willing to believe. Norm is thus a map user of the most religious type, he will explore only those features of the space explicitly represented by the map.

---

[23] Perhaps this means we have a new map entirely. Perhaps it means it is the same map so long as the features changed were not essential ones. Questions about the identity of maps aside, the important point is that with these kinds of cases, agents stay within the conceptual space or domain, but re-map it in some way or other. So the concepts normally applicable to and dependent upon that space are somehow enhanced or augmented, but none of those changes are inconsistent with the *space* as it stood before the change.

Type 2: Bach, in exploring a space, augmented the existing map for that space. He used the map that he began with, but showed how there were features of the mapped space not represented on the original map. This is an especially instructive example. We see on the one hand the importance of domain knowledge: Bach had to learn and know a great deal about music theory and technology. So we see the important constraint that P_1 may place on P_3 and P_4 in undestined conceptual exploration. On the other hand, we see how that knowledge—itself solely derivative from and thus constrained by the pre-existing map of the relevant domain—is insufficient for such creative advances. Thus something besides knowledge of the pre-made map had to be part of Bach's heuristic method.

Type 3: "Mental geography is changeable, while terrestrial geography is not" (Boden 2004: 61). Examples 5 and 6 exemplify one interpretation of Boden's claim: conceptual spaces may be difficult to change but they, unlike the earthy stuff more standardly mapped, are malleable. The innovations of e.e. cummings make a strong and simple case in point. cummings, in exploring the techniques of his trade, did more than just map or re-map the space for his domain. He showed how the domain itself, the theoretical framework in which all the poets had so far worked, could be expanded. He made a new space and then set to mapping that space. Here it is all the more clear how knowledge of the relevant conceptual domain would not be sufficient for the results achieved. cummings could not have derived his results (P_3) from only current literary technique and contexts (P_1), since such results were outside the scope of that conceptual space. So here again we see that the methods for solution—in particular the cognitive

processes—must involve more than doxastic or truth-tracking states. The candidate states for occupying the role P_4 in such cases will be the focus of chapter 5.

3.3.2 Dimension two: Creativity

Of the examples considered, which involve cognitive processes and thoughts that we should call creative? Here we draw on the conceptual analysis of minimal creativity proposed in chapter 1. We work in reverse order, since the question becomes more and more difficult to answer for some of the earlier examples. Note also that we have no way of knowing the *actual* nature and details of the thought processes of Bach, cummings, Kekulé or the rest. We thus must stipulate certain features of our analyses. This should be no cause for worry. The task here is to ask of certain *types* of problem scenarios, whether it is reasonable—given the principles with which we are working—to attribute creativity to the cognition involved. It should not matter if our analyses turn out, as a matter of fact, to misconstrue the examples.[24]

Type 3: Recall that *MC* consists of three conditions: the agency, novelty, and modal conditions. Consider cummings' innovations, focusing just on his syntactic manipulations for simplicity. Were the cognitive processes and thoughts (call them *Cp*) that resulted in such innovations creative ones? Cp satisfies the agency condition:

---

[24] Consider the cummings case. Were his syntactic innovations truly p-novel or h-novel at some time or other? Who knows? cummings may have been harboring these thoughts and keeping them to himself for some time. Maybe cummings ripped them off his postman. Maybe someone else had them and simply lacked the courage to voice them. Such considerations remind us that we must distinguish the metaphysics from the epistemology of these situations. When doing the metaphysics, we recognize that for some person or cognitive process, we have limited access to the facts of the matter. So much of our job is a conceptual one: we ask of a certain set of circumstances whether or not it is reasonable to conceptualize them in this or that way, given certain principles that we establish independently. When doing the epistemology, we begin asking normative questions about our access to the facts of creative processes, answers to which typically issue in some degree of skepticism. Though important, we are (presently) less concerned with the latter set of questions, and more concerned with the type-level metaphysics. And we should be wary of shifts from one set of concerns to the other, which will undoubtedly confuse any discussion of creativity.

cummings is non-trivially responsible for his formal inventions and the thoughts that engendered them. He undoubtedly did not have in mind from the start the particular results that he ended with. But he did approach a particular problem space in particular ways, and would not have gotten such results otherwise. Cp is thus the product of agency. Our inquiry into the second and third conditions requires more speculation, as we obviously do not have access to cummings' epistemic perspective. But here is how things look from where we stand. Cp satisfies the novelty condition: cummings' thoughts were novel with respect to his own mind. This seems reasonable given that the thoughts were historically novel, and historical novelty plausibly entails psychological novelty. Cp satisfies the modal condition: the thoughts could not have been had by cummings before. Assuming that before the relevant time t_i—whatever that time was—cummings remained as it were within the mapped features of the conceptual space he inherited, he could not have had those thoughts prior to t_i. According to that map and that space, manipulating syntax in the ways cummings eventually did was impossible; Cp involved thoughts that were not available given the rules of the current game. Or, recall the discussion in 1.3.3 of the individuation of actual cognitive profiles by appeal to heuristic paths or causal strings. In this case, the heuristic method which finally resulted in his innovation was not, before t_i, available for cummings. The satisfaction of these three conditions is sufficient for creativity. Thus, as we surely expected, we call cummings' innovations creative.

Type 2: Examples 3 and 4 are different from 5 and 6 with respect to the use of conceptual space: Bach and Gettier mapped and re-mapped parts of their conceptual domains, but stayed in the space already mapped. Kekulé and cummings went one

further, moving out of the space itself. But this should not be cause for withholding attributions of creativity to the processes and thoughts involved in the former. Take example 3, the Gettier case. Call the relevant thoughts and processes *Gp*. Was Gp creative? Gp satisfies the agency condition: Gettier's counterexamples and the thought processes that led to them were ones for which he was non-trivially responsible. Had he not approached that conceptual space in the ways he did, he would not have gotten the results that he in fact did. Gp satisfies the novelty condition: these thoughts were novel with respect to Gettier's mind. Why should we think this? Once again, we are without access to these thoughts, so we cannot know for sure. But it is reasonable to infer p-novelty from h-novelty, and the Gettier problem certainly was novel with respect to the history of ideas. Finally, Gp satisfies the modal condition: Gettier could not have had the relevant thoughts prior to the time in question t_i. Put another way, Gp—the enabling conditions or heuristic path for the relevant thoughts—was psychologically impossible for Gettier prior to t_i. What evidences this? Given the theoretical space that Gettier was working within, he could not have had the relevant ideas without some thought that did *not* logically derive from the explicit portions of that space. This is for the simple reason that the novel ideas were not on the map. The occurrence of Gp thus required something special, some new way of thinking about the space, which before t_i was not possible for Gettier.[25]

To be sure, application of the modal condition to the cummings and Gettier examples suggests the following question. Just what was special about their cognitive

[25] Again, the actual truth of the modal assessment does not matter. The task here is simply to question attributions of creativity to problem scenarios that meet a certain functional description. We thus must ask, whether this case, as an instance of a more generic functional type, plausibly meets the conditions of the conceptual analysis settled on.

processes and how did it enable a previously impossible thought? We will try to answer it in chapter 5. For now, we should note the following feature of Type 2 problem spaces. We should see, after even this brief analysis, that significant contextual constraints upon a thought process do not preclude that process from bearing creative fruits. We should also see, and this is just the same point from another angle, that an agent need not leave the conceptual space in order to have a creative thought.[26] So the proposition that creative ideas are not exclusive to radically innovative ones reveals itself again. This revelation is further clarified in the light of some Type 1 examples.

Type 1: My navigation from the westside to the eastside of Vancouver is about as constrained by context as it gets. I look at the map, read the route directly from it, and follow it carefully. These thoughts and processes seem not to meet the modal condition for minimal creativity: it is plausible to say that I could have had these thoughts before. Put another way, since all of the knowledge acquired and used derives solely from the map, and the acquisition and use of such knowledge does not require the acquisition of any particularly challenging or new content, I could have had these thoughts before. We thus should not (and would not) call my thoughts and processes creative.

What about Norm? Norm, like my navigation of Vancouver, is largely just following instructions. He has been given a map of sorts, and is doing his best to understand the content of that map. On the face of it—and so our intuitions tell us—Norm is not thinking creatively. But we should do the job right, and check Norm's

---

[26] In some ways, Boden makes just this inference, since she claims that conceptual exploration alone cannot yield what she calls *impossibilist creativity*. This richer sense of P-creativity requires one to either transcend or break the rules contained within some conceptual space. But if we opt for *MC*, we leave room for non-transcendent, P-creative exploration of conceptual spaces, where this involves agent-responsible, P-novel thinking with respect to the domain that could not, for that particular agent, have been done before. And this better accommodates our pre-theoretic intuitions: we are inclined to call acts or thoughts like Bach's and Gettier's richly creative.

cognitive processes (*Np*) against *MC* anyway. Are Norm's thoughts about quantum mechanics (partly) a product of his agency? The answer seems an obvious 'yes': Norm is non-trivially responsible for forming the various beliefs that he does about the topic. He has in fact forced himself to endure paper after paper on the stuff. Np thus satisfies the agency condition. Is Np novel with respect to Norm's mind? Here the answer seems a resounding 'yes': Norm has never thought about the stuff that puts a twinkle in Eigen's eye before aspiring to put a twinkle there himself. Np satisfies the novelty condition. Finally, could Norm have had these thoughts before? There are good reasons to answer in the negative. As we've already noted, prior to his Eigencrush, Norm had never entertained any thoughts about quantum mechanics (other than perhaps its name and how hard it must be). This fact alone implies that Norm could not have had Np before. This is because of the kinds of thoughts and thought processes Np involves, namely, ones very specific—some of them in fact exclusive—to a theoretical domain. Here we note an important difference between examples 1 and 2. Unlike thoughts about basic road maps, thoughts about quantum mechanics have contents that are unavailable to someone occupying Norm's cognitive position. As a fairly ordinary North American, I am accustomed to basic mapping concepts and know how to use them. But the concepts of quantum physics are entirely foreign to a social theorist the likes of Norm. So prior to t_i —the time when Norm began pouring over the quantum theory literature in the name of love—Norm could not have had Np. Consequence: Np is creative.

But can this be right? Hasn't something gone wrong if thoughts like Norm's turn out to be creative? Grant that we have applied our concept, via *MC*, correctly. Accept that this bears the following consequence: Norm's thoughts are creative. This

consequence motivates the following dilemma. Either we jettison *MC* to avoid the apparently dubious consequence or we maintain *MC* and reconcile the consequence. The first horn is straightforward enough: one or more of our conditions is wrong, or we need another one, or some combination thereof. The second horn apparently requires concept revision and an error theory. Since our pre-theoretic intuitions say that we do not call Np creative, we have a bullet to bite. We then must say why it is easy to swallow. I embrace the second horn: we have given principled reasons for settling on *MC* and should attempt to maintain it. What reasons are there for accepting instances like Np as creative?

3.3.3 'Creativity': Minimal and mundane

Jesse Prinz and Lawrence Barsalou defend a perceptual theory of concepts: concepts are symbol structures—representative of objects, properties, relations, events and the like—constructed out of more primitive perceptual symbols (Prinz and Barsalou 2002; see also 1997; Barsalou 1999). The latter are extracted from perceptual experience. "[P]erceptual systems must abstract away the irrelevant details of incoming signals, separate aspects of a signal that are independent, and introduce a significant amount of structure" (Prinz and Barsalou 2002: 109). Concepts are then built out of the symbols either as a natural consequence of perceptual experience, such as in biological taxonimization, or as a result of certain classificatory conditions or goals had by the agent, as in searching for something fit to serve as a door stop. Perceptual symbols are not accessible to consciousness, and thus the process of concept acquisition is not one entirely available to consciousness. We do not deliberately construct our concepts;

rather, they emerge given certain perceptual experiences, attention, memories, goals, and conditions.

Prinz and Barsalou claim that this process of concept acquisition is *mundanely creative*. Here they appeal to a distinction closely analogous to Boden's (as discussed in 1.2): they distinguish *mundane creativity* from *exceptional creativity* (2002: 106-7). Mundane creativity only requires a mental representation that is novel to the representer. By this criterion, concept acquisition clearly qualifies. Now, the perceptual concept model is controversial. The claim that concept acquisition—as situated in the perceptual symbol model—is creative is controversial. We need commit to neither thesis. Prinz and Barsalou make a convincing case for perceptual concepts, and for attributions of mundane creativity. But there are reasons to opt for a non-perceptual model of concepts, and there are reasons (given our *MC*) to deny creativity to the acquisition of concepts.[27] In any case, Prinz and Barsalou's view is instructive, compelling us to acknowledge the following simple point. Some of our most mundane cognitive processes display marks of psychological creativity. Most basically, the view says 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 for the mind of that agent. Granting that we form concepts (either in precisely the way the perceptual symbol theory proposes, or in the way a more traditional amodal concept theory

[27] Here is one: the mental phenomena in question may not be the product of agency. These are not actions that the agents perform—and the authors emphasize this—and so perhaps are not the kinds of events that we want to call creative. Concept acquisition results from certain attentional actions by the agent but the agent makes no effort to form concepts and has no conscious access to the parts that compose the concepts. By the same token, there is obvious wiggle room here: we might say that the perceptions, perceptual circumstances, attentional focus, beliefs, and other cognitive states are ones for which the agents are largely responsible, and this suffices for the agency condition. As we see above, we needn't make this decision. We need only recognize that there are compelling reasons for either assessment, and so mundane processes like concept acquisition cannot be dismissed as uncreative out of hand.

proposes), we are obligated to say that doing so is a psychologically novel act. This point generalizes to a variety of cognitive states. Belief formation, skill acquisition, and a variety of other kinds of learning and thought require the same novelty of their agents.

According to *MC*, psychological novelty is not sufficient for minimal creativity. Whether such processes qualify as creative depends upon the satisfaction of the agency and modal conditions. There is no way to make a clean demarcation here, but many cases of concept, belief, and knowledge acquisition naturally fit the bill. We have already seen that a mundane example like learning quantum mechanics involves thought processes that result from agency and could not have occurred before, given certain cognitive circumstances. The troubling consequence was that cognitive processes like Norm's thus qualify as creative. But if these cognitive achievements require of their agents anything like what Prinz and Barsalou propose, then this consequence should not appear troubling and *MC* should be vindicated.

Conceptual revision of pre-theoretic intutions requires justification. Our intuitions tell us that thoughts like Norm's are not creative, and *MC* tells us that we are wrong and should revise our concepts. Here we need to remind ourselves of the force of Boden's distinction between psychological and historical creativity. While admitting that only those thoughts that are novel with respect to the history of ideas are creative in the richest of senses, we can maintain a sense of psychological creativity for more mundane thoughts. Any surprise felt with respect to attributions of creativity to mundane instances derives from a failure to keep a finger on this distinction. We justifiably withhold 'creativity' from instances like Norm's only when richer historical creativity is the sense in mind. If we keep minimal psychological creativity in mind, where the relevant novelty

is relativized to the psychological agent in question, we should have no reservations about making the attribution. So an error theory is not really requisite here, since there is not much by way of conceptual revision going on. Rather, distinguishing p-creativity from h-creativity and focusing on the latter requires us not to use our concept differently but to disambiguate familiar senses of that concept and then focus only on one of those senses.

If this analysis is correct, we have learned that with respect to conceptual spaces and maps, not only is transcending the conceptual space (as in Type 3 examples) not necessary, but changing the conceptual map (as in Type 2 examples) is unnecessary. There are mundane types of cognitive processes that deserve the name creativity.[28]

3.4 Conclusions

A functional analysis of problem solving and creativity has proven a powerful exercise. We have identified the stuctural similiarities between problem solving scenarios and creative thinking scenarios. This partly vindicates Weisberg's claims, discussed in 2.3, regarding creativity and problem solving without endorsing his reductionism. Identification of the functional roles of problem solving elements reveals how those elements constrain one another, and this carries over to creative thinking. Importantly, we see that these constraints, even when rather rigid, do not prevent a problem space from engendering creative products: one need not transcend the relevant

[28] It should be apparent how much of the above owes to Boden's discussions of conceptual space and mapping. These metaphors are extremely powerful ones and to be sure, much of the present analysis derives from her use of them. However, there are many ways that the analyses are distinct. Here are two important ones. Boden distinguishes Type 3 cases from Types 1 and 2 with respect to radical or impossibilist creativity, where only the first qualify. She also distinguishes Type 1 cases from Types 2 and 3 with respect to creativity *simpliciter*. So only rule breaking or conceptual space transcendence involves impossibilist creativity, and rule-following processes simply are not creative. As our analysis has hopefully shown, there are reasons for doubting these clean distinctions.

conceptual space or map to think creatively. Mundane instances of psychological creativity abound.

The analysis also accommodates the points by a number of theorists, discussed in 2.2, regarding knowledge, memory, learning, metacognition or reflection, and incubation. All of these cognitive capacities are relevant to solving problems and to thinking creatively. Depending upon the occupants of P_1 and P_2, each of these capacities may be more or less important as heuristic methods (or components thereof) for completing the task: they potentially occupy functional role P_4. However, we have suggested that the heuristic methods required for creative problem solving must include more than truth-bound cognitive processes. The hunch is that imaginative processes are likely candidates to fit the bill. Chapter 4 thus provides a conceptualization of various modes and dimensions of imagination. Chapter 5, after consideration of some alternative candidates, turns to the role of imagination in creative thought.

Chapter 4
Modes of Imagination

The chapter following this one attempts an identification of the role of imagination in creativity. However, before anything interesting can be said about the connections between creativity and imagination, we have to do a bit of conceptual clarification. The aim of this chapter is to give a conceptualization of 'imagination.' There has been a wealth of recent literature on the topic, and so the exposition here will be relatively brief. Additional clarifications will of course be offered in the chapters that follow.

4.1 Distinctions

4.1.1 Propositional Imagination

One of the most appealing and intuitive ways to think of imagining is as a propositional attitude like belief and desire. To imagine Winston Churchill in the next room is to imagine that [Winston Churchill is in the next room]. It is to take an attitude, distinct from belief, with regard to some proposition *P*. As Gregory Currie and Ian Ravenscroft put it, imaginative states are *recreative* states, similar in character and content to certain *counterpart* states which they model or recreate (Currie and Ravenscroft 2002).

The most obvious and commonly discussed candidate for such an attitude is belief-like imagining—also dubbed 'make-belief,' 'belief mode of imagining,' and 'simulated belief,' among others—which takes belief as its *counterpart*. Kendall Walton is most responsible for championing the make-belief model of imagination (Walton

1990). Walton proposes that imagination be understood in terms of the kinds of make-believe games that children play. If a child pretends that a stump is a bear, the stump serves as a prop in her game of make-believe. Props function, according to Walton, to generate fictional truths: the truths *to be imagined* for that particular game of make-believe. Thus truths about the stump and its surroundings will generate fictional truths about the bear—that it is in a patch of poison ivy, that it is on a hill, that it is dark brown. Walton's thought is that we engage with representational artworks (and representations more generally) in the same ways. When we view a painting or watch a film, the work will serve as a prop in our game of make-believe and thus generate fictional truths—propositions to be make-believed—for that particular game. One important feature of the make-belief picture is that imagination is characterized not by type of content, since we can equally make-believe actual, non-actual, or fictional content, but (in part) by a certain kind of attitude with regard to that content.

Make-beliefs are similar to beliefs in terms of their inferential connections: a make-belief that 'If *P*, then *Q*' and a make-belief that '*P*' licenses a make-belief that '*Q*'—as does believing the same propositions. Make-beliefs, at least on some accounts, can guide and cause action.[29] Make-beliefs are also like beliefs in causing other imaginings, beliefs, desires, and emotions; they thus possess mental casual efficacy in virtue of various cognitive connections. (Inferential and cognitive connections are discussed in greater depth in 5.2.2.) Fiction writers, film makers, and other narrative artists ask us to employ, at least, this form of imagining. They ask us to make-believe the

[29] This is a contentious issue, separating Humean theories of motivation which take belief and desire to be the only causes/rationalizers of action, from non-Humeans who allow for other states to cause/rationlize action. See especially: Currie 2002b; Currie and Ravenscroft 2002; Funkhouser and Spaulding (in preparation); Nichols and Stich 2000, 2003; Velleman 2000.

existence of fictitious events, persons, things, worlds. When we comply with such requests, our thoughts lack any intrinsic commitment to truth or action: make-beliefs possess no epistemic commitment qua make-beliefs. (This feature of imagination is discussed at greater length in 5.2.1.) We may thus (safely) richly engage with the fictional worlds provided by films, novels, and other representational artworks. It is this potential richness and lack of epistemic commitment that sometimes makes our engagements with fictions so engrossing. Acknowledging, with Walton, the similarity between this cognitive activity and children's pretence is fruitful. At their own direction, as opposed to the artist's, a child may make-believe counteractual events, persons, things, and worlds. These features of make-belief in particular and imagination in general prove important in connecting imagination with creativity (see 5.2).

Cognitive states of this sort are typically termed *propositional attitudes*: they involve a relation between an agent and a proposition. They might be termed as such for one of two reasons. States like make-beliefs might be called propositional attitudes for their possession of propositional *content*. Alternatively, they might qualify as propositional attitudes for involving a certain *character* similar to their non-imaginative counterparts. That is, they might involve, like belief and desire, some kind of attitude. Distinguishing content from character helps to clarify the conceptual space.

4.1.2 Content and character: Non-imaginative mental states

Tokens of cognitive states like belief are, at least in part, individuated by their *content*. My belief that [It is raining in Vancouver] is distinct from my belief that [It is snowing in Chicago] is distinct from my belief that [It is sunny in Miami]. The

intentional content of each of these states—what they are beliefs *about*—does the work of distinction. All three are of the same type of state, beliefs, but their distinct content makes them distinct belief tokens.[30] I might also have a desire with the same content, say, a desire that [it is raining in Vancouver]. This implies two related facts. First, content is not sufficient to distinguish mental states at the type level. That is, beliefs, desires, intentions all have the same kind of content, so something else must demarcate them as different *types* of propositional states (i.e. as of the belief-type, the desire-type, etc.). Second, it is not sufficient to distinguish mental states at the token level. If A has two simultaneous occurrent states, one a desire that [It is raining in Vancouver] and one a belief with the same content, what makes them different tokens? The answer is that what does the extra work of distinction at the token level is the same thing that does the work at the type level: A's simultaneous belief and desire with the same content are tokens *of* different mental state types. But the question remains: what makes them different types, and thus A's two states different tokens?

Currie and Ravenscroft distinguish a mental state's *character* from its content.[31] Their emphasis is on forms of imagination, but the basic point generalizes.

> It will be important to distinguish between the content of a piece of imagining—what it is you are imagining—and the character of that imagining. Imaginings with a visual character need not have vision as part of their contents: their contents need not be such that specification of that content requires deployment of the concept *vision* (Currie and Ravenscroft 2002: 12).

---

[30] To qualify: their distinct content *and* their being tokened by a particular believer at a particular time. You and I could have the same type of attitude with the same content: we both believe that it is raining in Vancouver. Our states would thus be the same at the type level. But beliefs, like other propositional attitudes, are relational, and your being the relatum in one case and my being the relatum in the other renders the two states distinct tokens. For related issues, see Kenneth Taylor's discussion of what he calls attitude 'character' (Taylor 1998: 282-290).

[31] And note that Currie and Ravenscroft's use of 'character' is distinct both from the use made by David Kaplan (1989), as well as Kenneth Taylor (1998) (see fn. 30).

> If I have a visual image of a mountain, then the content of my imagining (assuming that imagery is a form of imagination) is the mountain or, if we want to make all contents propositional, there being a mountain somewhere. The concept of perception need play no part in specification of the content; it must, though, play a part in specifying the character (Currie and Ravenscroft 2002: 27).

This tells us what character is not: it is not part of mental content, nor is it part of the individuation of that content. But it fails to tell us what character *is*. There seem to be two, non-exclusive but distinct, options here: mental state types are characterized by distinctive phenomenologies and/or distinctive functional roles. The discussion of Currie and Ravenscroft is ambiguous on this point. They sometimes talk of phenomenology and, other times, of functional role.

The phenomenology of some mental states is easy enough to describe. There is something it is like to have a (human) visual experience. There is something it is like to have an auditory experience. Thus we describe a visual experience in terms very different from an auditory experience, even if the two experiences are experiences of the same thing. The properties that we experience and the concepts we use to characterize them will be distinctively visual and auditory, respectively.[32] And the same goes for the other perceptual modalities.

This is true of some non-perceptual states as well. Fear, anger, and sadness, among others, feel a particular way. For example, I can feel my pulse quicken and my body tense when I am afraid.[33] What about propositional attitudes? Do they possess *distinctive* phenomenologies? For some attitudes, the answer seems more straightforward

---

[32] Or at least, for any given modality, some of them. There are amodal perceptual concepts not distinctive of just one modality or another, e.g. direction or surface texture.

[33] Whether these feels are distinctive of particular emotions is debatable. For example, the physiological feelings or symptoms that accompany, say, fear and anxiety are very similar: both will issue in tensed muscles and increased heart rate. This has led some to posit cognitivist theories of emotion: what distinguishes fear from anxiety, as emotion types, are the kinds of cognitive judgments and evaluations that enable or cause the emotional states (see, among others, de Sousa 1987; Oatley and Johnson-Laird 1987; Oatley 1992; for alternatives to the cognitivist accounts, see Damasio 1999; Prinz 2004; Robinson 2005).

than for others, especially for those that are called emotions as reasonably as they are called attitudes. For example, it is intuitive to say that hope feels a certain way, remorse another way, and disappointment another.

The case is trickier for beliefs and desires. We might take belief to involve a feeling of conviction or certainty with respect to the proposition believed. This feeling varies in degree: I am more and less certain about varying propositions (Cohen 1989, 1991). Desires might be characterized by a feeling of drive or pull towards their objects. I feel pulled towards the beer in the fridge or to have another cookie. I feel driven to the woman in the red dress or the sleek motorcycle. This feeling also varies in degree: lusting over the woman in the red dress *feels* different from a long-standing wish for a motorcycle, or a mild desire for a second cookie. Perhaps women in red dresses have my number, while the motorcycle is a fading fancy and I can generally pass on extra sweets.

These points are contentious. Many philosophers of mind take attitudes like belief and desire to be lacking phenomenologies altogether, let alone distinctive ones. There is nothing it is like to be in a state of belief, desire or intention. Qualia, if we posit them at all, are reserved for perceptual experience, pain, and perhaps emotions, but not for propositional attitudes (Dennett 1988; Kim 2005). Others argue that propositional states, just like perceptual states, have distinctive phenomenologies which are in fact inseparable from the intentionality of such states—that is, such mental states have distinctive phenomenologies both in virtue of their being attitudes of one kind or another and in virtue of their having the content that they do (see especially Horgan and Tienson 2002, and Flanagan 1992; see also Jackendoff 1987; Loar 2003a, 2003b; Strawson

1994).[34] According to this view, a desire for a ripe tomato will have a distinctive quale (or qualia) no less than a visual perception of a ripe tomato. So although using phenomenology to individuate mental state types is controversial—for the simple reason that mental state phenomenology is controversial—we should at least grant that the issue(s) remains open.

We might instead distinguish mental state types like belief and desire by their functional role. Beliefs carry various consequences for action, as well as for other mental states like desire, intention, and various emotions. Beliefs dispose us to act in certain ways: if situated in the appropriate context, I will assent to the truth of a proposition believed. Or, if coupled with the appropriate desires, I may perform actions in accordance with my beliefs. These are observable symptoms of the functional role of belief. In Robert Nozick's apt phrase, beliefs track truth (Nozick 1981). It is the function of belief, qua belief, to do so.[35] This fact about belief underwrites how beliefs stand in relation to other elements in the larger cognitive-behavioral system. Our beliefs *guide* our inferences and our actions: I (at least tacitly) take my belief that *P* to be true (else it is not a belief I have) and this fact informs what I infer about or from *P*, how I might act on *P* or consequences of *P*, how I might try to get or prevent *P*, and so on. Moreover, this is a function unique to beliefs and, so we might think, provides a stable mark for distinguishing belief-type states from other types of state.

---

[34] In fact, C.S. Peirce, to whom the singular term 'quale' is often credited, explicitly used the term broadly, to denote the various ways things seem—their distinctive phenomenology—from a subjective point of view. "The *quale*-consciousness is not confined to simple sensations. There is a peculiar *quale* to *purple*, though it be only a mixture of red and blue. There is a distinctive *quale* to every combination of sensations—There is a distinctive *quale* to every work of art—a distinctive *quale* to this moment as it is to me—a peculiar *quale* to my whole personal consciousness. I appeal to your introspection to bear me out" (Peirce 1935 [1898]).

[35] Again, 'function' denotes a Cummins-function: dispositions or systematic relations that an element bears with regard to other elements in some larger system (see ch. 3, fn. 16).

Desires do not function to track truth. They do, however, serve a motivational function. And this fact determines the nature of desire's relation with other elements in the cognitive-behavioral system. Desires, when coupled with beliefs (or perhaps with other states: it depends upon whether or not one opts for a Humean theory of action) will motivate us to act. They may also cause or affect other mental states, like belief, intention, and emotion. Desire too, then, occupies a unique functional role in cognitive systems. This provides a mark for distinguishing desire-type states from other types.

It may be tough going, but it is plausible that distinctive functional roles for the other propositional attitudes—intention, hope, regret, and so on—are identifiable. Different types of attitudes function differently as elements in a cognitive system and it is this difference that (at least in part) makes them different types. A belief is a type of mental state different from a desire in virtue of its functioning differently. And the same can be said for all manner of propositional attitudes.

The point is not specific to propositional attitudes. Perceptual states also serve particular functional roles. Different perceptual modalities provide us with information about properties of our environments. (Of course there is some overlap here: vision and touch may both inform us of an object's texture, olfaction and taste, of an object's sweetness. But we don't hear colors; we don't smell weight; we don't see sonic properties; and so on.) They will also connect with behavior in unique ways. For example, the visual system acquires information that guides motor performance.[36] Suffice it to say—and this will come as no surprise to most—that perceptual states

[36] In fact, so much so that some have modeled a motor-guiding visual system distinct from the descriptive visual system. For philosophical discussion, see Matthen (2005) and Carruthers (in preparation).

function in distinctive ways. And (part of) what distinguishes them as states of one type or another is their respective function(s).

The functions of mental states thus provide another way of individuating those states as different types. The functional role of belief, for example, requires us to characterize belief differently than other states or, more strongly, it might *just be* the character of belief. And so on for other states. As we stated above, it is unclear if Currie and Ravenscroft intend something like this in using 'character.' They also sometimes speak of the character of a state in phenomenological terms. We saw above that for some states—perceptual experience and emotion—this is no problem. For others, it is more contentious. So we have a decision point here. We can understand character as (1) state function or as (2) state phenomenology.

Perhaps these do not exhaust our options. There is good reason to think, as our concepts imply, that (1) and (2) are not exclusive. Our auditory concepts, for example, are both phenomenological and functional. We speak of qualities of "tone", "timbre," and "volume". And it is the function of auditory experience to provide us with information as to these and other qualitative properties. The same is, at least to some degree, true of our attitude concepts. If I taunt you by saying "what does it feel like to want?", I am implying to you that desiring x is as close to x as you are going to get. You aren't going to get x so do the most that you can: relish in what it feels like to want it. This pop phrase is not arbitrary: included in our concept of desire are phenomenological features of wanting something that you can't have, or that you have had but that is now gone, or that you are very close to getting but haven't yet gotten. These features of the concept of desire are instructive, since they isolate what it is like merely to desire,

without the world yet matching one's desire. They indicate that part of our concept of desire involves what it feels like to be in such a state. Making the conceptual case for the function of desire is easy work: much of traditional action theory requires that desire function to connect with belief and behavior in particular ways.

All of this opens the door for two options in addition to (1) and (2). First, we can identify mental state character with (3) *both* the phenomenology and functional role specific to the state in question. There is nothing *prima facie* inconsistent with mental states possessing both phenomenologies and functional roles. Quick consideration of perceptual states shows this. And if one is willing to accept the phenomenology of attitudes, the same option is available for states like belief and desire. The thought here is that both features are constitutive of our mental state concepts and thus both do the work of demarcating mental state types. We might go one further: (4) all mental states possess distinctive phenomenologies and functional roles and the second is true (in part) *because of* the first. Mental states function the way they do in virtue of their subjective feels. Now the naturalist might rest content if this were just about content, that is, if this were just a point about intentionality and/or the objects of perception.[37] But it is not. Of course it is smell of *that* bakery that disposes me to cross the street for something sweet and gooey. But it is also *that smell*, that particular quality of experience, that carries such consequences. The phenomenology of the perceptual experience thus contributes to the way that it functions in the cognitive-behavioral system. We can say the same for belief. What else are we doing when responding to the question, "Are you sure?", but introspecting for our degree of conviction? We are checking to see how sure we feel.

---

[37] The 'and/or' accommodates the fact that, with respect to perceptual experience, some identify intentionality with objects of perception and some do not.

Sometimes we come up with certainty, sometimes we do not, and the two feel very different. We will act accordingly: I will tell you " I am absolutely certain" or, with less commitment, "I'm pretty sure, but I wouldn't bet the farm." It is, at least in part, the phenomenology of my state that disposes me to report in one way or the other. So (4) goes one further: it claims, like (3), that phenomenology and function work to individuate mental state types, and then offers a story about the connection between the two.

Many philosophers, especially those of a naturalistic bent, will remain suspicious of (2)-(4). For them, only (1) is viable for understanding a mental state's character as contrasted with its content. In any case, this is not a choice we are forced to make. Through the above analysis we have identified the available move space: there are a variety of ways of understanding character, and all of them serve to distinguish mental states at the type level. We have also revealed that our basic concepts of belief, desire, vision, audition, etc., encourage something like choice (3) if not choice (4). It is a conceptual fact that we distinguish mental state types both in terms of phenomenology and function. These may, of course, be unsophisticated folk concepts in need of revision or elimination. Maybe not. Either way, they are standard and so the burden of proof (and the obligation to an error theory) lies with those who insist on option (1).

At the very least then, both phenomenology and function should be kept in mind when appealing to the character of a mental state. And both potentially do the work of individuating mental state types, and thus mental state tokens with the same content (when simultaneously tokened by the same agent); both could individuate my belief that [It is raining in Vancouver] from my desire that [It is raining in Vancouver]. Nonetheless, the two features should be kept distinct: phenomenology is not function,

however the two should turn out to be related. Our use of 'character' thus must be non-ambiguous between the two; Currie and Ravenscroft fail to identify such a use. All of this should inform considerations of the nature of imaginative states.

4.1.3 Content, character, richness: (Standard) propositional imaginative states

Some imaginings, make-beliefs, take belief as their counterpart. They will thus share some, though not all, of the characteristics of belief. We argued above that the character of belief could be identified with functional features of the state, or phenomenological ones, or both. Our brief discussion in 4.1.1 of the inferential connections of make-beliefs with one another and with other states, suffice to show one similarity in function between make-belief and belief. Make-beliefs can be used in reasoning alongside other make-beliefs, or alongside beliefs (and the conclusions will of course be "tagged" accordingly). They also connect with action in similar, though weakened ways. My make-belief that the banana is a telephone will dispose me to use the banana in particular ways, at least for the duration of my imaginative project. Finally, make-beliefs can be cognitively causally efficacious like belief. A make-belief that there is danger down the dark alley or behind the shower curtain can cause fear, or worse, a belief with the same content. (More discussion of the cognitive and inferential connections of imagination generally are offered in 5.2.2). Make-beliefs are thus, in terms of function, similar in character to beliefs.

Other imaginative states take desire as their counterpart. In a way analogous to make-beliefs and beliefs, make-desires have cognitive and behavioral effects similar to desires. They may contribute to pretend decision: a make-desire to change careers and a

make-belief that the Billionaire Club is accepting applications encourages me to pretend that I will go ahead and apply (this feature is especially clear in the context of mental simulation and folk psychology; see below). Make-desires may cause (pretend) action: a make-desire to eat mud pies may cause me to motion as if to put the clumps of mud in my mouth (or if I am an inexperienced pretender, to stuff my face with mud).[38] They may cause other imaginative mental states, non-imaginative states, and emotional states. My make-desire to move to Hawaii might cause a make-belief that I just won the lottery; my make-desire for my neighbor's wife might get me into real trouble; my make-desire that Anakin Skywalker fight off the dark side will result in disappointment and pity. (More on function of belief and desire-like modes of imagination in the discussion of simulation theory below.)

How do these imaginative states, and others, *feel*? Do they have characteristic phenomenologies and do these characters match the ones of their non-imaginative counterparts? Put another way, is such a match what motivates (in part) our calling some states 'make-beliefs' and others 'make-desires'? It was argued above that the candidate phenomenologies for belief and desire are, respectively, a feeling of conviction and a feeling of drive. We also argued that these feelings can vary in degree, and that our introspective identifications of such variations further evidence belief and desire phenomenology. If the imaginative counterparts of these states have phenomenologies, they should be similar: a make-belief should involve some (imagined) conviction and a

[38] Two points: first, 'pretend' is parenthesized here because what renders an action pretend or not is unclear. It could just be context that does the work: an action is pretend if situated in the context of make-believe. It could be mental causes: an action is pretend if motivated and/or guided (in part) by imaginative mental states. It could be relative to the describing agent: if I don't know your mental states, your turning your back on a stump is just your turning your back on a stump (which is an action if anything is); but from your perspective turning your back on a stump is turning your back on a bear (which, at that level of description, is a pretend action). Second, whether imaginative states can cause actions unmediated by beliefs and desires is a contentious issue (see fn. 29 above).

make-desire should involve some (imagined) drive. It is hard to say definitively that this is the case, but here are a few considerations.

If part of our concept of belief and desire is that they feel certain ways, then when trying (either at our own discretion or by invitation, say from a fiction) to make as if to believe *P* or make as if to desire *Q*, we will attempt to imagine ourselves in the corresponding states of belief and desire. We will attempt to imagine the feel of really wanting and being drawn towards *Q*, or of being absolutely certain of *P*. This encourages us to think of imaginative states as varying in *richness*. Richer imaginings are ones that involve not merely supposing *P*, but imagining that one stands in a particular relation with *P*, like belief or desire, where that relation involves a distinctive phenomenology. Consideration of the machinery of simulation theory will help to bring this out.

According to (some) simulation theorists, we employ our belief-desire system *offline* in order to put ourselves in the mental perspectives of others, in order to model their beliefs and desires (Currie 1995, 1996, 1997; Currie and Ravenscroft 2000; Goldman 1989, 1992; Gordon 1986; Heal 1995, 1996a, 1996b; see also Davies and Stone 1995a, 1995b). This employment is offline in the sense that we do not actually form the relevant beliefs and desires, and thus do not actually act on them (at least not immediately). Instead, we simulate them: we make as if to possess certain mental states given certain environmental situations. It is in this way that simulation, according to some simulation theorists at least, is a kind of imagination: simulated beliefs are imagined beliefs. Note however, that simulation theory as such outruns a mere make-belief theory, positing not only simulated beliefs, but simulated desires, and perhaps other simulated attitudes.

Currie makes a good case that imaginative projects often involve desire-like imaginings in addition to belief-like imaginings (Currie 2002a; Currie and Ravenscroft 2002). Consider the simulation account in relation to folk psychology. Simulation theorists explain tasks of mindreading—tasks of explaining or predicting another person's actions and/or mental states—in the following way. If I wonder why Smith is performing some action or other, or wish to predict how he will act, I might simulate being in his environmental and social circumstances and having his mental states. Suppose that Smith is rifling bricks at Jones's house. Suppose Smith believes that Jones has been stealing from Smith's family, and Smith desires some kind of revenge. And suppose I somehow know all of this about Smith. By simulating his belief, his desires, and being in Smith's particular circumstances, I reason (again offline or hypothetically) that I (Smith) should act accordingly. (Maybe this particular method of revenge wouldn't be the obvious choice for me, but perhaps there are some other facts about Smith that I build into my simulation to get me this far.) I thus have an explanation for Smith's action. This explanation was enabled or facilitated, per the simulation theorist, by imagining being Smith in Smith's position. The important point here is that in order to construct this explanation, I had to simulate not only Smith's beliefs, but his desires and perhaps other states.[39]

It is reasonable to think that my successful explanation of Smith's action is partly enabled by really getting myself, as it were, into the mood of the simulation. Merely supposing that I have a certain desire for revenge is unlikely to get me anywhere near a

[39] Note that one only need to grant here that we *could* mindread in the described way, not that we *must* or even *do*. That is, the simulation account is one possible (and perhaps plausible) account of mindreading ability and as construed provides some motivation for positing imaginings of varied character and richness. But we don't need to endorse simulation theory as such to learn this lesson.

hypothetical decision to begin launching bricks at someone's house. I imagine what it must feel like to believe with absolute conviction that someone has wronged you terribly, what it must feel like to bitterly desire revenge and be driven to extremes to get it. The targeted mental states will not, of course, always be so extreme. But that is just the point: we simulate states that vary in phenomenological quality.

This implies that propositional imaginative attitudes are partly characterized by distinctive phenomenologies (which is consistent with possible notions of character (2), (3) and (4) from 4.1.2). Or this is true at least some of the time, namely, when they achieve some significant degree of richness. Imaginings are *rich* when they involve more than what we might call *bare make-belief* or mere supposition. We may suppose, much like we would for a task of counterfactual reasoning, that such-and-such is the case. We can, alternatively, form imaginings that are rich in quality or quantity.

Imaginative projects may be more or less rich in terms of the *quality* of the particular imaginings. I might just suppose that Smith seeks revenge, or I might make-desire to seek revenge in a more engaged manner. The difference results in differences in the (pretend and actual) decisions made, and thus affects how the imaginative states function in the larger system. Even if one rejects these points about phenomenology, we can still identify differences in richness vis-à-vis functional role. States that answer to 'make-belief' and 'make-desire' and, as will be discussed in 4.2.1, 'visual image', 'auditory image' and the like, possess a kind of cognitive richness that consists in their connections with other elements in the cognitive system. A make-belief that *P*, unlike a mere *entertainment* that *P* and to a degree greater than mere *supposition* that *P*, will stand in a complex set of relations with other cognitive states and with behavior (more on these

differences in 5.2.2). As we have argued, these relations will resemble the functional roles of their counterpart states.

We can make a similar point about the quantitative richness of imagination. If we are to imagine *P*, we may imagine not only the proposition in question, but the surrounding circumstances, consequent states of affairs, and so on. In this way, an imaginative project is rich in *quantity*: how much content we imagine. Compare: if asked to imagine that pigs can fly, I might baldly imagine that the proposition 'Pigs can fly' is true. Alternatively, as is usually the case when engaging with fictions for example, I might visually image pigs flying, pawing their hooves through the clouds, squealing in airborne delight, snorting at the windows of skyscrapers; I might imagine havoc near airports, Iowa farmers in fits of confusion and outrage, a new kind of "game bird" for our Thanksgiving dinners.

The points about richness in quality, then, are about the character of imagination. They allow for a number of options: we can understand character of imaginative states in terms only of function (option (1) from 4.1.2), only of phenomenology (option (2)), or in terms of both (option (3) or (4)). The points about richness in quantity are about the content of imagination.

One final clarification that emerges from our distinction between character and content. Imaginative states can be at odds with our actual states. This is evidently clear in the context of simulation: I can maintain my simulated belief that Jones stole from Smith *and* my actual belief that Green, not Jones, is the culprit. Considerations of fiction further illustrate this point. Walton and Currie are both careful to argue that our imaginative attitudes with regard to a fiction can oppose what we actually believe, desire,

and so on. I might generally dislike happy endings, desiring that fictions end unhappily, or at least, in unorthodox ways. However, when engaged with a fictional narrative—when playing the relevant game of make-believe in Walton's terms—I make-desire that the events turn out well, that the protagonist is saved and the antagonist gets his due. My actual desire is thus *backgrounded*, and I am able to form imaginative states with the contrary *content*. I can also form imaginative states whose *character* differs from my counterpart states with the same content. As an agnostic, I might believe with slight conviction that God exists. But, perhaps because I want to see if Pascal's trick really works, I can make-believe the same proposition with complete conviction. This make-belief is phenomenologically different from my actual belief with the same content. (Note also that it is phenomenologically richer than a supposition with the same content.) My actual desire for beer disposes me to take a trip to the pub, but my make-desire for beer doesn't get me that far (not, at least, by itself). My make-desire thus functions differently than a desire with the same content. (More on the *cognitive quarantine* and *inferential isolation* of imaginative states in 5.2.2.).

These facts about backgrounding should be no surprise to anyone familiar with the pleasures of novels and films: much of this enjoyment derives from the opportunity for trying on different perspectives, attitudes, worldviews, for seeing what it is like to be someone else, if only for a short while. It is, nonetheless, an important feature of imaginative capacities and the discussion of character and content, and of simulation theory, helps to bring this out.

4.2 Perceptual imagery

Philosophical discussion of imagination divides on the following point. Going back at least as far as Descartes, imagination is understood in terms of mental imagery: to imagine Winston Churchill is to image or mentally picture him, where the content of this state involves certain perceivable properties had by Winston: his wearing a charcoal grey, three-piece suit, his smelling of cognac and cigars, his passionately discussing current plans for the British government. As just discussed, imagination has more recently been understood as a propositional attitude: to imagine Winston Churchill is to imagine, for example, *that* Winston is in the next room. Some think the second conceptualization subsumes the first. Others think that imagery is importantly distinct from propositional imagination and thus deserves distinctive conceptualization.

The *imagery debate* in cognitive science, now nearly 30 years old, roughly mirrors this divide. We can understand the core of the debate as follows. Some think that there are features of imagery that go unaccommodated if imagery is construed as a propositional state. Others think that the propositional construal will do just fine: all imagining, imagery and otherwise, is propositional. This implies a dilemma: if imagery is propositional then perhaps we can do with just one kind of mental representation. But this may imply the consequence of poorly accommodating the phenomenological features and functional role(s) of imagery. If imagery is not propositional, then we seem forced into positing two kinds of mental representation, with potential loss of theoretical parsimony and unity. Theorists involved in the *imagery debate* embrace one horn or the other. We of course do not have to solve this debate, but considering some of the

relevant conceptual space and the occupying camps serves to clarify imagery and, more generally, imagination.

4.2.1 Character: Imagery and perception

Perceptual images are phenomenologically similar to perceptual experiences. What it's like to have a visual image of a tomato is very similar to what it's like to have a visual perception of a tomato; the first type of state, insofar as it *feels like* another type of state, feels like visual perception. The properties we experience, identify, and report—the redness, the roundness, the apparent texture—are the same (or close) for the two states. The same is true for other forms of imagery. We attend to and describe an auditory image of a melody in the same ways, and using the same concepts, as an auditory perception of the same melody. (And the points about verbal identification and report is not just a point about disposition or function. Rather, the fact that we identify and report in the way that we do implies that the states possess certain subjective, qualitative feels.) These facts about phenomenology are no doubt (part of) what motivated Hume to say that images, and ideas more generally, are copies of sense impressions: the first feel a whole lot like the second. In terms of phenomenology, perceptual imagery is thus like perceptual experience in character.

Images possess similar character as perceptions along another dimension, namely, functional role. The same point as above can be made here, this time with different emphasis. A visual image of a tomato will *dispose* me to describe it in ways just like (or near enough) a visual perception of a tomato. Studies in the neurosciences support this thought. In a number of experiments involving mental imaging of maps, subjects

scanned and focused on particular features of the mental image in ways and at speeds proportionate to scanning a physical (external) map.[40] People sometimes mistake mental images for perceptions, or have trouble determining if a perceptual memory is one of something perceived or something merely imaged (Segal 1970; Kosslyn 1994: 55; Reisberg *et al.* 1986). Imagery thus functions in ways that are importantly similar to perception, serving to represent the perceptual properties of some actual or non-actual object(s) and disposing the imager to other mental states and actions.

The functional similarities between imagery and perception go deeper. There is substantial evidence that imagery and perception depend upon the same or some of the same neural structures. A number of experiments have shown interference with or degradation of visual imaging capacity when visually perceiving (Brooks 1968; Segal and Fusella 1970; Hampson and Duffy 1984), and interference with visually perceiving when visually imaging (Craver-Lemley and Reeves 1992). Patients with damage to one side of the brain who thus suffer neglect of the opposite side of the visual field, will similarly neglect one side of an imaged scene (Bisiach and Luzzatti 1978). Parkinson's patients unable to recognize facial expressions of emotion are unable to image such expressions (Jacobs *et al* 1995). This kind of evidence has motivated Stephen Kosslyn, among others, to argue for an identification between at least some of the neural underpinnings for the two mechanisms. For Kosslyn, what's common to vision and visual imagery are regions in the occipital lobe of the brain responsible for distinguishing figure from ground, what he calls the *visual buffer*. As he puts it, "Once a pattern of activation is evoked in the visual buffer, it is processed the same way, regardless of

[40] See, among others, Metzler & Shepard (1971); Kosslyn *et al* (1978); Finke and Pinker (1982,1983).

whether it was evoked from input from the eyes (perception) or from memory (imagery)" (Kosslyn 1994; 74).[41]

Imagery and perception are thus similar in character, both in terms of phenomenology and functional role: perception is the *counterpart* of imagery. That said, imagery is not perception. There are important, though perhaps obvious, distinctions between images and perceptions. Here are three. (1) Mental images are not causally efficacious in the same ways as perceptions. Perception, it seems safe to say, functions to provide us with information about our current environment and assist us in our navigation of that environment.[42] It thus possesses commitments to accuracy and action. Images possess no such function. This difference plays out in differences in mental and behavioral efficacy. Images may cause the formation of some belief or desire; images may cause one to act without the mediation of belief and desire. But it is no part of what it is to be an image that it function in these ways. Thus images less often *do* function in this way.[43] (2) Images also seem to be cognitively penetrable in ways or at least to degrees that perceptions are not. We cannot, desire and believe though we may and try as we might, make the lines in the Muller lyer illusion *look* the same length. But we can manipulate the properties of images almost entirely at will.[44] (3) Finally, we can image

---

[41] Currie and Ravenscroft (2002: 71-88) offer a critical discussion of Kosslyn's claims regarding imagery and the visual buffer. They also discuss some of the empirical studies mentioned above.

[42] And here we have in mind something stronger than a Cummins function (see ch. 3, fn. 16 and this chapter, fn. 35). We have in mind some kind of biological function.

[43] These claims are quick: there are many controversial issues here and the conceptual space allows them to be addressed in a variety of ways. One may doubt that images are less efficacious in this respect; one may instead think that images are equally efficacious but simply cause different types of mental states and/or actions; one may worry that there is no principled way to quantify over the causal effects of images and percepts; one may think that the question of biological function is irrelevant; and so on. If after addressing each of these issues, one determines there to be no causal difference between images and perceptions, then points (2) and (3) will have to be sufficient to distinguish the two kinds of states.

[44] Or, given evidence against the cognitive manipulability of image tokens, perhaps we should say it is new images that we may form almost entirely at will (see Lopes 2003). The point needed here is a general one: images are far more amenable to cognitive manipulation than perceptions.

things that we can never perceive, say, golden mountains, mermaids, and Sherlock Holmes, as well as things we cannot perceive at present, say, the face of a friend on the opposite end of the globe, the car keys that we cannot find, and the smell of mom's cookies right out of the oven. Contrast this with a standard conceptual understanding about perception: a perceptual experience requires a stimulus external to the perceiver as the object of perception. In relational terms, a perception requires a perceiver and an external relatum; there is no such requirement for imagery.

4.2.2 Content: Imagery and perception

Imagery is thus like perception in character. It is also like perception in content: our images are about the same kinds of things as our perceptions. This provides a useful constraint upon any theory of imagery: whatever we say about perceptual content, we should say about imagistic content. One might exclude imagery as non-propositional for its lack of propositional content. This is in fact a central, if not *the* central, point of contention in the imagery debate. The question then becomes: is the content of perception, and thus imagery, propositional? Some important conceptual clarifications reveal different ways of answering it.

We need to distinguish propositions from *propositional content bearers*—the things or states that express propositions. One might understand a proposition in terms of sets of possible worlds, namely, the worlds where the proposition is true. Equivalently, one might understand a proposition as a function from possible worlds to truth values. Or one might understand a proposition as a state of affairs, possible or actual. No matter which alternative one opts for, a safe assumption is that propositions are the fundamental

bearers of truth: a necessary feature of propositions is truth-evaluability. When we assign a truth value to 'Vince is drinking beer and eating fries' and 'Vince boit de la bière et mange des frites', we assign that truth value to one and the same proposition, not to the respective English and French sentences. The latter merely express a certain proposition.

We need also to ask what is meant by 'propositional content.' One common gloss is to say that content is propositional if and only if it is sententially describable, exhausted by content sentences. At bottom, this notion suggests that propositional content is just sentential content, possessing the semantic and syntactic features that sentences in some natural language possess. Call this *propositionality*$_d$.

An alternative to propositionality$_d$ is to say that something is propositionally contentful just in case it expresses a proposition. If propositions are fundamental bearers of truth values, then they will stand in certain logical relations with one another: relations of consistency, inconsistency, and implication. If an item or state expresses a proposition, then we should expect it to mirror these logical features. Paul Churchland recommends a notion of propositional content that accommodates these basic intuitions well. According to Churchland, a bit of content is propositional if and only if it functions in a larger inferential or cognitive economy (Churchland 1999: 31, 63-6). The crucial thought here is that what qualifies content as propositional is not its representational medium: it needn't be particularly sentential. It simply must play a distinctive role in an integrated system, standing in inferential relations with other elements of that system. Thus propositional elements, in virtue of membership in some larger system, imply other elements, will be so implied by other elements, are consistent with some elements, and inconsistent with others. Call this *propositionality*$_c$. On this view, not only can a

proposition be realized or expressed by linguistic items from different languages, but it can be expressed by items which are altogether non-linguistic, for example pictures or brain states.

Many parties in the imagery debate have something like propositionality$_d$ in mind. The fact that one camp in the debate, identified foremost with Zenon Pylyshyn, is dubbed *descriptionalism* and the other, identified foremost with Stephen Kosslyn, *pictorialism*, is telling. As Kim Sterelny claims, a representation is propositional according to descriptionalists like Pylyshyn only if the representation is sentential—syntactically structured and generative, involving semantic features like truth and reference (Sterelny 1986). Descriptionalists understand imagistic representation in terms of *structural descriptions*, where the relevant structure is sentential. As Michael Tye puts it, "A structural description of an object is simply a complex linguistic representation whose basic nonlogical semantic parts represent object parts, properties, and spatial relationships" (Tye 1991: 61). For Pylyshyn, such representation is couched in a kind of unconscious inner language (Pylyshyn 1981a, 1981b, 2002; see also Tye 1991: 64-71). Pictorialists seem also to have in mind a sentential notion of propositionality. Thus Kosslyn and Pomerantz understand "propositional format" as abstract "but otherwise similar in structure to that of English sentence format" (1981: 158). They analyze whether propositional format, understood as sentential description, is the appropriate way to understand imagistic content. Elsewhere Kosslyn says, " A propositional representation is a 'mental sentence' that specifies unambiguously the meaning of an assertion" (Kosslyn 1994: 5). Of course, pictorialists like Kosslyn reject propositional representation for (among other reasons) failing to accommodate the unique functional

role of images, rendering them mere epiphenomena (Kosslyn 1980: 29). Pictorialists argue instead that imagistic representation is non-sentential or non-linguistic, and thus non-propositional.

Descriptionalists and pictorialists thus fall victim to the following mistakes. One, it is often assumed (perhaps tacitly) that propositions are syntactically structured like sentences. Images thus must possess a sentence-like structure in order to represent propositions. This is an inference from what is represented—underwritten by a contentious assumption—to the nature of the representation; call it the *proposition-to-bearer inference*. There is also a confusion running the other way. Based upon the sentential nature of "standard" propositional representations, namely sentences (including the content sentences embedded in attitude ascriptions), it is inferred that propositions are sentential. Call this the *bearer-to-proposition inference*. This motivates the same constraint as the first confusion: images must be sentential to express propositions.[45] Either inference may motivate either of the following conclusions: imagery requires a distinct, non-sentential kind of representation (à la pictorialism) or images must be reduced to sentential representation (à la descriptionalism). Both are rooted in a failure to distinguish representation from represented—propositional content bearers (e.g. mental states, sentences) from propositions.

Pictorialists and descriptionalists alike have another option in propositionality$_c$. Perceptual and imagistic contents play particular roles in a cognitive system: a percept or an image might be consistent or inconsistent with or may imply memories, beliefs, desires, associations, may inform the perceiver's conceptual framework, may entail or be

[45] And notice that it might be something like the second confusion (the inference from sentences to propositions) that motivates the assumption regarding propositions central to the first confusion.

entailed by various sub-personal responses in the system. Perceptual and imagistic content is thus propositional$_c$. Perceptions and images, like beliefs, desires and other standard propositional attitudes, function in an integrated cognitive system—though in characteristically different ways and with different cognitive roles—and are thus all propositionally$_c$ contentful.[46]

Choosing between propositionality$_d$ and propositionality$_c$ remains controversial. The important point for now is just this: it is not straightforward that imagery (and perception) are non-propositional states for lack of the right kind of content. It depends upon how one construes 'propositional content.' And one must be careful of the principles motivating this construal in distinguishing propositional mental types from non-propositional types.

4.3 Imagination as generic cognitive capacity: Conclusions and assumptions

What we haven't done: we haven't definitively established that propositional attitudes like belief, desire and their imaginative counterparts possess distinctive phenomenologies. We haven't settled between a sentential versus an inferential notion of propositionality. We haven't solved the imagery debate or determined if perceptual imagery possesses propositional content.

---

[46] Percepts and images involve sensory *phenomenology*—vision-like, audition-like, and so on. This fact is consistent with a propositional$_c$ construal of such states, since propositionality$_c$ is silent with regard to the way that propositional content is presented. Perception and imagery *function* in ways distinct from beliefs and desires. An image of Churchill, for example, will dispose me to certain states; a belief about Churchill will dispose me to others. This distinction in function is consistent with possession of similar content: propositional$_c$ content only requires integration in some larger cognitive system. Application of propositionality$_c$ thus satisfies two pictorialist desiderata: it maintains both the phenomenological and functional similarities of perception and imagery. It also satisfies an important descriptionalist desideratum, namely, commitment to just one system of mental representation. That is, at least, if *kind* of content individuates *kind* of mental representation.

What we have done: We have, through merely analyzing the above issues, made a number of clarifications of and distinctions between propositional imagination and perceptual imagery.

(a) Imaginings can be distinguished both in terms of content, at the token and type level, and in terms of character, at the type level and thus also at the token level. Content is content; though, as we have seen, what makes content propositional or not is an open question. Character can be understood in terms of state function, state phenomenology, or both.

(b) Imaginings like make-belief and supposition do not require an image as part of their character. This is a lesson we learn from Descartes. We can imagine, or *conceive* as he put it, the chiliagon, but we cannot (and, *qua* conceiving, need not) form a mental image of it. So there are some modes of imagination that lack images: I can imagine having a blood type different than the one I have, that Vancouver not Victoria is the capital city of British Columbia, that Reagan was born on May 5th, 1976.

(c) Imaginings of any stripe take various non-imaginative cognitive states as their counterparts: make-beliefs take beliefs, make-desires take desires, visual images take visual perceptions, auditory images take auditory perceptions, and so on. The first two then, take standard propositional attitudes as their counterparts. The second two, and other forms of imagery, take perceptual states as their counterparts.

(d) Point (c) can be made without distinguishing imagery as non-propositional from make-belief and the like. Plausible notions of propositional content are

available—propositionality$_c$ is one example—whereby it is reasonable to say that imagistic content is propositional content.

(e) Imaginative projects and states may be more or less rich. An imaginative project may be rich in quantity of content: we may merely imagine that *P*, or we may imagine an entire set of circumstances in which *P* is embedded. Imaginings may be rich in quality of character, involving more arresting phenomenology, or more complex and expansive cognitive connections, or both.

(f) A point implicit in much of the discussion of this chapter is that any mode of imagination enjoys cognitive freedom: make-beliefs, make-desires, supposition, and mental images all share a lack of commitment—*qua* make-belief, *qua* make-desires, *qua* mental image, etc.—to truth and action. We could imagine in ways that track the truth, or that lead to some action or other, but this is no part of what it is to be an imagining. This sets imaginative states apart from all of their counterpart states, from belief, desire, intention, visual perception, and so on.

This last point is most important. It implies two things. First, it isolates the one thread uncontroversially common to the various modes of imagination, despite any of their differences. It is the cognitive freedom characteristic of all of these cognitive states—make-belief, make-desire, visual imagery, etc.—that encourages us to categorize them as species of the same genus. Our default strategy will thus be to assume that so-called propositional imagination and perceptual imagery are of the same kind. We possess a general cognitive capacity, we call it 'imagination,' and this capacity includes make-belief, make-desire, supposition, visual imagery, auditory imagery, and a variety of

other imaginative states. There are important differences between these states, some of them already identified, some of them surfacing later. But none of these differences are, so we will assume, enough to force the positing of multiple *kinds* of imaginative states. Second, it is the cognitive freedom of imagination, of any of the modes discussed, that makes it ideal if not essential to creative cognition. This feature will be exploited in the chapter that follows.

Chapter 5
The Role of Imagination in Creative Cognition

Novelty should be taken seriously. Although we have argued that any legitimate ex nihilism is unparadoxical, there is nonetheless something special about novelty. Novel properties do, in some sense, emerge from certain organizations of properties and events in a way irreducible to those base entities. If creative thoughts have this property of novelty, as we have argued they do, then the epistemology and etiology of such thoughts will require special treatment. To further motivate this analysis of the phenomenon, we will first consider why knowledge—in a variety of forms—will never be sufficient for creative thinking.

5.1 Knowledge and other truth-bound cognitive capacities

5.1.1 Failures of knowledge

Recall that in certain paradigmatic examples of creative accomplishment all the propositional and procedural knowledge available is insufficient for such accomplishment. This point surfaced most clearly in 3.2.2 where we considered what we called *Type 2* problem space situations: ones where the agent is significantly constrained by the conceptual space in which he is operating but augments that space in some way. Bach might have known everything there is to know—both in terms of all of the facts and all of the relevant skills—about the clavier, tempered tuning, and the twelve-tone scale, and this would not have been sufficient for the creation of *The Well-Tempered Clavier*. Or better, an omniscient being might have complete knowledge of all of the music-theoretical space of the time and would not yet have knowledge of the musical structure

of *The Well-Tempered Clavier*. Bach was of course importantly constrained by this knowledge (or at least by what he had of it, since he presumably didn't have *all* of it). But this knowledge alone does not amount to, afford, or even imply the musical work in question. This is for the simple reason that there is nothing in this conceptual domain, or cluster of domains, that engenders (by itself) *The Well-Tempered Clavier*. The point then is not about Bach or about any agent, but about how far *truth-bound* cognitive states can take us towards thinking creatively. Any cognitive state that faithfully tracks the contents or information of some conceptual space—be it a true belief, propositional or procedural knowledge, an accurate memory—can at best play a necessary but insufficient role in the thinking required for an accomplishment like *The Well-Tempered Clavier*. (So 'truth-boundedness' targets not the epistemic justification component but the truth or accuracy component of mental states that constitute knowledge. 'Truth' here is used in a broader, stipulated sense, since it must accommodate procedural kinds of knowledge or skills—one learns more or less *accurately*, not truly, how to make an omelet or install a carburetor.) In cases like these, and *a fortiori* in Type 3 problem space situations where the agents transcend both the conceptual map and the space (e.g. Kekulé's discovery of the structure of benzene), knowledge of the relevant conceptual space is never enough for a creative solution of the task or problem.

We will now advance the following stronger thesis. Any instance of creative thinking will require more of an agent than a truth-bound cognitive state. Even mundane creative instances—say, Norm's learning of quantum theory and other examples in *type 1* problem spaces—will require non-truth bound states to enable the creative thought or thoughts. We will call this, for reasons made clear below, the *cognitive manipulation*

thesis. Before returning to Norm, we will consider two sets of studies from developmental psychology that support this thesis.

First consider three related studies on drawing capacities. Annette Karmiloff-Smith solicited drawings of nonexistent houses, people, and animals from children ranging in age from 4 to 11 (Karmiloff-Smith 1990; see also Karmiloff-Smith 1992: 155-61). Christina Cacciari and colleagues solicited drawings of imaginary houses and animals from children in the same age range (Cacciari et al 1997). Thomas Ward performed similar studies on creative tasks of adults, asking them to imagine and draw nonexistent creatures (Ward 1994, 1995). The data compiled motivates a number of interesting hypotheses. One hypothesis, common to all three, is that children and adults alike are highly constrained by their existing concepts: what a person knows about categories like HOUSE, PERSON, or ANIMAL will significantly constrain how that person is able to represent novel instances of such categories. Although frequency of cross-category combination increases with age—for example, older children are considerably more likely to anthropomorphize a house by adding sense organs and limbs to their drawings—the properties from any one category are relatively stable. Ostensibly then, individuals "retrieve a specific instance of a given category and pattern the new creation after it, regardless of whether they were required to imagine and draw an artifact such as a house or a natural kind such as an animal" (Cacciari et al. 1997: 157).

So even given invitations to create nonexistent things, the drawings were largely predictable—conventionally generated by the subjects in line with the relevant conceptual schemes. Nothing radically novel here. However, we should know by now not to withhold attributions of creativity on such bases. Instead, grant for now that many

of the drawings are minimally creative. We should then ask if any of these drawings were possibly enabled *merely* by the relevant conceptual knowledge? That is, we know that the subjects consistently employed their concepts of 'house', 'person', etc., to make their drawings; was this knowledge sufficient? The obvious answer seems to be no. The concepts of 'house' and 'person,' no matter how rich, will not (by themselves) enable a child to draw a house with eyes for windows, a mouth for a front door, and arms and legs. These cross-category changes require the child to cognitively manipulate the conceptual space in particular, though perhaps minimal, ways. And this point is not unique to children: adults may have all the information possible for some set of concepts or conceptual space and still not have the thought to combine them in ways needed to yield "humanized house" drawings. We thus have minimally creative cognition that requires more than truth-bound cognitive states.

The second set of studies concerns the development of linguistic competence, with an emphasis on the acquisition of figurative competence—that is, the comprehension of metaphors, idioms, proverbs and the like (Cacciari et al 1997; Cacciari and Levorato 1989; Gibbs 1987, 1991, 1994; Levorato and Cacciari 1992, 1995). One finding of these studies is that children as young as 7 years of age are able to understand and use figurative language. The development of this competence is based not in rote learning mechanisms but in what some call *figurative competence*, which consists in a suite of abilities responsible for general semantic competence and linguistic understanding (Cacciari et al 1997; Levorato and Cacciari 1992, 1995). Such abilities include the apprehension of a variety of meanings for a single lexical item, suspension of purely literal or referential linguistic strategies, awareness of linguistic conventions, and

the importance of the context of utterance (Cacciari et al 1997: 159). The interesting moral for our purposes is that learning a word or phrase, be it a colorful simile or a dry technical term, is rarely enabled by simple rote learning: we do not simply memorize a meaning and syntactic role for some lexical item. Instead, language acquisition typically requires more than truth-bound cognitive states, more than simply entertaining and assenting to the information contained in the relevant conceptual space. It requires some degree of cognitive manipulation: some consideration of the potential for multiple meanings (could 'bank' refer to this kind of thing *and* this other very different kind of thing?), multiple syntactic roles (can 'bank' function as a verb or a noun?), conventional and contextual factors (if we are talking about rivers, we mean this kind of bank, money, that kind of bank). These considerations involve more than the formation of true beliefs and accurate skills.[47]

Development of linguistic competence, figurative or literal, may or may not be creative cognitive activity. No matter. Even assuming that it is not, it nonetheless requires non-truth bound cognitive states. We thus have additional reason to think that cognitive activity that *is* creative requires something more from the cognizer: it requires more than rote memorization, more than simply comprehending and assenting to some set of information, more than mere knowledge or truth-bound states.

We may return now to poor Norm. Recall that for Norm, thoughts about quantum mechanics are novel. Grant also that Norm has successfully *learned* some of the basics of quantum theory, rather than simply memorizing them without comprehension. Assume for the sake of simplicity that quantum theory comprises just three propositions,

[47] Or better put, even if we take such formation to be the end result—the acquisition of propositional and procedural knowledge—this result involves as its means some non-truth-tracking states. Such formation is not enabled by "merely reading off" the information contained in the conceptual space.

A, B, and C. Assume further that the tenets of our toy theory are expressed by a simple text called *QT*. Given that these propositions are entirely novel for Norm, he cannot *just* read *QT*—the linguistic expressions of A, B, and C—to afford any kind of comprehension of quantum mechanics. Put another way, Norm cannot just entertain the thoughts with contents A, B, and C and thereby come to know how A, B, and C form the theory or come to know how to use the concepts that structure A, B, and C. Such novel concept and knowledge acquisition requires more substantial cognitive manipulation. This point is no different from the one above regarding general linguistic competence: the learning of many words and phrases, where this is most salient with figurative language, is not a matter of mere memorization. One must consider a variety of features before one is able to use a term or concept appropriately. And if proper use is a reliable indication of acquisition of linguistic competence (of learning), then these considerations are necessary conditions for such acquisition. This point applies to the learning of figurative, literal, and technical language—to learning 'break the ice', 'bank', or 'superposition.'

All of these proposals support the strong cognitive manipulation thesis with which we began: any instance of creative thinking, no matter how mundane, will require more of an agent than truth-bound cognitive states. We might think of the thesis this way. For any conceptual space C, which provides the intentional object/s for a set of thoughts T tokened by some agent, if at least some of the thoughts in T are creative, there must be a non-isomorphic mapping between T and C. There must be some mental tokens in T not immediately derivable from C; or, as we have been saying, T cannot be read directly off of C. Truth-bound states function in just this way: they aim to simply mirror a

conceptual space. If an agent employs *only* these kinds of states, the mapping from (her) T to C will be isomorphic, and T will not be p-creative.[48]

We might also think of the point in the terms of emergence discussed in 1.3.4 Creative thoughts are psychologically novel, and thus novel with respect to the cognitive agent in question. If some thought *t* is novel with respect to some cognitive profile, then *t* is not reducible to any one element or set of elements tokened in that cognitive profile. Rather, it emerges from that profile given some re-organization of, fusion of, or addition (or combination thereof) to the profile. If this is right, then the thesis under consideration is a natural consequence: mere appeal to knowledge and memory will not yield any emergent states or properties. Some manipulation of the contents of these states, plus perhaps some other contents, is needed.

We have thus identified another cognitive role essential for creative thought. Creative thought requires non-truth bound manipulation. Call this the *cognitive manipulator role*. In the terms of the functional analysis from chapter 3, for creative instances of (T_R), cognitive manipulation will be an occupant of P_4: the former is what we might call a sub-role of the latter. In theoretical language, cognitive manipulation will be part of the heuristic set leading to the creative solution (where without this heuristic, the agent *could not* token the creative thought; see 1.3). The obvious next question is this: what kind of cognitive state is best fit to serve this functional role?

---

[48] This is an oversimplification: we do form false beliefs and inaccurately learn certain skills. So employment of all and only truth-bound states does not entail an isomorphic mapping between one's cognitive states and the conceptual space they aim to represent. It is thus possible that one could token a creative thought by means of purely truth-bound states, namely, by forming some false beliefs. This is, however, an unlikely model for a significant amount of creative thought. The simplification will thus be taken for granted in what follows.

Before answering this question, we should clarify the importance of knowledge and truth-bound states. The point of this section is not that knowledge, concepts, or memory are not important for creative thought: it is likely that they are essential. The point is just that they are insufficient. This proposal is consistent with theorists who argue for the importance of knowledge in creative cognition (e.g. Bowers et al. 1995;Ward 1995; Weisberg 1986, 1995, 1999). It is, however, inconsistent with reductionist views, like Weisberg's (see 2.2), that argue that knowledge is the whole of the story. We thus have not identified an outright *failure* but instead a shortcoming of knowledge.

5.1.2 Candidates for cognitive manipulation

The Darwinian model of creativity, pioneered by Campbell and carried on by Simonton, provides a candidate mechanism for cognitive manipulation (Campbell 1960, 1965; Simonton 1999). Simonton advances an even stronger thesis than the cognitive manipulation thesis: he argues that education and knowledge reaches a threshold, after which point it becomes counterproductive to creativity. Whether or not this thesis is plausible, it implies that knowledge is insufficient for creative thought and thus something else must play a role in generating creative ideas. For the Darwinian model, the relevant mechanism is a variation-generator, which creates random (non-truth bound) variations on existing concepts and ideas. The selection and retention-mechanisms will then determine which of these ideas are endorsed and communicated.

Another option is to follow Martindale and his cortical arousal model (Martindale 1977, 1981, 1995, 1999; Martindale and Armstrong 1974; Martindale and Hines 1975).

According to Martindale, creative ideas are the products of creative minds, which more readily enter a special kind of cognitive processing—what Martindale calls *primary processing*. What's most important about this subconscious kind of cognition is that it involves a wider range of attention focus. That is, when an agent responds to a problem as such, her cortical arousal decreases, while the range of contents that she may entertain (again, largely subconsciously) increases. Primary processing is in no way bound to acquisition of knowledge and, according to Martindale, is what leads to creative insights. (See chapter 2 for discussion of both models. Further discussion of Martindale in chapter 6.)

Chris Thornton proposes a less revisionist nothin' special model of creativity (Thornton 2002). Thornton proposes that creative thinking might either reduce to, or at least be importantly tied to, more fundamental learning abilities. Understanding learning as a task of behavior and/or concept acquisition, acquiring some new skill or concept *c* involves a recursive process of aligning *c* with some set of data. For example, to properly learn the concept "greater-than", I must apply it consistently to the data, in this case integer values. Once such alignment takes place, the data is re-coded. Creative thinking occurs, Thornton proposes, when this recursive process carries on after the conceptual alignment: the data is continually manipulated and interpreted, often subconsciously, sometimes resulting in novel re-constructions of data and uses of concepts. Thornton calls this potentially creative process *runaway learning*.

We thus have three examples of cognitive mechanisms or capacities that might fill the cognitive manipulator role. However, all of these theories are inadequate for the following reason. Subconscious thought is no doubt an important feature of creative

thought.[49] It is this feature that has both inspired mystical accounts of creativity and deterred the more empirically minded from giving an account at all. The Darwinian, the cortical arousal, and the runaway learning models each suggest an empirical accommodation of this feature of creative thought: variation-generation, primary processing, and offline learning are largely subconscious. However, just as it is a mistake to overlook this feature, it would be a mistake to locate the mechanisms of creative cognition so as to render *all* (or most) such cognition subconscious. A quick consideration of the phenomenology of creative thought reveals the following fact: creative thoughts do sometimes come when one is actively struggling with a task or problem. We do, at least sometimes, *consciously* think creatively. None of the proposed mechanisms will accommodate this feature of creative thought. This problem for each theory, coupled with those earlier discussed (see chapter 2), should force us to keep looking for a suitable occupant for the role.

All of this motivates the following job description: We need a cognitive mechanism that fills the *active* cognitive manipulator role. This is essential to P_4 of creative problem solving. An incubatory mechanism is also desired: a cognitive manipulator that can also function at a largely subconscious or less attentive level would be ideal.

[49] Much of the next chapter aims to provide conceptual clarity and empirical support to "incubatory" cognitive processes.

5.2 Imagination and cognitive manipulation

5.2.1 Imagination as play

Imagination is traditionally characterized as being playful, answering to names such as 'pretence', 'pretend play', 'role-playing', and 'make-believe', among others. To engage in imaginative activity is to engage in a cognitive activity that, as it were, carries smaller stakes for (epistemic and bodily) action. In fact, on one defensible model of imagination, this is one of the very features that distinguish it from other cognitive states like belief, desire, and intention: imaginings do not, qua imaginings, immediately cause action.[50] An instance where this seems most clear is in certain theoretical responses to the paradoxes of fiction. I fear the slime that comes at me as I sit in the movie theater, but this fear is underwritten by imaginings rather than beliefs that I am in danger. This is evidenced by the fact that in spite of my fear, I do not flee the scene in self-defense but rather squirm in cinematic enjoyment. This indicates both that imaginative states are responsible for emotional responses to fiction and that those states are not connected with action in the same way as their doxastic and orectic counterparts (Currie 1997; Walton 1990). Even if this is a special case, we nonetheless have evidence that imagination is relatively decoupled from action at least some of the time.

Berys Gaut takes just this feature of imagination to be the one that makes it especially well-suited for creative thought (Gaut 2003). According to Gaut, imagination is a *vehicle* for active creativity. As contrasted with beliefs and intentions, imaginings lack intrinsic commitment to truth and ends. It is a conceptual fact about belief that beliefs aim for the end of truth, and a conceptual fact about intention that intentions aim

[50] Although this very issue is a point of contention in the philosophical literature on imagination. See Currie and Ravenscroft 2002; Funkhouser and Spaulding (in preparation); Nichols and Stich 2000,2003; Velleman 2000). See dicussion in 4.1.3.

for the end of action. They are thus, in some sense, committed to these ends or results: to believe P is to be committed to the truth of P and to intend to do Q is to be committed, *ceteris paribus*, to doing Q. Imagination bears, by itself, no such commitment. Gaut offers a compelling test case to this end. Moore's paradox tells us that it is problematic to assert "I believe that it is Tuesday, but it isn't Tuesday." Analogously, an assertion like the following is problematic: "I intend to take a vacation, but I won't take a vacation if given the chance." However, there is nothing paradoxical about either of the following: "I imagine that it is Tuesday, but it isn't Tuesday"; "I imagine taking a vacation, but I won't take a vacation if given the chance." Imagination thus enjoys a freedom that other states lack. It is this freedom from commitment that, according to Gaut, makes it an ideal vehicle for creative thought: for trying on various hypotheses, considering counterfactual situations, teasing out possible implications and consequences, etc.

Peter Carruthers also recognizes the importance of the playfulness of imagination (Carruthers 2002). According to Carruthers, the playfulness of childhood pretence provides practice for adult creative cognition. In fact, he takes childhood pretence to serve the evolutionary function of preparing and enhancing adult creative capacities. Carruthers endorses the basic pretence model of Shaun Nichols and Stephen Stich, which proposes that pretence requires two cognitive mechanisms: a *supposition-generator* and a *possible worlds box* (Nichols and Stich 2000, 2003). The first mechanism does just what its name suggests, and the second is a memory system for storage of the generated suppositions and their elaborations. These suppositions and their elaborations, which may vary in richness and complexity, bear no intrinsic commitment to truth or action. So

again, we see that imagination is especially well-suited for generating thoughts with novel content.

Gaut and Carruthers each offer a handful of additional suggestions, and are in many ways onto different projects. The point to note, however, is what's common to the two. Both recognize the importance of one particular feature of imagination, namely, its being decoupled from truth and action. Imagination, unlike belief and intention, allows one to entertain contents without entailing (by itself) epistemic or behavioral commitments. Imagination is thus distinct from belief, memory, and skill-acquisition in the relevant way: it is not what we have been calling a truth-bound state.

This point alone does not suffice to show that imagination fills the cognitive manipulator role. Grant that imagination is a playful kind of cognitive state, unbound to action and truth. If imagination does involve such apparently free-wheeling cognitive engagement, why should we think that it is *serious enough*, as it were, to engender thoughts whose bearers (and their received audiences) will take seriously? If imagination is non-truth bound, is it bound in any way such that it will provide or enable products of significant cognitive value? Our task, then, is to determine the cognitive purchase, if any, of imagination.

5.2.2 Imagination as work

It proves useful to consider some of the connections between imaginings and other elements of cognition. There are at least two types of connection between

imaginings and other mental states and capacities: *inferential* and *cognitive*.[51] The first includes *fictional inferential elaboration*, *actual inferential elaboration*, and *inferential isolation*.[52] Let's briefly consider each. First, we often subject our imaginings to ordinary inferential practice; we draw inferences from our imaginings in the same ways we would if we had beliefs with the same contents. If we imagine that P, and we imagine that if P then Q, then we will be disposed to imagine Q just as we would in actual circumstances of logical reasoning. This is one way we might inferentially elaborate a fictional project. Here is another. We will often supplement our imaginings with actual beliefs in order to render the imaginative project coherent and consistent. So for example, if I am told by a narrative that Jones has blood all over his shirt and the story has been (explicitly or implicitly) a realistic one about humans, I may well infer that Jones has *red* stains all over his shirt, given my beliefs about the color of human blood in the actual world. This is a detail that the author can safely leave out, and that I am warranted in inferring.

Inferential elaboration can also run the other way. Imaginings can be used to supplement reasoning about the actual world. This direction of inferential trading is in many ways the crux of (some forms of) simulation theory. The thought here is that in order to mindread, we use simulations (imaginings) to draw inferences about the mental states and actions of others. So this is one plausible candidate for actual inferential elaboration in ordinary practice. In the theoretical domain, there is a long tradition of taking conceivability as a guide to possibility; we might infer possibilia from

---

[51] Note, I am following others in using these terms. The terms are misleading in some ways, given that the inferential connections are obviously cognitive as well. The mark that distinguishes them is inference: cognitive connections need not and often do not involve any kind of inference.

[52] Both Nichols and Stich (2000, 2003) and Gendler (2003) use 'inferential elaboration.'

conceivabilia. Here again, we are using imaginings alongside beliefs to reason about the world, or at least to reason about the modal status of objects of the world.

Inferential isolation opposes inferential elaboration. Our imaginative projects are sometimes immune or blocked from normal inferential practice. Such projects might tolerate inconsistencies, incoherence, or incompleteness—all unequivocal no-no's by the lights of traditional logic (see Walton 1990: 57-67). Moreover, our actual beliefs are often barred from playing any role in the inferences drawn within the imaginative project—sometimes called *backgrounding*, see 4.1.3—perhaps because they are inconsistent with certain imaginings. So I might imagine of myself that I am both rich and poor (inconsistency/incoherence), or might leave out the name of the woman about whom I fantasize (incompleteness), or might hypothetically suspend a handful of beliefs about physics in order to imagine traveling through time (backgrounding).

With these connections in mind, we see better how supposition is a kind of imaginative state (though perhaps impoverished in certain ways, see 4.1.3 and below; see also Currie 2002; Currie and Ravenscroft 2002; Gendler 2000: 80-1), while entertaining is not. The proposed inferential connections seem to hold between suppositions *and* beliefs and ordinary inference, but not between entertainings *and* beliefs and ordinary inference. Like imagination, supposition involves doing something with a proposition. We can suppose that certain counterfactuals are true and then *elaborate* them via ordinary inferential practice. Alternatively, we can also *isolate* the supposed propositions from ordinary inference and consider them only with respect to a limited set of additional propositions or inferential rules. This is not the case with entertaining: entertaining is not yet at the requisite level of cognitive activity. Entertaining a proposition simply doesn't

involve any kind of commitment enabling it to work into inferential practice (Walton 1990: 19-21). So with regard to inferential connections, supposition, but not entertainment, is like to imagination.

Now for the cognitive connections between imaginings and other cognitive states. Here there are at least two kinds of connections: *cognitive quarantine* and *cognitive contagion*.[53] The first can be understood as follows. In spite of the fact that beliefs and imaginings can token the same propositional content, we keep the attitudes with regard to such contents distinct. In short, we generally do not confuse imaginings with beliefs (or other propositional attitudes). Moreover, imaginings do not usually cause actual beliefs or desires or a change in intentions, values, or other mental states. It is in these senses that imaginings are kept quarantined.

But, one might quickly object, imaginings are sometimes not quarantined as such. Imaginings often relate to other mental states in precisely the ways mentioned. So if we keep imaginings quarantined, there must be some leaks. There is a distinction here that proves useful. We should distinguish the causal *influence* of imaginings—an imagining causing some other mental state(s)—from the *identification* of imaginings with other states—an agent's confusing an imagining for some other mental state(s). Barring imaginings from having causal influence would involve *strong quarantine*; keeping imaginings individuated from other mental states only requires *weak quarantine*. Weak quarantine seems to be more effective: we rarely mistake a make-belief or a make-desire for their respective counterparts.

Currie and Ravenscroft propose an interesting exception to even this weaker kind of quarantine. They argue that schizophrenics and other patients experiencing so-called

[53] These terms are Gendler's (2003).

delusional beliefs suffer from a failure to make precisely these distinctions. Such patients fail to recognize imaginings as imaginings, treating them in ways more like belief (Currie and Ravenscroft 2001: 161-84). This is a candidate violation of weak quarantine, what we will call *cognitive contagion.* Currie and Ravenscroft also cite evidence for contagion in perceptual imagery. People sometimes mistake mental images for perceptions, or have trouble determining if a perceptual memory is one of something perceived or something merely imagined (Segal 1970; Kosslyn 1994: 55; Reisberg *et al.* 1986). A more common kind of contagion involves a violation of strong quarantine. Imaginings influence or cause other mental states. There are a number of ways that this might occur. In a phenomenon known as *filling in*, a visual percept is completed by concepts, memories or images had by the perceiver. Filling in is common in persons who suffer from visual scotoma—large blindspots in the visual field (see Ramachandran 1998: 85-112). I might make-believe that God exists, and come to actually believe that God exists. I might make as if to desire a beer, and come to actually desire a beer. I might visually image the couch fitting through the doorway and determine that the couch will fit through the doorway (and then act upon this judgment). In spite of weak cognitive quarantine (i.e. keeping imaginings individuated from non-imaginings), imaginings are incredibly efficacious in this respect.

These cognitive connections partly distinguish supposition from richer imagination, and further distinguish entertaining from imagining. If entertaining just involves grasping a propositional content, or getting it "before the mind's eye", as we say, then it lacks any of the commitments requisite for the above mentioned causal efficacy. We do not confuse entertainings for beliefs or other committal states, and the

former do not cause (independently) the latter. Entertaining the proposition that God exists will not, by itself, get me any closer to believing in God. Entertaining the proposition "I am going to lose my job" does not issue in the anxiety that believing or imagining the same proposition does. Supposition misses other cognitive states in the same way: supposing that the boogey-man is behind the shower curtain will leave me emotionally unaffected; this is one sense in which supposition lacks the richness of more engaged imagination. It does, on the other hand, bear to other cognitive states some of the same cognitive connections born by richer imaginings. Although strong quarantine of supposition will rarely if ever be violated, weak quarantine will. Supposing that P may cause (or at least be causally relevant in some non-trivial way) beliefs and other more committal cognitive states.[54] These are results we should embrace: given the complete lack of commitment of entertainment and the ostensible importance of supposition in the generation of theoretical hypotheses, we should want to deny the first and maintain the second as candidate cognitive manipulators.

5.2.3 Imagination as candidate cognitive manipulator

Clarification of these inferential and cognitive connections is instructive: we now see that in spite of its playfulness, imaginative states play important, cognitively valuable roles in human cognition. Imaginings, including suppositions, integrate with other states in inference, and causally influence our cognitive commitments to the world around us. At the same time, they are not entirely unconstrained. Our use of imaginings in inference

[54] Whether this is cognitive contagion or actual inferential elaboration is unclear. No matter. These distinctions are not absolute: they are merely intended to clarify the ways that imaginative states relate to other cognitive elements. The point here is that imaginings, including suppositions, influence our beliefs and inferences about the actual world in important ways.

and belief formation, among other cognitive operations, may be constrained by the same kinds of considerations that non-imaginative inferences and beliefs are, namely, considerations of validity, soundness, evidence, reliability, and prudence. This point is important, since it reveals that imagination is not all play and no work.

But neither is imagination all work and no play: it *is* playful. Gaut and Carruthers are right to concentrate on this feature of imaginative capacity. It is this playfulness—lacking the epistemic and behavioral commitments that other states bear—that enables cognitive manipulation. It allows for safe consideration of counteractual states of affairs, for consideration of hypothetical actions, for revealing the implications of a variety of theoretical alternatives, for trying on different solutions to the same problem or task. At the same time, it is a cognitive capacity the outputs of which we may take seriously: this is the moral of the inferential and cognitive connections identified.

Imagination, as a cognitive state that fuses playfulness with cognitive purchase, and may be fused with what we might call cognitive responsibility (i.e. it is at least a capacity that can be used more or less responsibly), is ideally suited to serve the role of cognitive manipulator. It is not truth-bound, and so provides the freedom needed to generate novel cognitive states. It allows agents who, when working from within a conceptual space, to augment and transcend the map (as Bach did with the *Well-Tempered Clavier*) or to move off the map and the space itself (as Kekulé did with the benzene molecule). As we have already argued, these achievements require some non-truth bound cognitive manipulation: neither Bach nor Kekulé could read the available maps, as it were, and then token the creative thoughts in question. They might have possessed all the knowledge possible for the relevant conceptual spaces but this would

not have sufficed. The reasons are simple, though slightly different, for each type. For Bach, in spite of his self-imposed constraints—which derive from the conceptual spaces of 12-tone scales, clavier instrumentation, and tempered tuning—there was no map of these spaces and their combinations in their entirety. Bach made the map, and we have the *Well-Tempered Clavier* to enjoy as a result. For Kekulé, there was not only no map, but insufficient conceptual space: Kekulé's discovery of the ring-like structure transcended the conceptual space of organic chemistry (up to that point). So in both cases, counterfactually, had either thinker relied solely upon states that accurately tracked the maps and space available, they would not, in fact *could not*, have made their creative breakthroughs. Their imaginations occupied the additional cognitive role.

The same is true for Type 1 problem spaces. We have argued above that Norm's learning quantum mechanics—in a way analogous to language acquisition and cross-categorical representation in both children and adults—involved something more than just entertaining the relevant information, as expressed by our quantum text, *QT*. Granting that some of Norm's thoughts about quantum theory are creative, then they must be novel with respect to his mind. He thus had either to re-organize or combine thoughts already possessed or acquire some new thoughts about new information (and given Norm's quantum ignorance, we should bet on the latter). Learning of this sort, like the child's learning of figurative expressions such as 'crying over spilled milk' and 'the last nail in the coffin', is not gotten by rote memorization. Rather, one has to consider a variety of linguistic and conceptual features. This kind of consideration involves asking questions of the sort: how does this connect with that? could this mean something like this? what does this imply? what if this word is being used differently than before? In

short, it involves cognitive manipulation. Imagination serves this role well. One can suppose, imagine that, and mentally image in order to ask and answer these kinds of questions. And one can do so without significant commitment to truth or action. It is likely then that all of these examples of Type 1 problem spaces—Norm, the Caciarri/Karmiloff-Smith/Ward cross-category drawings, and figurative language acquisition—involve the use of one's imagination. We use our imagination, at varying degrees of richness, to perform the cognitive manipulation necessary for these kinds of tasks.

Situating these claims within the functional analysis of chapter 3, we get the following point. P_4 (the role identified with the theoretical term 'heuristic set') varies depending upon the other three roles (those roles identified with 'context,' 'problem,' and 'solution set') and the agent in question. We have been arguing that if P_3 is to be named creatively (a creative solution secured) then P_4 must involve some non-truth bound states. Imaginings of various sorts seem to be just the right kind of states. The upshot: problem solving, creative or not, will often involve the use of one's imagination.

But we should quickly note, imagination is not enough to accomplish such cognitive tasks, be it the learning of quantum mechanics or the construction of a musical masterpiece. As we claimed in 1.2 and 1.3, creative cognition emerges from particular cognitive profiles. These profiles will comprise a variety of states and capacities: knowledge, concepts, skills, memories. Creative cognition is not creation (strongly) *ex nihilo*.

5.2.4 Imagination: Conscious and subconscious

Recall the job description from the end of 5.1.2. We identified one essential and one desirable feature for the cognitive mechanism needed, active cognitive manipulation and subconscious cognitive manipulation, respectively. We have spent more time above and will spend more time here discussing the cognitive architecture of the first. While the second will receive treatment in 6.3.

Berys Gaut makes a relevant distinction, between what he calls *active* and *passive* creativity. "Passive creativity occurs when the subject is unaware of the creative process, if any, which has occurred to produce the creative outcome." "[A]ctive creativity occurs when the subject actively searches out various solutions, consciously trying out different approaches, and in the course of this activity comes upon a solution" (Gaut 2003: 192-3). There are two dimensions to Gaut's distinction: phenomenology and conscious intention. In *actively* creating, one deliberately attempts to solve a problem and is conscious of one's doing so. While in passively creating, one is not consciously working on a problem, or even intending to solve it (if we interpret the first quotation strongly), and the insight comes in a seemingly unbidden flash. There are problems with the latter understanding. Either it isn't creativity we are talking about or the passivity is just overstated. If one is not aware of "the creative process, if any" then one is hardly responsible for the outcome in a way that would render the action creative. So perhaps the point is just overstated. Gaut might be making the same mistake that many have made before him: inferring from flash phenomenology that the engendering actions are non-deliberate or unconscious.

We can let these worries pass: Gaut's core distinction is a good one. There is a distinction between actively thinking in creative ways and passively thinking in creative ways, even if this difference is merely one of degree and not kind. Gaut also has tradition at his side: theorists of creativity have made subject matter of both the conscious efforts that creative individuals exert in their projects and the bursts of insights that, ostensibly, fall in their laps. How is imagination situated on this continuum?

As discussed in 5.2.1, Gaut takes imagination to be particularly well-suited to serve as the *vehicle* for creativity.[55] He isolates this importance to active creativity. Given the discussions of both the playfulness and the workfulness of imagination in 5.2 above, the point here should be obvious enough. Coupling its lack of certain epistemic and behavioral commitments with its mental causal efficacy, imagination serves well the active solution of problems. This is especially true if those problems are solved in creative ways.

We thus actively, though sometimes without thinking about the fact that we are imagining, use our imagination for a variety of cognitive tasks. This should be no surprise. But just in case there is any doubt, follow these instructions. Imagine a circle. Now add a straight line to the bottom of the circle. Bisect that line at its midpoint with another line running perpendicular to it. Now add an equilateral triangle to the bottom of the first line. Then remove the bottom segment of the triangle (which should parallel the second line). If you followed me, you should be left with a mental image of a common symbol for man, a stickman. If not, who knows what you ended up with. It doesn't

---

[55] Gaut makes a distinction between the *source* and the *vehicle* for creativity (Gaut 2003: 195). He argues that imagination cannot be the source of creativity, since other cognitive states are required—knowledge, skills, motivation, etc. Gaut's claim seems to be a causal one, and is true as far as it goes. However, it only entails that imagination is not the *only* source of creativity. It remains consistent with the thesis that imaginative capacity is a necessary but non-sufficient cognitive condition for creative thought and action.

matter: either way you were using your imagination in an active and conscious manner. In this instance, someone else was issuing the directions, but the execution is no different in principle from instances when we direct ourselves.

On to subconscious imagining. For now, the following claims coupled with a promissory note for further treatment in 6.3 will suffice. Our imaginative capacities are not limited to deliberate or even conscious employment. Our imaginative projects often take on a life of their own, even when they begin under deliberate, strict mandates. They also continue when we are not, phenomenologically (i.e. from the perspective of our own experience), thinking at all. Or at least, they may continue when we are not consciously working on some cognitive task or other: when we are dreaming, daydreaming, or even working on an altogether different cognitive task. This proves especially important since it gives us a familiar cognitive capacity to accommodate subconscious and incubatory creative insights, a phenomenon that has garnered a lion's share of attention in creativity literature since the days of Plato. Imagination is, we might say, *incubation-functional.*

5.3 Imagination as necessary to creativity

The considerations in 5.2 imply that imagination is an important component in most creative cognitive processes. If this caps the strength of our thesis, we have done well. We might consider though, the following stronger thesis. Imagination is necessary for creative thought, even when minimally creative. There is no creative thought without imagination playing a role. What reasons do we have for endorsing such a thesis?

In advancing the cognitive manipulation thesis, we argued that creative thoughts are of a sort that requires for their tokening non-truth bound states. Given the novelty of

creative properties, some cognitive manipulation is needed to engender them. Knowledge thus fails. Memory thus fails. In 5.1.2. we rejected the candidate mechanisms of the Darwinian theory, cortical arousal theory, and runaway learning theory. Mere entertainment fails for reasons given in 5.2.2. What is left? Imagination seems to be the only viable option and is especially well-suited to fill the role of cognitive manipulator. This gives us: (a) If creative thought requires cognitive manipulation, then creative thought requires imagination.

Roughly the same point can be made in simpler conceptual terms. Creative properties emerge from a cognitive profile, but cannot be reduced to the base states of that profile. In thinking creatively, we thus must either cognize the contents we already have in some new ways (re-organizing or combining them) or we must cognize some novel content. Either way, we must engage in some degree of counterfactual thought. We must consider alternatives not already known or thought of, non-actual and counterfactual possibilities, new ways of combining information. Call this broad range of tasks *counterfactualization.* Imagination, including supposition, is the most obvious candidate capacity for counterfactualization. We thus have the following: (b) If creative thought requires counterfactualization, then creative thought requires imagination.

The cognitive manipulation thesis can be employed from additional angles: it also applies to other, potentially non-creative, cognitive capacities. In considering the developmental experiments on children's and adults' drawings, we found that cross-categorical drawings (e.g. fusing the properties of a house and a person) required something more of the agent that mere conceptual mastery of the relevant domains. The child had not only to understand what a house is and what a human is, but then had to

combine the properties of these categories in order to get a novel representation, a humanized house. This motivates the following argument: (c) If cross-categorical representation requires cognitive manipulation, then it requires imagination (granting the cognitive manipulation thesis). If creative thought requires cross-categorical representation, then it requires cognitive manipulation. If creative thought requires cognitive manipulation, then creative thought requires imagination.

The same move is available with respect to figurative and much of literal language comprehension. Cacciari et al. argue that figurative language acquisition, like much of literal language acquisition, requires the formation and honing of certain linguistic competencies not gotten via rote learning. We argued that this indicates a need for cognitive manipulation. This motivates (d) If figurative (and literal) language competence requires cognitive manipulation, then it requires imagination. If creative thought requires figurative (and literal) language competence or employs the same cognitive mechanisms, then it requires imagination. So if either of these capacities—cross-categorical representation or figurative language competence—is basic to creativity, then creative thought requires imagination.

Paul Harris argues that imagination is necessary for mental model building. Mental model building is essential, according to Harris, for functioning in a community where we exchange information about hypothetical situations. We provide testimony in the form of texts and verbal discourse about the not-here-and-now. Imagination serves to mentally model such hypothetical situations, and thus contribute, participate, and understand discourse about events that happened elsewhere, at some other time, in the future, or merely possibly (Harris 2000; see Carruthers 2002: 235-7 for criticism). This

motivates (e) If creative thought requires hypothetical thought, then it requires mental modeling. If creative thought requires mental modeling, then it requires imagination. Thus if creative thought requires hypothetical thought, then creative thought requires imagination.

As mentioned, Carruthers argues for an intimate connection between imagination and creativity (Carruthers 2002). He argues that all of the cognitive capacities requisite for creative thought have been in place since the emergence of the anatomically modern human species. Except one: imagination. Creative problem solving and pretend play involve the same cognitive mechanisms. And imagination in children serves, according to Carruthers, the function of preparing and practicing our adult capacities for creative problem solving. Sometime between the appearance of the modern human anatomy and c. 40,000 year ago, there was selection for imaginative abilities.[56] What might such selection pressures be? Carruthers leaves this question open, but there are a number of options. We might follow the general strategy of Geoffrey Miller, taking the relevant pressures to be sexual (Miller 2000). The most plausible story here would be that children more practiced in imaginative play would grow up to be better problem solvers, and this quality would be perceived and sought by potential mating partners. Alternatively, we might, following Miller more closely, assume that humans have an innate preference for creative, and thus imaginative, sexual partners. The pressures might instead have been environmental: imaginative play was selected for because it enables more successful problem solving across a range of environmental variables. Or finally,

[56] What has been termed the 'creativity explosion' is typically dated at around 40,000 years ago. Around this time, there was a massive increase in art, technology, and culture: body ornamentation, artistic representation, evidence of religion, extravagant burial practices, and significantly more sophisticated tool-making and hunting techniques (Boyer 1994; Mellars 1973, 1989; Mithen 1996; Stringer and Gamble 1993; White 1982, 1993).

we might fuse the sexual and environmental pressures. Whatever the case, Carruthers' proposals motivate the following: (f) If imagination is an (evolutionary) functional prerequisite for creative thought, then creative thought requires imagination.

We thus have a wide range of independent arguments intended to support the claim that imagination is necessary for creative thought. Now consider the following. Take the antecedent clause for the (first) conditional of each of (a) through (f). Combine these clauses as a disjunction. If this disjunction is true, then the consequent proposition that imagination is necessary for creative thought is true. And if any one of the disjuncts is true, the disjunction is true. Now consider, is it not likely that at least one of these disjuncts is true?

Some will be persuaded by such an argument, some will not. In any case, nothing crucial stands or falls with its success. Even if it fails, we have nonetheless identified a handful of independent reasons for thinking that imagination is essential for creative thought. And we have, at the very least, made a strong case for the link between imagination and creativity. Thus imagination is undoubtedly an important if not essential component in much of human creativity.

5.4 A worry about mundane cognition

Much of the above is an argument that creative thinking requires non-truth bound cognition; it requires cognitive manipulation and thus, we have argued, imagination. We might worry, however, about the scope of this proposal. We might think that the fact that creative thinking requires cognitive manipulation derives simply from the fact that novel cognition requires cognitive manipulation. Any mental state that is novel with respect to

an agent (and we are just talking novelty here: set aside satisfaction of the modal condition) requires cognitive manipulation for the simple reason that the state is psychologically *novel* and thus not a component in the cognitive profile before the present time. But this point generalizes to all novel cognition: any learning, concept acquisition, belief formation, desire formation, new intentions, etc., would require the same. And so we have two problems: One, we have the consequence, following the reasoning offered above, that a great deal of rather mundane cognition requires imagination. Two, the point about creativity and imagination, if it is true, is parasitic on a more general point about novel cognition.

The first problem is no problem. It may seem surprising upon first pass that mundane cognitive acts like forming a belief or desire, acquiring a concept, or learning a simple skill require imaginative activity. However, we should keep in mind a few things. First, imagination varies in richness: we can baldy suppose some proposition or other, richly imagine a proposition and its various entailments, image an entire multi-instrument musical work. To propose that much of cognition requires imagination is only to propose that it involves, at least, the first of these.

The second point extends from here. Consider the formation of a simple mental state, say, the belief that "The Cubs will win the Series." We might form this belief rather immediately, say upon reading the current league standings very late in the baseball season, or over a longer period of time, say only after watching all the Cubs games from April until October. Either way, the belief will involve some consideration of hypothetical circumstances in its formation: what if Kerry Wood's elbow problems return? what if the Cubs hadn't traded Sosa? what if the Cubs drop the first two games to

the Giants? what if it rains during the home games? And so on. Set aside both the truth and the justification for the belief. Its formation, if it involves any of this kind of consideration at all, involves the cognitive manipulation that imagination offers us. The degree of imaginative engagement shifts of course with both the epistemic agent and the context: given certain contexts some epistemic agents let their imaginations rip, as David Lewis puts it (Lewis 1996). Given other contexts, the imaginings are fewer in number and narrower in scope. Some agents let their imaginations rip all the time; some just aren't very imaginative, ever. The point for us is that we should be happy to accept that much of belief-formation involves imagination.

This point generalizes: there is nothing particularly special about belief in this regard. Desires and other propositional attitudes are often the result, in part, of imaginings. Cross-categorical concept application require imagination—this was one of the morals of the Caciarri/Karmiloff-Smith/Ward studies. Learning figurative language, and even much of literal language, depends upon imaginative engagements of various sorts (see discussion of Caciarri et al above, 5.1.1). So in short, we should not take the mundaneness of a state to be an indication that it did not require for its formation the use of imagination.

This pushes us directly into the second problem. Grant that creativity requires imagination. But this is for the trivial reason that creativity involves novel cognition. And if this is right, in establishing a conceptual connection between creative cognition and imagination, we have failed to distinguish it from mundane cognitive acts like learning and forming beliefs and desires. Our cognitivist analysis is thus uninformative with respect to creativity in particular, since our point is really a general one about

novelty. The response to this criticism is two part. First, we must recall that we are offering an analysis of *minimal* creativity. It has already been acknowledged in 1.3.1, 3.3.2 and 3.3.3, that much of mundane cognition will qualify, for *that agent*, as minimally creative. Second, we must recall that such states, like more radical or preconceived instances of creativity, qualify as creative by satisfying not only the novelty condition, but also the agency and modal condition. A cognitive state, no matter if it is a belief, a skill, a desire, or an imagining, must also be the product of agency and previously (nomologically) impossible with respect to the cognitive profile in question; novelty will not suffice. It is this feature of our analysis, carefully argued in chapter 1, which distinguishes creative cognition from mere novelty.

Chapter 6
Learning Without Looking: Incubated Cognition and Creativity

Creativity scares naturalists. Traditionally, this has been the case for a number of reasons, sometimes disparate sometimes connected. One such reason consists in an extreme analysis of creativity that insists upon locating creative thought in unconscious, free-associative thought.

Based largely upon the introspective reports of Hermann von Helmholtz, psychologist Graham Wallas distinguished four stages of creative cognition (Helmholtz 1896; Wallas 1926; these distinctions have also been credited to Poincaré [1902-8]1984, and Hadamard 1954). *Preparation* involves acquisition and application of the relevant skills and knowledge to some problem or task. *Incubation* occurs when conscious attention is diverted away from the problem. The third stage involves a moment or moments of *illumination.* The name here says it all: after incubation creative insight flashes into sight. The fourth stage is one of *verification.* Here the initial insight is subjected to evaluation, criticism, and eventual improvement. This model is fine as far as a superficial description of the phenomenon goes.[57] What it lacks most is expanded treatment of just what goes on in each of the four stages. The second of these is of present interest: just what is incubated cognition? Sections 1 and 2 of this chapter attempt an answer to this question. Section 3 attempts to connect that answer to the discussion of imagination in the previous chapter.

[57] As we will see, it is not great even by this standard, since it is apparently at odds with creative insights that do not involve any incubatory thought, that is, ones that are gotten via or during conscious attention to the problem or task. Moreover, the model seems to posit flash phenomenology for all creative thought, which is surely overstated.

6.1 Incubation and unconscious processing

Brief consideration of the above four stage model informs our first pass on the the term 'incubation.' It involves an input stage, comprising conscious preparation, an output stage, comprising some conscious insight, and a mediating process that is causally affected by and causally affects the first and third stages, respectively. 'Incubation' denotes this mediating period: a period of cognitive activity where an agent ceases to deliberately attend to some problem or task *x*. Incubation could thus involve a shift in attention (away from *x*), decreased attention (e.g. in dreaming or daydreaming), "no" attention (e.g. during deep sleep), a lapse in concentration, or some form of distraction. In any case, an incubatory period with regard to some *x* is a period where *x* is not part of one's *conscious* experience. Incubated cognition thus ostensibly involves some kind of unconscious processing. So our first step is to clarify the nature and plausibility of unconscious cognition.

6.1.1 Spooky beginnings

Arthur Koestler cites a lecture given by Henri Poincaré in 1908, where Poincaré considers the following dilemma (Koestler 1964: 164-5). During creative thought processing, ideas are combined in novel ways, and this combination is performed largely unconsciously, by what Poincaré calls the *subliminal self.* For Poincaré there are only two ways we might think of the unconscious. One, we might think of the unconscious as capable of careful and fine discernment and, importantly, distinctions and combinations that the conscious mind fails to make. This implies that the unconscious mind is superior to the conscious. Poincaré doesn't like the sound of this, and so opts for what he takes to

be the only other option: we should think of the unconscious as an automaton that mechanically runs through various combinations of ideas.

> Figure the future elements of our combinations as something like the hooked atoms of Epicurus. During the complete repose of the mind, these atoms are motionless, they are, so to speak, hooked to the wall. During a period of apparent rest and unconscious work, certain of them are detached from the wall and put in motion. They flash in every direction through the space...as would, for example, a swarm of gnats, or if you prefer a more learned comparison, like the molecules of gas in the kinematic theory of gases. Then their mutual impacts may produce new combinations (cited in Koestler 1964: 164).

Incubated cognition on this account is thus an automatic, free-associative combination of ideas and concepts. At some point, certain of these combinations surface for conscious reflection.

An alternative, and familiar, model of the unconscious comes from Freud. Freud's theory is no doubt a complex and extensive one, but central to his pioneering techniques of psychoanalysis was a robust theory of the unconscious self. According to Freud, we all have unconscious states that are repressed for their being too painful, socially unacceptable, or self-destructive to indulge. These urges and drives manifest themselves in mildly obsessive behavior, verbal and behavioral "slips", and dreams for most. For others they manifest themselves in more consistent disorders like schizophrenia and depression (Freud 1940). Freud famously argued for a three-part structure of the mind, postulating the *id*, the *ego*, and the *super-ego*. The first involves unconscious drives and urges, while the third comprises the social, conventionally acquired, conscience. The ego then, is the conscious self that results from the struggle between the first and the third.[58] Of this power struggle, Freud says the following:

---

[58] In fact, this point is contentious. Some take the ego to be free-floating between conscious, preconscious, and unconcious experience.

> The poor ego has a still harder time of it; it has to serve three harsh masters, and has to do its best to reconcile the claims and demands of all three. These demands are always divergent and often seem quite incompatible; no wonder that the ego so frequently gives way under its task. The three tyrants are the external world, the super-ego and the id. When one watches the efforts of the ego to satisfy them all, or rather, to obey them all simultaneously, one cannot regret having personified the ego, and established it as a separate being. It feels itself hemmed in on three sides and threatened by three kinds of danger, towards which it reacts by developing anxiety when it is too hard pressed. Having originated in the experiences of the perceptual system, it is designed to represent the demands of the external world, but it also wishes to be a loyal servant of the id, to remain upon good terms with the id, to recommend itself to the id as an object, and to draw the id's libido on to itself. In its attempt to mediate between the id and reality, it is often forced to clothe the unconscious commands of the id with its own preconscious rationalisations, to gloss over the conflicts between the id and reality, and with diplomatic dishonesty to display a pretended regard for reality, even when the id persists in being stubborn and uncompromising. On the other hand, its every movement is watched by the severe super-ego, which holds up certain norms of behaviour, without regard to any difficulties coming from the id and the external world; and if these norms are not acted up to, it punishes the ego with the feelings of tension which manifest themselves as a sense of inferiority and guilt (Freud 1940).

Unless you are a neo-Freudian, so much unconscious thinking and strife comes as a surprise.

Considerations of views like the above have no doubt motivated one of two reactions to talk of incubation and unconscious thought in theories of creativity. Many theorists of creativity have endorsed this basic framework and put heavy emphasis on the incubation stage. According to these views, creativity requires incubation: a cognitive system can produce a creative thought *t* only if some of the processing that enabled *t* is incubated. Call any such theory *incubation essentialism* (Koestler 1964; Martindale 1977, 1981, 1990, 1995, 1999; Mendelsohn 1976). These views have often been coupled with mystical, neo-Freudian, or supernatural treatments of creativity.

These considerations motivate an argument, call it the *Argument from incubation essentialism*, which captures the cause for the naturalist's fright.

(1) Creativity entails incubated cognition.
(2) Incubated cognition entails a subliminal self (understood (a) as a Freudian unconscious or (b) as an automaton.
(3) (a) and (b) are out of the naturalist's purview.
(C) Incubated cognition is out of the naturalist's purview.
(C_2) Creativity is out of the naturalist's purview.

The first conclusion (from (2) and (3)) encourages *incubation phobia*: if creative cognition requires *that kind* of cognition, then there is little if anything that scientifically minded philosophy can say about it. There is ample room for response here: one might deny (3) by offering a naturalistic model of (a) or (b). One might deny (2) by modeling incubated cognition in a way that requires nothing like a subliminal self (this is to argue through Poincaré's dilemma). Either move is sufficient to bar the inference to (C). The model offered below is probably best understood as making the second move: a denial of (2).

The second conclusion (C_2) follows from (1)-(C). Thus one simple way to bar this inference, in addition to those offered above, is to weaken (1) for something like (1′) Creativity involves incubated cognition; this is simply to deny incubation essentialism. But a conclusion like (C_2) is supported in other ways, for example by acknowledging some other purportedly spooky features of creativity.

Here are two such features (also discussed, among other places, in 1.1.2, 1.2.1, 1.3.4). Creativity requires, as a conceptual point, genuine novelty. To be creative, an *F* must be new with respect to some system: social, cognitive, environmental, biological. Genuine novelty implies *ex nihilism*: creative *F*'s emerge from nowhere. And science has got nothing on nowhere. Second, creative ideas often come to their bearers unbidden

like bumps on the head. We describe such ideas as ones that "just happen" or "just come to us" unwilled in flashes or bursts of insight. This *flash phenomenology* mocks naturalism: unwilled creative insight inspires inspirationalism which, argued since Plato, is outside of the naturalist's purview. These considerations conjoined with arguments like the one offered above have proven sufficient to keep the naturalist away, for the most part, until now.[59]

But these two extremes can be avoided. We can split the difference between incubation essentialism and incubation phobia, maintaining unconscious incubated cognition explained naturalistically. And, as a bonus, we will exorcise some additional spooks in the process. A few distinctions provide a better start.

6.1.2 Clarifications and distinctions

We should first distinguish the unconscious *self* or *person* from *unconscious mental processing*. Freud clearly was more concerned with the former: the id is commonly understood as the deeply buried self, manifesting only in obsessive-compulsive behavior, dreams, and other semi-conscious states. The same is true for Poincaré, whose dilemma forces a choice between an intelligent and deliberating unconscious self or an automatic unconscious self. Both the Freudian and Poincaréan views entail unconscious processing, but the entailment does not run the other way: we can posit unconscious mental processing without positing any kind of unconscious self or person. Use of 'unconscious' will thus be understood as reference to mental states or

[59] There are exceptions: Boden 1994, 2004; Dartnall 2002; Finke et al. 1992; Gabora 2000, 2002; Simonton 1999; Smith et al. 1995; Sternberg 1999; Weisberg 1986, 1995, 1999.

processes and not to selves. Considerations from cognitive neuropsychology, among other fields, will be offered below to this end.

We should also distinguish what Ned Block calls *access consciousness* from *phenomenal consciousness* (Block 1995). The notion of phenomenal consciousness derives from Thomas Nagel: a system or organism is phenomenally conscious if and only if there is something it is like to be that system or organism. A state or process of an organism is phenomenally conscious if and only if there is something it is like to be in that state or process (Nagel 1979). Access consciousness is more tricky, but the following gloss is sufficient for our purposes. A state or process is access conscious if it is available for judgment and inference—it can be attended to, can be the subject for belief and other states, can be evaluated, can figure into reasoning, etc. Such states and processes are thus *potentially* the object of current attention, but need not *actually* be attended to. (See Davies 1995: 359-364 for further nuances on the notion of access consciousness.) Our interest in incubated cognition is thus in states and processes that are *not* phenomenally conscious.[60] Whether or not incubation is access conscious will remain an open question: some of the processes we will consider appear to be access conscious in the sense we have clarified, others do not.

Commitment to incubated cognition thus only involves a commitment to non-phenomenally conscious mental states or processing. This is consistent with deflationary accounts of consciousness like Daniel Dennett's. According to Dennett, we can think of consciousness as involving multiple drafts of stimulus interpretation.

---

[60] What would it mean for a *process* to be phenomenally conscious? A process is phenomenally conscious if the states (and their relations) that comprise the process are phenomenally conscious. This admits of more and less precise formulations: perhaps all of the states must be conscious but not all of the relations, perhaps a certain number of states and relations is sufficient, perhaps it is only the states and not the relations that matter, etc. For our purposes, the first pass gloss will do.

> [A]ll varieties of thought or mental activity—are accomplished in the brain by parallel, multitrack processes of interpretation and elaboration of sensory inputs. Information entering the nervous system is under continuous "editorial revision."
>
> These editorial processes occur over large fractions of a second, during which time various additions, incorporations, amendations, and overwritings of content can occur, in various orders. We don't directly experience what happens on our retinas, in our ears, on the surface of our skin. What we actually experience is a product of many processes of interpretation—editorial processes, in effect.
>
> [O]nce a particular "observation" of some feature has been made, by a specialized, localized portion of the brain, the information content thus fixed does not have to be sent somewhere else to be rediscriminated by some "master" discriminator (Dennett 1991: 111-3).

The implication is that there is no threshold or "finish line" for consciousness: some content discriminations appear in our stream of conscious experience and some do not.[61] The important point for Dennett is that there is no "moment of consciousness" and thus the conscious/unconscious distinction is, at best, an arbitrary one. Our commitment to unconscious cognition is consistent with though non-commital to Dennett's metaphysics. Unconscious processing, as we are understanding it, is non-phenomenally conscious but causally efficacious cognition: states that contribute to the overall consciousness of the system in some non-trivial way but which do not appear as part of the "stream" of

---

[61] Freud says something in a similar vein, "Some processes become conscious easily; they may then cease to be conscious, but can become conscious once more without any trouble: as people say, they can be reproduced or remembered. This reminds us that consciousness is in general a very highly fugitive condition. What is conscious is conscious only for a moment. If our perceptions do not confirm this, the contradiction is merely an apparent one. It is explained by the fact that the stimuli of perception can persist for some time so that in the course of it the perception of them can be repeated. The whole position can be clearly seen from the conscious perception of our intellective processes; it is true that these may persist, but they may just as easily pass in a flash. Everything unconscious that behaves in this way, that can easily exchange the unconscious condition for the conscious one, is therefore better described as "capable of entering consciousness," or as preconscious. Experience has taught us that there are hardly any mental processes, even of the most complicated kind, which cannot on occasion remain preconscious, although as a rule they press forward, as we say, into consciousness. There are other mental processes or mental material which have no such easy access to consciousness, but which must be inferred, discovered, and translated into conscious form in the manner that has been described. It is for such material that we reserve the name of the unconscious proper" (Freud 1940).

conscious experience. In Dennett's terms, these are just the states that serve in the process of editorial revision, but do not themselves register in experience.

6.1.3 Empirical foundations: A test for scientific plausibility

We can thus settle on the following simple thesis: some information processing in a cognitive system occurs at an unconscious level. These processes are not part of phenomenal conscious experience, but nonetheless may causally affect such experience. Call this the *unconscious cognitive processing thesis*. Notice that this is silent with regard to both the question of access consciousness and the details of the processes involved. Having given then a rough conceptual understanding of unconscious processing, the next step is to check the empirical support for the thesis.

Disputes over scientific plausibility are aplenty. Nonetheless, there are two commonly accepted conditions useful for testing such plausibility. Karl Popper identified *falsifiability* as requisite for theoretical legitimacy: serious scientific consideration is due only to theories with identifiable falsifying conditions—possible outcomes that would provide evidence against the theory (Popper 1934).[62] A theory that has passed the tests that would falsify it is said to be *corroborated*: the theory has thus survived refutation (Popper 1934; Putnam 1979). What then would falsify our unconscious cognitive processing thesis, and do the empirical results corroborate the thesis?[63]

---

[62] To be more precise, theories that fail to satisfy this criterion are, by Popper's lights, simply not scientific and thus not even candidates for scientific plausibility.

[63] Owen Flanagan uses just this method of testing Freudian psychoanalytic theory broadly understood—taking unconscious mental processes as the core of such theories. He suggests the results of hypnotic suggestion as a potential falsifier which plausibly corroborates the theory through a variety of experiments (Flanagan 1991: 74-8).

Some tests: We can look to cognitive neuropsychology for a number of empirical tests.

(a) The phenomenon known as *blindsight* was first observed by George Riddoch in 1917 and has since received extensive study (Riddoch 1917; Barbur, Ruddock, and Waterfield 1980; Bridgeman and Staggs 1982; Weiskrantz 1986). Patients with lesions to the primary visual cortex suffering partial or total blindness can nonetheless sometimes detect and identify motion and other features of stimuli in the otherwise blind portions of the visual field. These abilities are typically identified using forced-choice situations: the patients are forced to make guesses about various features of objects in their blind visual fields (of which they are thus unaware). The degree and nature of these abilities vary between patients. For example, one patient with an entirely destroyed right primary visual cortex who thus suffered total blindness in his left visual field, nonetheless maintained the ability to identify the motion, location, and orientation of stimuli, and to discriminate shapes like 'X' from 'O' in the left field (Poppel et al. 1973; Weiskrantz 1986; Weiskrantz et al. 1974). Color vision discriminations are sometimes preserved (Brent et al. 1994). And, remarkably, presentation of certain words in the blind field have been shown to consistently bias interpretation of later presented ambiguous words (Marcel 1998; see Farah 2001 for further discussion of these and other studies).

(b) Patients suffering from various forms of *associative visual agnosia* have impaired abilities of object recognition despite healthy perceptual abilities and overall general intelligence. *Prosopagnosia*—an impairment of recognition of familiar or previously learned faces—has been studied extensively. Covert recognition has been identified in such patients by a number of researchers. de Haan and others showed that

prosopagnosics are better at learning face-name pairs for familiar faces (which they fail to overtly recognize) than unfamiliar ones. They also found that processing time in reaction tasks co-varied with the familiarity of the faces used as stimuli (de Haan et al 1987a). Finally, de Haan et al found that photographs of faces could evoke associations with those faces, despite the fact that the patients would deny any such knowledge when explicitly questioned (de Haan et al 1997b; see also de Haan 2001 and Farah 1990, 2001: 167-72.)

(c) Patients suffering *apperceptive agonosia* fail to consciously recognize and discriminate object features like shape, size, and orientation in virtue of a failure in the perceptual system. These patients—as contrasted with associative agnosics who fail to retrieve stored information or knowledge about objects— fail to properly perceive objects. A number of studies have, however, identified preservation of *motion-guiding vision* in such patients.[64] One such patient, DF, was able to properly shape the grip of her hand to grasp objects in spite of her inability to describe the various features of the same objects. In fact, her motor performance was wholly intact: she was fully capable of grasping a pencil or doorknob, or putting a card through a slot (Milner et al 1991; Goodale et al 1991; Milner and Goodale 1995). These results have motivated dual systems models of vision, comprising a system for motor guidance and a system for descriptive vision.

(d) *Pure alexia* is a disorder where patients, despite normal abilities to write and speak language, require abnormal amounts of time to read language or lack the ability altogether (Beringer and Stein 1930). A number of researchers have identified "implicit

[64] For philosophical discussion of motion-guiding vision, see Matthen 2005: 293-324, and Carruthers (in preparation) 2.2.

reading" abilities in such patients. One subject was able to perform lexical decisions—distinguishing words from non-words—with considerable accuracy when presented with letter strings for two seconds (which was not enough time for him to make explicit discriminations) (Shallice and Saffran 1986). More strikingly, subjects in another set of studies performed better on implicit reading tasks when the stimulus was even shorter, at 250 milliseconds (Coslett and Saffran 1989, 1992; see also Farah 2001: 176-8.)

The studies discussed in (a) through (d) all concern patients who have lost the ability to consciously process a particular kind of information or complete a particular kind of task. To simplify: all of these patients lack the ability to consciously do *x*. An outright failure to do *x*, whatever it should be, would constitute evidence against the unconscious cognitive processing thesis. Enough of such evidence, we could plausibly suppose, would in fact falsify our thesis. These considerations are sufficient to show that the thesis meets our first condition: we have identified conditions that would falsify it.

The results: Do any of these falsifying outcomes obtain? The obvious answer is no. In each class of studies, the patients have retained some ability to do *x*: patients with blindsight still make some visual discriminations; prosopagnosics retain some ability to recognize faces; apperceptive agnosics retain normal motion-guiding vision; pure alexics can perform lexical identifications. This data corroborates our thesis: complete failure by any of these patients to perform the relevant tasks or process the relevant information would count against the thesis, but in no case does such failure occur. We thus have good reason to infer that the unconscious cognitive processing thesis passes the corroboration condition.

But, one might worry, surely these tests are *ad hoc*: you have taken empirical data and then set up the tests for falsification afterward. This is true, the test formulation did in fact move from the data backwards. But the procedure is in no way unprincipled. The empirical results obtain independently, and the fundamental tests for scientific plausibility do as well. And the fact of the matter is that that data fails to falsify the thesis in question. So there should be no serious concern for ad hockery.

We might also set aside concerns about falsification and corroboration, and use these studies in a more direct manner. So while Popper's claim is that we should ask not what would prove a theory but what would disprove it, it is still reasonable to consider direct evidence *for* a theory. Surely the fact that these patients retain the ability to do *x* provides evidence for unconscious processing. We reach this conclusion by a simple inference to the best explanation. We know that agnosics, for example, do not process the information consciously—a prosopagnosic will consistently deny to recognize previously known faces. We have then only to ask how they *do* process it. The only viable answers support the unconscious cognitive processing thesis.

One final concern worthy of brief mention goes like this. The data in question provide evidence that we process sensory information unconsciously. But who in the 21st century would doubt that? It is no surprise that much of the sensory processing responsible for my visual experience is not actually experienced by me. What we need is evidence for unconscious *cognitive* processing. And unless these retained abilities can be explained in terms of cognitive processing, then the data do no relevant work. This forces an important but unharmful qualification: some explanations of these phenomena are in terms of perceptual, sub-cognitive processing. But most of them are not. Much of

the processing is cognitive if any processing is: describable features of objects, learned associations with faces, linguistic discriminations. So even if some of the studies in question only evidence sub-cognitive processes of the perceptual or nervous system, there are enough of them that clearly evidence something at the cognitive level. That's all we need.

6.1.2 and 6.1.3 should suffice to clarify the nature of unconscious cognitive processing generally, and to lend scientific plausibility to the positing of such processes. Also recall the intuitive observations motivating the claims and theories of Helmholtz, Wallas, Poincaré, and Hadamard with which we started the chapter, and the introspective reports of Archimedes, Kekulé, and Coleridge, among others.[65] All parties observe some kind of unconscious processing as (at least potentially) part of creative cognition. This processing takes place during a period of incubation: good things apparently happen when we abandon conscious attention to a problem. The inference we should draw is that work on the problem does not stop during this incubatory period: incubated cognizing of x involves *cognizing* of x, albeit at an unconscious level. Grant then *that* we unconsciously cognitively process information, the question now becomes *how* we do it. And an answer to the second will lend additional support to a positive answer to the first.

6.2 Incubation, attention, and learning

In a series of experiments, psychologists Steven Smith and Steven Blankenship studied the phenomenon of incubation, focusing on what they call *incubation effects*:

---

[65] Archimedes was reportedly having a bath when his insight into the proper measurement of the volumes of irregular shapes came to him. The famous report of Kekulé was mentioned in 1.1.2. Coleridge was reportedly enjoying an opium holiday when the inspiration for his *Kubla Khan* came (see Boden 2004: 25-8).

instances where subjects have greater success solving an initially unsolved problem after setting it aside for a period of time. Their project had two primary motivations. First: show that incubation and incubation effects do occur during problem solving. Prior experimental research had mostly failed to establish this thesis (Olton 1979). Second: establish a reliable methodology for testing incubation (Smith and Blankenship 1989, 1991; see also Smith 1995).

The 1989 study was largely successful in meeting the first motivation, and constructed a methodology that would be maintained with success in the later study. Smith and Blankenship begin with the hypothesis that failed problem solving often depends upon *fixation*: subjects retrieve or construct incorrect strategies for and solutions to the problem and then suffer a mental block from the correct one/s.[66] The fix for fixation? Forget it. Smith and Blankenship propose and test the *forgetting-fixation* hypothesis which suggests that overcoming fixation is crucial to making unsolved problems solvable. After initial presentation of a problem, they induce fixation in subjects by priming them with incorrect solutions. The subjects are then either retested immediately or after a period of time. The second group, those who presumably had time to forget the fixated (incorrect) solutions, did consistently better than the first group upon retesting. (Also worth mentioning is that the second group were given "filler tasks" during the incubatory period. These tasks were reportedly very difficult and were stressed as no less important to the subjects. So even during the incubatory period, the subjects were still cognitively engaged.) These studies are instructive: they identify attention as an important cognitive dimension in considerations of incubation. They also lend some empirical support to the merely introspective reports with which we began.

[66] Fixation was first studied by Woodworth and Schlosberg (1954).

6.2.1 Attention, arousal, and association

Recall Martindale's cortical arousal theory, discussed in 2.2 and 5.1.2. We can use some of his points of emphasis as a starting off point for our own model. Central to Martindale's theory are two important insights borrowed from psychologists who worked on creativity before him. Gerald Mendelsohn stressed the importance of attentional capacity to creative cognition (Mendelsohn 1976). His innovation was to recognize that it is the range of attention rather than the strength or duration that is crucial. Mendelsohn's method was to use the Remote Associations Tests (RAT) frequently used to test creative problem solving capacity (Mednick 1962). The task in these tests is to identify an associative (semantic) link between distinct words of "distant associative clusters." For example, the solution to 'apple', 'family' and 'house,' is 'tree.' Common explanations of the varying success at such tasks appeal to vocabulary, general intelligence, and ability to maintain focused concentration on the task. Mendelsohn does not deny that these elements may be important, but isolates the importance of the range of attention. "The ability to receive and store in accessible form a broad range of information from the environment would serve to increase the range of elements, including unusual, peripheral, or incidental elements, which could be evoked during the process of thought. The ability to maintain several streams of cognitive activity simultaneously, i.e., in parallel, increases the likelihood that otherwise separate sequences of thought will be brought into contiguity and combined" (Mendelsohn 1976: 363). The issue then, according to Mendelsohn, is how many associations are activated, not how closely we attend to one or another. Moreover, quantity of associations stored is not as

important as the quantity of associations activated. This feature depends both upon the level of arousal and the associative connections between concepts.

The latter notion, called *associative hierarchies*, was at the center of S.A. Mednick's theory of creativity, and has been adopted by a number of other theorists (Mednick 1962; Gabora 2000, 2002; Martindale 1977, 1981, 1995, 1999; Mendelsohn 1976). The notion goes further back to the *habit-family hierarchies* of psychologist Clark Hull (Hull 1943). The basic idea is this. When presented with a stimulus, say, a word, persons respond in one of two ways. The stimulus will, for some persons, strongly activate a small number of associated concepts. These persons have *steep* associative hierarchies and will typically give predictable responses to the stimulus. For other persons, the stimulus will activate a larger number of associated concepts, but each to a lesser degree. Such individuals have *flat* associative hierarchies. Flat associative hierarchies yield unexpected responses. Mednick's thought was that creative people typically have flat associative hierarchies and uncreative people, steep ones (Mednick 1962). Such a sharp demarcation is artificial and our interest is not in modeling particularly creative persons but creativity generally; we can nonetheless learn from Mednick's basic insight.

Following Mednick, Martindale emphasizes the importance of associative hierarchies. Following Mendelsohn, Martindale emphasizes the importance of the range of attention focus. Decreased cortical arousal, according to the theory, results in *defocused* attention (Martindale 1995). When a person enters this stage of cognitive processing (what Martindale, following Kris 1952, also calls *primary* processing), the range of activation increases. And thus the range of concepts that may be *attended to*

increases. (And here, as above, 'range' refers to quantity or density of networks, not to the spatial area of the brain.) Martindale sometimes puts the point in the terms of neural networks: defocused attention results in a greater activation of nodes in a network, though the strength of the activation of each node is relatively equal. By contrast, when attention is sharply focused, a few nodes are highly activated (Martindale 1995; see also Rumelhart et al 1986). In other words, decreased but broadly spread attention increases the concepts that figure into the cognitive processing of a stimulus or task.[67] This condition, coupled with flat associative hierarchies, dramatically increases the chances that the output will be unusual and, possibly, creative.

Following Mendelsohn and Mednick, Martindale uses these points to infer a sharp distinction between creative and uncreative people. But we can recognize the importance of attention range, quantity of conceptual associations, and degree of cortical arousal without making such a sharp demarcation (more on this in 6.2.5). The question, if we want to distance ourselves from a theory of creative *persons* and instead offer a theory of minimally creative thought is: why should we think that an average human being is capable of this?

6.2.2 Hebbian cell assemblies and neural plasticity

Donald Hebb famously coined the term 'cell assembly' for clusters of neurons that constitute a loop or circuit for continued (post-stimulus) neural activity. These assemblies form as a result of the synchrony and proximity of the firing of individual cells. Hebb describes their formation as follows: "any two cells or systems of cells that

[67] *Spreading-activation* theories of various cognitive capacities—for example, semantic processing and memory—are closely related. See Collins and Loftus (1975) on semantic processing, and Anderson (1983) and Gabora (2002) on memory.

are repeatedly active at the same time will tend to become 'associated', so that activity in one facilitates activity in the other" (Hebb 1949: 70). These assemblies are the basic elements for Hebb's broader theory of mind, crucial for learning and memory.

Hebbian learning takes conceptual root in the notion of cell assemblies. "Let us assume then that the persistence or repetition of a reverberatory activity (or 'trace') tends to induce lasting cellular changes that add to its stability. The assumption can be precisely stated as follows: *When an axon of cell* A *is near enough to excite a cell* B *and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that* A's *efficiency, as one of the cells firing* B, *is increased*" (Hebb 1949: 62). These *reverberations* and the changes they cause ultimately result in the formation of closed, semi-autonomous systems of neurons—cell assemblies—which can activate in the absence of the initiating stimulus. Neural structures thus change with the learning of new concepts, skills, and information.

Memory recall, on the Hebbian model, involves activating neurons in specific patterns along the suggested pathways. This has prompted many psychologists to endorse a model of *distributed* memory (or cognitive processing more generally). On such models, memories are not found at particular locations in the brain, but rather via particular patterns of neural activation (Hinton and Anderson 1981; Hinton et al. 1986; Kanerva 1988; Willshaw 1981; see also Gabora 2002). The ability to recall a memory thus depends upon the connections (and their strengths) between neurons and the nature of the present stimulus—that is, whether the present stimulus evokes a similar pattern of neural activation. The first—the strengths of the connections constituting the cell assembly—are determined by the principles of Hebbian learning briefly articulated

above. Finally, neurons activated during memory recall will of course connect with other cells and assemblies: we thus have the experience of one memory leading to another.

There have been numerous variations on Hebb's basic ideas in the 50 plus years since their initial publication. The conceptual fundamentals are still widely accepted by a range of physiologists, psychologists, and computational scientists (Amit 1995; Braitenberg 1989; Kaplan et al. 1991; Miller and Wickens 1991; Mishkin 1993; Palm 1982; Sakurai 1999; Spatz 1996). Many have been critical of, among other things, the notion of reverberation (Milner 1957, 1999), of the capacities for discrimination of associations, memories, and concepts given the closed nature of Hebbian assemblies (Hopfield and Tank 1986), and of the apparent lack of constraints on the growth of cell assemblies and their threshold for activation (Milner 1999). Others have opted for distinct, albeit importantly similar, concepts to serve the same (or nearly the same) role: e.g. 'mneme' (Semon 1921), 'engram' (Lashley 1950; Milner 1999), and most recently, 'neural network' (Marr 1969, 1976, 1982; Hinton and Anderson 1981; Sejnowski 1986). Nonetheless, these criticisms and amendments maintain the core features of Hebb's proposals: there are vast networks of connections between the brain's individual neural cells, these connections are essential to learning, memory, and other cognitive functions, and they are continually changed by our interactions with the world.[68]

This last feature of the Hebbian legacy is most important for our purposes: we change our brain by interacting with the world. More specifically, continued attention to

[68] Hebbian terminology will be used throughout. In some ways this is just useful shorthand, since there are several contentious issues that attach to the respective terms and many terminological variations on offer. An analysis of the logical space, how to navigate it, and who occupies what space, is orthogonal to our concerns. And, as already suggested, most of the basics of Hebbian theory are widely maintained today (see Spatz 1996; Orbach 1998; Fentress 1999; Sakurai 1999; Seung 2000). It will thus be assumed that Hebb's basic substantive framework adequately models the network structure of the brain (at least at the level of abstraction needed here).

a problem, what some have called *cerebral effort*, causes changes in the networking of the brain's cortex (Donald 2001: 175-8). In Hebbian terms, continued attention to a stimulus strengthens the connections between neurons in existing cell assemblies and/or creates new connections and thus new cell assemblies. In popular terms, the brain's hardware is effectively re-wired in important ways in virtue of thinking about the same task or problem. Such functional re-wiring may be temporary or may have more lasting effects. Thus attention not only affects neural activation but also neural structure.

This consists with a more general point about *neural plasticity*. Current neuroscience models the brain as an organic structure—rather than as a rigid "knowledge-independent, hardware construct"—stable in genetic material but constantly undergoing functional change and development in neural networking in response to external stimuli (Young 1951; Rosenblatt 1961; Von der Malsburg 1973; Pettigrew 1974; Changeux and Dauchin 1976). The term 'plasticity', originally employed to describe the behavioral recoveries of patients suffering various forms of brain damage, has been adopted to describe this developmental feature of the brain. Some have taken the degree of plasticity to its extreme. Neural Darwinists argue that:

> The brain is part of an organism belonging to a species that has evolved (and is still evolving) in geological time-scales according to Darwinian mechanisms at the level of the genome. But the complexity of the brain is such that it may itself be considered as a system evolving *within* the organism with, at least, two distinct time scales: that of embryonic and postnatal development for the process of organizing neuronal somas and their connectivity networks, and that of psychological times for the storage, retrieval and chaining of mental objects and for their assembly into higher-order motor programs, behavioral strategies and schemas. The extension of *selectionist* mechanisms to all these levels breaks down the rigidity of the strictly nativist or Cartesian schemes by introducing, at each level, a degree of freedom linked with the production of variations (Changeux and Dehaene 1989: 100; see also Edelman 1987; Calvin 1989).

So the use of 'plasticity' is broad to say the least. We will use it simply to refer to the neurophysiological capacity for functional arrangement and re-arrangement. Attention thus changes not only *neural activation* but *neural structure*, and it does the second by doing the first. We have now to ask how attention can effect such change.

6.2.3 Attention and automaticity

"Everyone knows what attention is. It is the taking possession by the mind in clear and vivid form of one of what seem several simultaneous objects or trains of thought" (James 1892). James was right: attention is a perfectly understandable folk concept. However, to understand the connection between attention and neural plasticity, we need to think of the former in terms of its neural correlates. Most neuroscientific research on attention focuses on object-based perceptual attention, e.g. on perceptual *alerting*, *orienting*, and *feature integration* (see Posner and Bourke 1996; Umiltà 2001). Our interest is obviously in something broader, not limited to perceptual experience: we are concerned with cognitive attention or what some call *executive attention*, which is often inclusive of but not exclusive to perceptual attention (Norman and Shallice 1986). Following James, we can attend to objects external or internal, to a moving figure in the visual field, a train of thoughts, or a cognitive challenge. According to neurobiologist Peter Milner, attention to some stimulus *o* (where *o* could be an external stimulus or a mental state) correlates with the facilitation of the neurons or neural networks normally excited by *o*'s or *o*-like stimuli (Milner 1999: 33). Say I am attending to a thought about cherries. Attention to this thought could, most simply, be prompted by a perceptual experience as of cherries. Or, it could be prompted by some other related stimuli (small

fruit-bearing trees, home-baked pie, a terrible song by a terrible heavy metal band in the early 90s) or some other related mental state (an intention to buy cherries at the market). Any of these cherry-like stimuli could excite the relevant pattern of neural activity. This reveals the connection between plasticity and attention: attention involves continuous neural activation which strengthens synaptic connections and thus contributes to the continued shaping of the brain. Thus neural plasticity allows that cognitive attention, whatever the objects of that attention should be, will contribute to the working structure of the brain, since attention correlates with the kind of activation that determines functional networking.

This motivates a simple hypothesis—that attending to and performing cognitive tasks affects neural networking—confirmed by a variety of experimental research (Posner et al. 1997; Posner and Raichle 1994; see also Karni et al. 1995; Nudo et al. 1996). They focus on the acquisition and execution of mundane but high-level skills like reading, arithmetic, and object recognition. Their work provides behavioral data for the proposals of Hebb, Changeux, Milner and others. Using fMRI and PET scanning and imaging techniques, this research confirms a variety of changes in neural structure and activity corresponding both to the practice and eventual acquisition of cognitive skills, and to the attention and effort required for these processes. What kinds of changes are we talking about?

We can look to William James once more. "But actions originally prompted by conscious intelligence may grow so automatic by dint of habit as to be apparently unconsciously performed. . . . Shall the study of such machine-like yet purposive acts as these be included in Psychology?" (James 1890: 6-7). The answer, at least if the question

is descriptive, is yes. A behavioral consequence implied by Hebbian learning and born out by the behavioral studies of Posner and the various others cited shows us why. With continued effort on and attention to a problem—what we would more commonly call 'practice' when referring to skill acquisition—portions of the task or skill become, as it were, internalized. We begin to perform segments of the task automatically. This *automaticity* results, partly, from the strengthening and expanding of neural networks, changes in Hebbian cell assemblies. These automatic processes are typically understood as involuntary and unconscious (Kahneman and Treisman 1984; Posner 1978).

Thus what begins as an activity involving highly focused executive attention, becomes one that is (partially) automatic. Learning another language or a musical instrument, for example, like just about any cognitive task, will begin with conscious attention to each stage of the task. One has to consciously distinguish tenses in irregular verb conjugation or the finger placement on a guitar for an A versus an $A^{\#}$. With time, however, conscious control makes way for automatic performance. There is an intuitive evolutionary reason for this: if we didn't internalize elements of a task, we wouldn't perform many of them. "If an act became no easier after being done several times, if the careful direction of consciousness were necessary to its accomplishments in each occasion, it is evident that the whole activity of a lifetime might be confined to one or two deeds" (Maudsley 1876). In fact, any lifetime spent thus would be a short lifetime indeed and any species with such limitations would not survive long. The move towards automaticity is thus a move towards efficiency.

This efficiency is largely enabled by a reduction in the potential for interference in processing. Conscious attention requires, well, attention: we must focus our efforts on

one stimulus or set of stimuli and keep it there. Automatic processing, by contrast, is less prone to interference. This is for the simple reason that it involves fewer regions of the brain; it involves a decrease in cortical arousal. This is a fact observed by neuroscientists since the early part of the last century. The overall electrical brain activity in organisms is considerably less when the stimulus is conditioned versus when it is novel. Put another way, the area of cortical arousal decreases when experience with an object or task increases (Durup and Fessard 1935; more recently, see John and Killiam 1959; Pigarev et al. 1997). These results are observable in human brains. Using PET imaging, the brains of children before and after learning computer games were compared. After just a few weeks of practice, the area of cortical surface arousal decreased in spite of a sevenfold improvement. In fact, the study showed that the greater the improvement, the greater the decrease in activation area (Haier et al. 1992; for similar studies see Petersen et al. 1998; Karni et al. 1995; Buckner et al. 1995; Desimone 1996). So the more we practice doing something, the less of the brain we will use in doing it. This cuts down on interference, since less irrelevant information will cloud performance of the task (Edelman and Tononi 2001: 58-61). Subjectively, a task feels much easier, if not effortless, when we no longer must attend to, for example, the difference between certain verb tenses or guitar chords.

These proposals are consistent, we must be clear, with Hebbian learning. Even though the regions of the activated cortex become fewer—the activation *area* decreases—with practice and automaticity, the connectivity in those regions may continue to increase. Decreased cortical arousal does not entail a decrease in the formation and augmentation of cell assemblies. In fact, it is likely that just the opposite is true. Since the activation is more localized, the likelihood of proximal neurons firing

synchronously is increased, and with it the likelihood of the creation or strengthening of a synapse between those cells. Edelman and Tononi offer the following metaphor. "It is as if, at first, an initially distributed and large set of cortical specialists meets to try to address a task. Soon they reach a consensus about who among them is best qualified to deal with it, and a task force is chosen. Subsequently, the task force recruits the help of a local, smaller group to perform the task rapidly and flawlessly" (Edelman and Tononi 2001: 61). A shift from conscious attention to partially automatic performance thus increases the efficiency with which we can perform cognitive tasks. It accomplishes this *both* by decreasing the area of activation and increasing the networking complexity in the brain.

6.2.4 Back to incubation

We see how all of this speaks to incubation, and creativity more generally, by considering the subjective consequences of plasticity and attention. Automaticity, as mentioned above, makes a cognitive task easier to perform. This is why experts don't just make it look easier, it *is* easier. Second, automaticity frees up cognitive resources to take on other parts of the task. If a task requires a particularly creative solution, then the more work done automatically the better, since we can continue to practice not-yet-mastered components of the task and ultimately secure a solution or complete the task. Finally, it contributes to that sense that "your brain is working for you" and you aren't working it. Initially, we might be consciously aware of various associations or memories relevant to the task at hand, and then move on to others. But the relevant associations, correlating with cell assemblies, remain active. This activation can contribute to

additional connections, or strengthening thereof, between assemblies and other cells. Given the right stimulus, these new or newly strengthened assemblies may be activated (or re-activated) while one is attending to some other feature of the task, some other cluster of associations. Sometimes the associations that surface are creative or lead to others that are. All of this then, is consistent with and (at least partially) explanatory of the introspective reports with which we began.

But now we incur a worry. We have the makings for a model of incubated cognition, but they underdetermine just how the model should look in the following way. We might opt for (at least) either or the following two options.

(IS) Incubatory solution thesis:

The incubation stage is a stage of lessened or weakened attention to some elements of a task or problem *x*. During this period, activation and strengthening of cell assemblies continues (after the initiating stimulus.) Some of this activity results in a solution to the problem (or something near it). When one returns conscious attention to *x*, the original pattern of cell activation occurs plus activation of the newly formed or newly strengthened ones.

From a subjective point of view, a solution (new association(s)) just comes to us when we return to the problem. (IS) implies that (creative) solution occurs during incubated cognitive processing.

Alternatively, we might opt for the:

(IP) Incubatory preparation thesis:

The incubation stage is a stage of lessened or weakened attention to some elements of a task or problem *x*. During this period, activation and strengthening

of cell assemblies continues (after the initiating stimulus). (Note that this is the same as (IS) up to this point). During the incubated period, cognitive effort can be directed elsewhere. In the meantime, much of the work is "done for you" so that when conscious attention is returned to *x*, new or newly strengthened associative connections have been formed. Some of these associations prove relevant to *x*. With (post-incubatory) attention paid to *x*, including to the newly developed or strengthened associations (i.e. we keep at the problem), a solution may be secured.

From a subjective point of view, a solution comes much easier when we return to problem since we are much better prepared to solve it. (IP) implies that (creative) solution is enabled by the preparatory work that occurs during the incubatory period.

Again, the data and theory we have considered supports both theses, how then do we choose? We don't have to. We have provided a conceptual and neuropsychological basis for *incubated cognitive processing*, and a choice between (IS) or (IP) makes little difference to this fact. Moreover, the two theses are not exclusive: it is likely that sometimes (IS) is a true description of how incubated cognitive processing yields an *incubation effect*, and sometimes (IP) is the true description.

We have thus sketched a model for incubated cognition. We have identified the importance of attention and of unconscious, automatic processing. Using Hebbian theory, coupled with empirical studies from the neurosciences, we have provided a neuropsychological basis for initial conceptualization. In brief, the explanation goes as follows. Attention to a stimulus increases the number and density of connections between neural cells and assemblies. Continued attention to that stimulus (i.e. practice)

decreases the area of cortical arousal, resulting in some degree of automatic information processing. At this point, attention is diverted and post-stimulus activation in these assemblies continues. This is *incubated cognitive processing*. This processing *sometimes* enables or results in useful or novel, perhaps even creative associations. These are *incubation effects*.

Note then that we are distinguishing incubated cognitive processing from incubation effects. The first, though it is unclear precisely where to demarcate it, is in line with our initial characterization: it is a period where conscious attention is removed from some stimulus, but unconscious cognitive processing continues with regard that stimulus. An incubation effect is a conscious mental upshot of that period: a solution, insight, or novel thought. Thus an incubation effect entails, by definition, that incubated cognitive processing took place. But the entailment does not run the other way: incubated cognitive processing could take place and, for a variety of reasons, no conscious upshot surfaces. Note also that an incubation effect needn't be novel and, *a fortiori*, needn't be creative. One could have a thought which resulted from incubation, but which was not, as a matter of fact, novel with respect to one's own mind.

Now recall the choice between incubation essentialism and incubation phobia. How does our model balance between the two? Incubation essentialism, recall, requires incubation for creativity, without the first, you don't get the second. Even without our model, we can introspect counterexamples against this view. Assuming you've had a creative idea or two, haven't some of them come when you were consciously attending to the problem? If not, isn't this surely possible? The answer seems an obvious yes and so, at least phenomenologically, incubation essentialism looks false. We can also use the

basics of our model to show the neuropsychological implausibility of essentialism. Our analysis certainly supports the hypothesis that *some* creative thoughts result from incubation, but it does not evidence the claim that *all* such thoughts are so explained. Practice, attention, and effort may induce decreased cortical arousal and automaticity, but not necessarily before cognitive benefits can be gained from those efforts. Neural networking can change very quickly (assuming that is even necessary for a creative thought), so surely a creative solution to a task or problem may be secured *before* a decrease in attention and conscious effort, that is, before incubation. Thus incubation is not, as modeled, essential to creative thought.

Now for incubation phobia. Our model offers at least two advantages. First, we have maintained naturalism, but have done so without eliminating our target explanandum: we have modeled incubated cognition as consistent with the introspective and behavioral conceptualizations and have done so in scientifically responsible ways. Second, incubation as we have modeled it is *not* basic or specific to creativity but to *cognitive novelty*. Incubated cognitive processing is a basic capacity—enabled by neural plasticity and the effects of attention—fundamental to functional working of a variety of cognitive abilities (memory, learning, and mastery of cognitive tasks, skills, and information). So there is some reduction here: incubated processing is important to creativity *because* it is important to how we learn, practice, and engage with novel tasks, skills, and information—how we are capable of cognitive novelty. But note the reduction is not one to learning *simpliciter*, but to what we might call *learning without looking*. Incubation is important for these reasons *even if* it results in no creative output.

It is thus an explanandum for any naturalist: it is not something the naturalist should be frightened of but rather should consider herself obligated to explain.

We thus split the difference between incubation essentialism and incubation phobia.

6.2.5. Martindale, Mendelsohn, and Mednick reconsidered

We further clarify our positive proposals by reconsidering some of the psychological theories we used as starting points. The neuropsychological considerations of 6.2.2 and 6.2.3 vindicate some of these theoretical posits and cast doubt on others. First the good, then the bad.

The fundamentals of Hebbian learning provide a neural basis for the importance to creative cognition of associative hierarchies (Mednick) and the importance of attention range (Mendelsohn). The formation of cell assemblies involves, in response to stimuli, the building of connections between neural cells and the strengthening of already existing connections. The denser the connections between cells and other cell assemblies, the more extensive the associative hierarchies become. Mendelsohn argued that crucial, perhaps most crucial, to creative thought is not high excitatory strength of activation but the quantity and density (the range) of associative activation. Mednick's thought was that flat associative hierarchies were needed for a greater range of activation. And some people have *flat* associative hierarchies while others have *steep* ones; the first people are creative, the second, uncreative. But carving up the terrain this way is mistaken. We all may have dense networks between neurons and cell assemblies. What determines if the density is maximally activated, thus flat rather than steep, is the level of attention.

Sharply focused attention will decrease the likelihood of the activation of more weakly connected cells and assemblies which are perhaps promising to the task or problem at hand.[69] This is the insight that Mendelsohn seemed to have his finger on, but that Martindale pins down more explicitly.

Martindale is right to emphasize the relevance of cortical arousal to creative thinking. We can divide his basic thesis into a weaker and stronger version. The weaker thesis (W) says that a decrease in cortical arousal, given its effects on cognitive processing, may contribute to or enable creative thought. The stronger version (S) draws a negative correlation between high cortical arousal and creative thought: increased arousal entails less potential for creativity (Martindale 1989, 1990, 1995). Martindale typically opts for the stronger thesis, which ultimately *requires* of creativity decreased cortical arousal and the enabled mode of primary processing. The neuroscientific data only support (W), and cast considerable doubt upon (S).

According to Martindale, less cortical arousal corresponds with a defocusing of attention. This mode of attention allows for a weaker, but denser activation of available associations and concepts. Activation of these associations is likely to lead to conscious associations, thoughts, that are less predictable and perhaps novel. Starting with something like the Hebbian framework, coupled with the data on attention and automaticity in 6.2.3, we see why Martindale's basic insight is very plausible. Upon

---

[69] This claim is not to be confused with a commitment to *active inhibition* (also sometimes called *lateral inhibition*) of cells and assemblies, which require inhibatory connections between the (inhibited) cell/assembly in question and the inhibiting cell/assembly. Martindale sometimes talks of associateve hierarchies and attention in this way (e.g. Martindale 1995: 256-7) but offers little evidence to support his application of the notion (and moreover slides between that notion and the following, weaker notion that we are commited to). The present claim only commits us to the favoring of some connections in virtue of the strength of those connections (which thus results in less chance for activation of more weakly connected cells). This is consistent with there being a relevant role for active inhibition, but does not require it.

initial encounter with a problem, especially when the problem is a relatively novel one, attention and effort will be high. This correlates with high or widespread cortical arousal. After some time the brain, as it were, settles down and finds its bearings. (And the time required can be very short: in some studies, practice of less than 15 minutes dramatically reduced the cortical activity. See Peterson et al. 1998) Cortical activation reduces (in area) and becomes more concentrated. In the regions that remain stimulated, activation feeds on activation: cells continue to fire and networks increase in number and strength. The cognitive processing becomes more and more automatic, thus involving less potential for interference and requiring less focused attention. This is enough to vindicate (W), Martindale's weaker thesis: decreased cortical arousal may contribute to or enable creative thought. In such stages of processing, the thresholds for the activation of cell assemblies is lower (as they are strengthened first by conscious attention and second by the shift to more concentrated ranges of activation), and the attention required is lower since much of the processing has been made automatic. All of this heightens the probability that an unusual association or concept application will be made.

Some of the same considerations, however, tell against (S), Martindale's stronger thesis. (S) requires of creative thought, primary processing—which itself requires decreased cortical arousal. He proposes that there is a negative correlation between high cortical arousal and creativity. So creative thoughts will only come when we are in a semi-automatic, less-attentive stage of cognitive processing. But there is no reason to think this is true, and good reasons to think that it is false. The first reason is one already given against incubation essentialism above. Our model acknowledges that attention to a stimulus changes the brain in ways that enhance associative connections; attention thus

re-wires the brain (at least temporarily) to improve functionality. Continued practice, attention, and effort may induce some degree of automaticity and something like Martindale's reduced cortical arousal, but perhaps not before cognitive benefits can be gained from those efforts. Brain circuitry can change very quickly, and so surely a solution to a task or problem may be secured before the proposed decrease in attention and conscious effort. This criticism is consistent with the criticism—offered on largely intuitive grounds—made against candidate cognitive manipulators in 5.1.2 (and against incubation essentialism above). Although some creative insight clearly comes from unconscious processing, not all of it does.

Here is a second reason to doubt (S). The stronger thesis assumes that with decreased cortical arousal always comes decreased or defocused attention. But the neuroscientific evidence does not show this. Instead, it shows that decreases in cortical arousal enable automatic processing of certain parts of a task. This may result in Martindale's defocused attention, but it is not entailed by automaticity. It is possible that with the introduction of automaticity attention is just shifted to other parts of the task, namely, those parts that have not already been mastered by the system. And surely a creative insight could come in this cognitive stage. The more automatic a task or skill, the more free cognitive energy to spend elsewhere.

Finally, the neuroscientific literature considered undercuts the creative/uncreative person demarcation endorsed by Mednick, Mendelsohn, and Martindale. As we have shown, novel conceptual associations are not a matter of some people having flat associative hierarchies and some steep ones. Rather, they are a matter of the kind of attention that one pays to a stimulus. If one pays very close attention to a few

associations, say when initially approaching a new problem, then other, more weakly connected, associations will not be activated. In terms of cell assemblies, when the strongest of connections between cells are excited, the weaker ones are less likely to activate. With lesser attention, however, the chances of weaker connections being activated increases. Of course, some persons may have denser networks of associations about some domain. To use Mednick's terminology, a given person could have a flat association hierarcy for some domain, say sports, and a steep association hierarchy for another, say physics. This is just to acknowledge the potential for expertise (and our commitments are consistent with this). But none of these points do the work of separating two kinds of person: creative and non-creative.

So an emphasis on the role of attention, while acknowledging the importance of associative connections, is appropriate. But Mendelsohn and Martindale also go wrong in distinguishing creative from uncreative people, their mark of distinction being the ability to defocus attention, to enter Martindale's stage of primary processing (e.g. see Martindale 1995: 259.) Brains are neurally plastic and our engagement with the world changes them in functionally important ways. One of these ways is the evocation of automaticity: with practice of a task, parts of the task are internalized, thus not requiring conscious attention. This shows how someone could enter Martindale's stage of primary processing. But it does not show that some people are able to do it and others are not.

One might wonder at this point, just what Martindale (and Mednick and Mendelsohn) mean by 'creativity'? And given both our analysis and the data presented by these researchers, is the distinction between creative and uncreative persons just an imprecise or shorthand way of making a distinction in degree not kind? So what they

really mean is that creative people tend towards one end of the arousal/association spectrum and less creative ones towards the other end (low/flat and high/steep, respectively). Call this the *degree interpretation.* This, after all, would be the natural way to interpret a figure like 6.2, and would indeed be charitable.[70]

It is too charitable. First, it is simply unclear what they mean by 'creativity.' There is just no consistent notion of creativity throughout these discussions. At times, creativity is defined as ideas that are both (historically) novel and useful, (Martindale 1989: 89, 1995: 250; Mednick 1962: 221).[71] This definition is inadequate: we argued for the non-necessity of historical novelty in 1.2 and the non-necessity of utility in 1.4.1; even granting the utility (constructive) condition, the two conditions are insufficient even for minimal creativity given our analysis in chapter 1. At other times, creativity is defined as admitting of degree, where that degree goes up with the remoteness of the conceptual association or application (Martindale 1995: 257; Mednick 1962: 221). And at yet other times, creativity is understood in terms of the four-stage process (see introduction-this chapter) of Wallas and others (Martindale 1989: 90, 1995: 251). So it is not clear what an agent either has or does not, or has to some degree or not in order to be

---

[70] In fact, in at least one paper Martindale (1999: 140) makes an initial demarcation in terms of degree: arguing that there should be a parallel continuum between low/high arousal and primary/secondary processing (the latter, for Martindale, corresponds to activation of flat/steep associative hierarchies). He favors the hypothesis that creative ideas come at lower levels of cortical arousal and a shift to primary (free-associative) processing (and the opposite for uncreative ideas), but admits that the data underdetermines whether creativity should be located on the continuum in this way. He then appeals to self-reports of very creative persons, geniuses, to solve this problem of underdetermination. These reports of course favor his hypothesis: all the most famous reports (think of the examples discussed previously: Archimedes, Coleridge, Kekulé) discuss flashes of insight that come at low levels of attention. The problem is that these are the most famous reports for a good reason: they are the most memorable, striking us as remarkable if not bizarre. But there is no good reason to think that they provide the standard for creative insight, and so no justification for invoking them to motivate controversial claims and distinctions about creativity, attention, and arousal.

[71] Note also that Martindale explicitly commits to a denial of psychological-novelty, claiming that "Were someone to rediscover the theory of relativity, we would think the person quite clever but not creative because the idea has already been discovered" (Martindale 1995: 250).

on one end of the spectrum or other. Finally, even if these various conditions can be worked into one definition, we still need guidance on the application to persons (e.g. do *only* creative persons have these kinds of ideas? are creative persons the ones with some significant quantity or frequency of these ideas? etc.) No such guidance is offered.

One might respond by suggesting we take Martindale at his word: "Creative people can get themselves into primary process states of defocused attention" (1995: 259). Following Kris, "creative individuals are better able to alternate between primary process and secondary process modes of thought than are uncreative people (Martindale 1989: 90). "There is a good bit of evidence that creative people are in fact physiologically over-reactive" (1989: 99). "[C]reative subjects exhibited low arousal during inspiration and high arousal during elaboration" (1995: 260). "Creative people obtain high scores of psychoticism" (1989: 91). So if we want to apply the degree interpretation, we just say that creative people are the ones with more of these features and the uncreative people, less.[72]

This won't do. One cannot say that creative people *just are* the ones with the above symptoms—e.g. a creative person is one who can flatly associate—since that both uses these symptoms to *define* a class of individuals and to test whether those individuals are members of the defined class (by testing for those very symptoms). This is a non-starter: it will never draw a principled line, sharp or fuzzy, between classes of individuals and it will never vindicate a theoretical hypothesis or definition. Alternatively, we might assume, given the language being used here (e.g. 'creative people have F', 'creative people can do G'), that the line is already drawn. Creative people are already identified

[72] Note, however, the sharp distinction-language use. These uses are not exceptions: they are ubiquitous in these discussions. So if it is just terminological shorthand, it is anything but innocuous. This too suggests that the degree interpretation is too charitable.

as creative prior to testing for potential for primary processing, low cortical arousal, etc. But then we are back where we started: we must either appeal to the cluster of inadequate definitions offered, or to the introspective reports of geniuses.[73] And neither option will do.

This discussion has been brief, and there is a great deal more to consider. Suffice it to say that the notion of creativity in the work of Martindale and others is underdeveloped, and with it the theoretical conclusions they draw (including the creative/uncreative person divide). Nonetheless, this research does provide useful insight on attention and association, and we can employ those insights without settling these other disputes.

If minimal creativity is something we are all capable of, and is basic to more radical creativity, then we should not expect radically creative persons to have some basic cognitive capacity that the rest of us lack. They may, of course, be better at "automatizing" or at defocusing attention on a task. They might be better at gaining expertise in some domain. This may be two ways of putting the same point. In any case, these are differences in degree not kind. This is an important moral to clarify. It supports the thesis that creative capacity is basic to human cognition: it is something we are all capable of—though perhaps to varying degrees—and is plausibly essential to our survival.[74]

---

[73] As discussed (see fn. 70), Martindale, at times, makes the second move. Mednick makes something like the first. His definition of creativity, which involves novelty, utility, and some degree of remoteness, is assumed and then structures the remote associations test used. So the test, as noted by Mendelsohn (1976: 341), is only adequate insofar as the definition is.

[74] More on this point at the close of 6.2.6.

6.2.6 A final note to the other spooks

This model also accommodates flash phenomenology, at least when creative thoughts result from incubation. If we return to a stimulus after periods of incubation (or even just maintain our focus on that stimulus for a period of time in a non-incubated case) newly established or strengthened associative connections may be activated. If we take seriously the basic posit that memories, concepts, associations, etc., are "located" via certain patterns of activation, when we attend to the right stimulus, the cell assemblies correlative with these mental entities may be activated via some network or another. From a subjective point of view, this translates to the experience of an idea popping into our heads. There is nothing mysterious about this: it is a simple, though no less remarkable, feature of cognitive processing.

Recall *ex nihilism*, which says that given their novelty, creative ideas emerge from nowhere. Our model shows that, one, they come from somewhere and, two, they are not unlike lots of uncreative ideas in this respect. Creative thoughts are bound to particular cognitive profiles. Consider a thought which is an incubation effect. This thought depends (in part), according to our model, upon new and newly strengthened connections between neural cells and cell assemblies. These neural changes depend upon previous stimuli and the resultant cell assemblies (and their strengths), upon current stimuli, upon attention, upon the degree of automaticity involved in processing. Some of these new connections surface in consciousness as novel ideas. Assume that whatever other conditions one puts on creativity are satisfied: the thought in question is creative. Note then that this maintains a genuine novelty without invoking *ex nihilism*: the thought hardly came from nowhere, it depended upon a number of states or properties of the

cognitive profile. If we can offer this explanation in the incubated case, there is little reason to think that we cannot offer one in the non-incubated case.

We close with the following simple moral: cognitive novelty is possible given the plastic nature of the brain. There is nothing paradoxical about it, as some have suggested. In fact, it is plausibly essential to human survival, which requires, if it requires anything, the acquisition of new skills, the ability to learn new information, augment and apply concepts in multiple circumstances, and so on. The cognitive processing modeled in the last several sections explains (in part) how we are able to do such things. Both novelty and multi-tasking are enabled by neural plasticity and the functional effects of attention on that structural feature of the brain. This is true of the cognitive function of human brains generally and thus we all *do* process information in the ways suggested.[75] If the latter was not a fact, we simply would not survive, as we would learn little if anything, be unable to adapt to new circumstances, and be unable to complete more than one or two simple tasks at once.

And note that this feature is not specific to creativity, and is not specific to incubated cognition: it is a general feature of cognition. But focusing on incubated cognition as one way creative thoughts are tokened makes the general lesson salient. As a bonus, it reveals the scientific tractability of some other traditionally spooky features of creativity. This explanation should give the naturalistic philosopher and cognitive scientist another good reason to return serious attention to creative cognition.

[75] It is undoubtedly true of the brains of many other species. But this issue will not be treated here.

6.3 Incubation and imagination

We turn, finally, back to imagination. At the close of 5.2.4, we proposed that imagination could likely serve the role of unconscious cognitive manipulator: it is *incubation-functional*. We have now offered a more robust conceptual and neuropsychological model of incubation. To make good on our suggestion, we must ask whether any part of incubated cognition thus understood is rightly called imaginative.

Upon consideration of the conceptualization of imagination and its various features offered in chapters 4 and 5, one might take a quick look at our model of incubated cognition and infer that none of that kind of cognitive processing, whatever it is, is imagination. The model of incubation centered, on the one hand, on attention and associative connections and, on the other hand, on neural cells, synapses, chemical excitation, etc. None of this meets any of the descriptions of imaginative character or function, content, inferential and cognitive connections, or the rest. So incubated cognition is not or does not involve imagination.

To make this inference is akin to the category mistake of Cartesian dualism, famously articulated by Gilbert Ryle. According to the Cartesian, Ryle proposed, our mental concepts are not properly applied to mechanical bodily processes.

> The differences between the physical and the mental were thus represented as differences inside the common framework of the categories of 'thing,' 'stuff,' 'attribute,' 'state,' 'process,' 'change,' 'cause,' and 'effect'. Minds are things, but different sorts of things from bodies; mental processes are causes and effects, but different sorts of causes and effects from bodily movements. And so on. Somewhat like the foreigner expected the University to be an extra edifice, rather like a college but also considerably different, so the repudiators of mechanism represented minds as extra centres of causal processes, rather like machines but also considerably different from them. Their theory was a para-mechanical hypothesis (Ryle 1949: 19).

Thus the mistake of the dualist is to assume that mind is some mechanism over and above the observable bodies, another individual in the category of observables like various elements of the nervous system and the brain. He is like the spectator at the military parade who, after viewing various military personnel, vehicles, and weaponry, asks when he will see the battalion. The dualist, upon not finding the mind like he finds individual elements of mechanism, infers that it must be something else: it is the ghost in the machine.

There are a number of options for the materialist here: she can opt for type or token identity, supervenience, emergence, or some kind of functionalism. This choice is of no matter to us. What is important is that we can also respond to the Cartesian's mistake by distinguishing levels of description. At one level of description, we may describe the mind in behavioral terms: in terms of publicly observable action. At another level, in folk psychological terms: in terms of beliefs, desires, intentions, thoughts, emotions. At another, in neuropsychological terms: in terms of cortical regions, neural cells, assemblies, activation pathways and patterns. And so on. This is an option available to any materialist (well, except for the logical behaviorist).

Is the same true of our present concern? We must ask if our imagination concepts, at one level of description, and our neuropsychological concepts, at another level, compose descriptions of the same phenomenon. When asking this question, we should be careful to think of ourselves as faced with a decision point rather than a challenge to make some discovery. We have a great deal of the empirical science in place: we must assume that the fundamentals of neural plasticity and networking are correct. We have some independent conceptual machinery: notions of imaginative

character, content, and relations to other mental states. We must decide, given that both conceptualizations are of our own making whether the one could plausibly be a description of the same thing as the other. In other words, we must decide if the descriptions are merely different in level of description and are actually (at least to some degree) co-extensive. What issues and questions should constrain this decision?

One place to start is to unpack the proposals in 5.2.4, which indicated that imagination seems capable of functioning in unconscious ways. Our imagination often takes on a life of its own, yielding imaginings that we didn't expect, no less ask for. Walton proposes that imaginings, just like beliefs and other propositional attitudes may be *occurrent* or *non-occurrent* (Walton 1990: 16-18). Non-occurrent mental states, as contrasted with occurrent ones are states not presently attended to but ones which nonetheless have causal effects on states that *are* presently attended to. For example, my belief that '2+2=4', is maintained (i.e. is still a belief possessed by me) even when I am not consciously thinking it, and perhaps even using it in some basic arithmetic calculations. If this were not a feature of belief, I would be re-forming thousands of beliefs all the time.[76] Imagination possesses the same feature. For example, when I begin a project of imagination, I will begin with one imagining, which leads to another, which leads to several more, and so on. When I begin reading *The Brothers Karamazov*, I start with some basic imaginings, say, that the father, Fyodor Karamazov, is a menace

[76] Another common, but mistaken, way to clarify the notion of occurrent mental states is to distinguish them from *dispositional* states. My belief that campus is east of me disposes me to take the appropriate bus east, even though I am not consciously thinking the former thought. The belief is thus dispositional and not occurrent. This characterization is unhelpful: the dispositional feature fails to effect the right discriminatory cut, since even fully conscious beliefs *dispose* me to act one way or another. (For instance, if I am new to Vancouver I have to make a conscious mental note of the direction of campus.) Put another way, the characterization implies that occurrent states are non-dispositional (and non-occurrent ones are dispositional). But that cannot be right: beliefs, desires and the like are all dispositional, no matter if they are conscious or not. This is part of what it is for a state to be a belief or a desire, namely, for it to function as an element in that system, disposing it to behave in certain ways as a result.

and a fool; the eldest son, Ivan, is a somber and calculating intellectual; the youngest son, Aloysha, is a free-spirited, generous theist; and so on. If I am to get far in Dostoevsky's 800 plus page masterpiece, these imaginings must be maintained but cannot occupy my conscious attention. They are thus formed as occurrent imaginings, but become non-occurrent early in my engagement with the novel. (And of course, they will sometimes be drawn back to occurrent status throughout my reading: they can shift between being occurrent and non.) Currie and Ravenscroft's counterpart model lends further support to this proposal: if beliefs and desires, by their nature, may be both non-occurrent and occurrent, then their counterpart states, make-beliefs and make-desires, will follow suit. The same goes for other imaginative states.

This is fine for a start: non-occurrent imaginings provide one way to think about imagination as incubated cognition. However, it only takes us so far. We need now to consider a couple of issues: operational constraints and cognitive output. What kinds of constraints operate on incubated cognition as we have modeled it? And what kind of output does incubated cognition yield?

First the constraints. One thing that Freud got right is that unconscious cognition is pretty wild. Bizarre conceptual combinations and thoughts surface during and after daydreams and dreams. The same is largely true of incubated cognition. There is little restriction on how associative networks form and strengthen. There are, of course, physical constraints on how the brain networks: the proximity of cells and the synchrony of their firing, patterns of activation, current attention and stimuli, among other things, determine the changes in functional networking. But there is no restriction on change as far as truth is concerned. That is, if we take seriously the suggestion that memories,

concepts, and associations correlate with certain neural assemblies or certain patterns of activation of assemblies, then we see that there is no veridical or world-tracking restriction on how those mental entities connect. If you happen to experience Nixon's face repeatedly with 'Clinton', then that association will be formed and your brain is networked accordingly, in spite of the fact that the former object does not answer to the latter name. This is the very reason the Freudian term 'free-association' has stuck: we are free to associate just about anything with anything.

The same fact is observed at the cognitive-behavioral level in the kind of output that incubated processing yields. Given time with a mundane input—say mention of a word, an everyday perceptual experience, or some thought or association related to a problem to be solved—very unexpected and (ostensibly) unconnected thoughts sometimes result. Sometimes these thoughts are useful for solving the problem with which one began, sometimes they are not. The important point for now is that the output states of incubated cognition sometimes have no obvious conceptual resemblance to the input information, and preserve nothing like truth in relation to the input. This indicates that incubated cognitive processing, at least some of the time, is non-truth bound.

We learn something else from looking at the output of incubated cognition. What goes in need not come out: sometimes the output is vastly different from the input. (This is the same basic feature as just discussed, but we will now tease out a different point.) Following a period of incubated cognition as we have construed it, a problem may "come to us" re-framed, as being approachable from a different angle, as amenable to a different kind of analysis, or sometimes, as being already solved! What does this indicate? It indicates that incubated processing "does something" with the input: it involves what in

chapter 5 we called *cognitive manipulation*. Incubation yields novelty. This is the datum that intrigued Smith and Blankenship, motivating them, among others, to perform studies on incubation effects. Putting this point together with the point of the previous two paragraphs, we get the following: incubated cognitive processing involves non-truth bound cognitive manipulation.[77]

In chapter 4 we catalogued the candidate states for conscious creative cognition. Given the need for non-truth bound cognitive manipulation, we ruled out capacities like belief formation and skill acquisition in virtue of their truth-boundedness. (That is, we argued that these states would not be *enough* to yield creative results.) We also ruled out mere entertainment of propositional content, since it lacked the manipulative feature needed. We also considered and rejected a number of possible candidates from other theorists—variation generation, primary processing, runaway learning. We finally argued that imagination, given its nature both as a workful and playful kind of cognitive processing was best suited for the role of (active) cognitive manipulator. We have now found that incubated cognition, at least when it yields novel output, requires the same kind of processing (where the processing functions, at least partly, unconsciously). Given the constraints on processing and the kinds of output yielded, imagination is most appropriate to serve this role. It is non-truth bound, manipulative, and can function unconsciously or non-occurrently. Imagination is thus a mental concept appropriately applied to incubated cognition as we have modeled it.

What does this claim commit us to? We have not claimed that incubation *is* imagination. We have not claimed that *all* incubation involves cognitive manipulation

[77] Given the discussion of 6.2.3-6.2.6, this is no surprise. But it is useful to make these points clear so as to compare the conceptual features of incubation and imagination.

nor that *all* incubation is non-truth bound; we thus have not claimed that *all* incubation involves imagination. There are no doubt a variety of other kinds of states and processes constituting incubated cognition: imagination concepts are not the only applicable mental concept at the folk-behavioral level of description. But imagination concepts *are* appropriate. And they are especially appropriate when incubated cognition yields novel output. Our aim was only to show that our two conceptualizations adequately describe, at different levels (one folk psychological, one neuropsychological) the same (or part of the same) phenomenon. We have done that. By conjoining the morals of this chapter and the previous one, we get the following: imagination plays a role in creative cognition, conscious and not.

Conclusion

As the target analysandum requires, our analyses and the model it motivates are complex and varied in treatment and discipline. The first section of this conclusion is an attempt to connect and clarify that complexity and variety. The second section considers the model, and offers a handful of general observations. The final section selects and considers, out of an overwhelming number to choose from, some issues and questions for further research.

The model

We began our analysis with a definition of creativity as a psychological phenomenon. The definition comprises three necessary and conjointly sufficient conditions for attributions of *minimal creativity.*

> *MC:* Some x is minimally creative if and only if, for some agent A, x is the product of the agency of A; x is psychologically novel; and x could not have been tokened by agent A before the time t_i when it actually was tokened by A.

These conditions *describe* the basic phenomenon of creativity. The agency condition requires non-trivial responsibility. The novelty condition requires a certain actual cognitive profile, such that the thought in question had not, as a matter of fact, been previously tokened by the agent. The modal condition requires a certain modal cognitive profile, such that the thought in question could not, as a matter of fact, have been previously tokened by the agent. This, we found, already sets the analysis apart from a variety of traditional approaches: inspirationalism, irrationalism, and strong ex nihilism, to name a few.

Cognitivist in ambition and approach, the analysis attempts to explain the psychology of creativity, as much as possible, in terms of more basic cognitive mechanisms and capacities without going so far as reduction. A framework for this approach comes in the form of a functional analysis of problem solving and creativity. The general lesson this analysis affords is straightforward: problem solving and creative thought are functionally similar, constituted by the same kinds of functional roles in spite of the potential for widely different occupants of those roles. But the purchase of our functional analysis does not stop there, as we see by recalling the individual roles identified, and the treatments of those roles that follow.

We started with a simple theory of problem solving:

> (T) Given some agent *x* in some context, the context presents *x* with a problem, which requires for its completion some solution set achieved via some heuristic set.

We then translated (T) to its Ramsey sentence:

> (T_R) There exists some P_1, P_2, P_3, P_4, and some *x*, such that *x* is in P_1, and P_1 presents *x* with P_2, which requires for its completion P_3 achieved via P_4.

The Ramsey translation eliminates all theoretical terminology in favor of functional roles. And as we name the variables in various ways, we begin to see more clearly how the roles relate and what kinds of things may serve what kinds of roles.

P_1 correlates with context. We learned a few important lessons about how this role may be occupied. In general, a context constrains what heuristics and what solutions are available (thus how P_3 and P_4 can be occupied) *for a particular agent* addressing a task or problem (the occupant of P_2). This was especially salient when the occupant of P_2 is less determinate. We considered instances of conceptual augmentation and re-mapping (recall Bach's *Well Tempered Clavier* and the famous Gettier cases) and instances where

new conceptual space is carved (recall Kekulé's work on the benzene molecule and e.e. cummings' innovations in poetry). In circumstances like these, the task just becomes an exploration of the conceptual space (or what there is of it), shifting much of the burden of constraining the solution set and heuristic set from the problem role to the context role.[77] Thus in many radical cases of creativity, like any of the four just mentioned, P_1 is vitally important: context significantly constrains the creative outcomes that result.

The discussion of P_1 just offered serves double duty, also revealing certain facts about P_2. In some problem-solving circumstances—perhaps the ones commonly called 'problem solving'—the occupant of P_2 is determinate: add these figures and then determine the median and mean; fit this sofa through that door; pay all of these bills with this paycheck. These problems may or may not be solved in creative ways: that will depend upon the other roles (context and heuristic set) and the agent in question. Other times, the occupant of P_2 is less determinate or non-determinate: write a prelude and a fugue with only the following parameters—explore each key on this scale, write for this instrument tuned this way; consider the standard grammar and diction of poetry and see what can be changed, tweaked, thwarted. These variations in determinacy affect how the other roles may be occupied, constraining what kinds of heuristics are available to secure what kinds of solutions, and determining which and how much of a context is relevant. This is an example of the reciprocal relation between functional roles in instances of (T_R): P_2 constrains P_1 and vice versa. Bach's choice to explore the 12-tone scale and tempered tuning informed the music-theoretical space that would be relevant, and that space informed the nature of the problem Bach had made for himself.

---

[77] Sliding between theoretical language and functional role language is important to avoid, but is done here as a shorthand to aid recall and comprehension of chapter 3.

Cases where the occupant of P_2 is less determinate encourage another observation. When the problem or task is less determinate, the candidate heuristic methods are "freed up" (at least typically, since this too depends upon how the other roles are filled). Cases of conceptual exploration make this clear. In a situation like this, say cummings' exploration of poetic techniques, there is considerable cognitive freedom, for what kinds of cognitive capacities are employed and for what kinds of contents are tokened. This, as we will see below in our discussion of P_4, sits well with our emphasis on cognitive manipulation, imagination, and incubated cognition.

Occupancy of P_3 is constrained most straightforwardly. It, more than the other roles, is generally just a consequence of the agent and how each of the other roles are filled. Of course, to use T-language, if the problem is less determinate, say in the case of conceptual exploration, then what may count as a solution will also be less determinate. However, once all the other roles are filled, the set of potential occupants of P_3 narrows considerably. And this set is narrowed further by the psychological properties of the agent, identified with that agent's actual and modal cognitive profiles. These properties, whatever they should be, are the ones that matter for creativity. That is, assuming we are talking about a responsible agent, it is the actual and modal cognitive profiles of that agent vis-à-vis the mental token in question (i.e. which is identified with or partly constitutive of whatever occupies P_3) that determine the creative status of that token *for that agent*. If they have tokened the thought before or could have tokened the thought before, then it is not creative for failure to satisfy the novelty and modal conditions, respectively.

As the above discussion reminds us, the occupant of P_4 is especially constrained by the occupants of P_1 and P_2 and by the cognitive profile of the agent. In theoretical terms, what heuristic methods an agent can use to secure a solution to a problem as situated in a context, depends on those features and on the actual cognitive profile of that agent. The present thesis aims for a model of creative cognitive *processing*, and so we have said the most about this functional role, about what cognitive means are employed for creative ends. In 3.2.2, the seed was planted for what was to come in chapters 5 and 6: creative thoughts require non-truth bound cognitive states in the causal path that leads to their tokening. Bach's *Well Tempered Clavier* is once again an illustrative example. All of the relevant knowledge available—of tempered tuning, the clavier, and the twelve tone scale—is important but insufficient for creation of the musical structures embodied in Bach's masterpiece. Instead, we suggested, some cognitive manipulation was required: Bach had to explore and manipulate the space in non-truth bound ways to get the result he did (to occupy P_3 in the way it was occupied).

This suggestion was generalized in 5.1.1 in the form of the *cognitive manipulation thesis*. Any state that faithfully tracks the contents of some conceptual space can, at best, play a necessary but insufficient role in thinking creatively in that space. This is as true of mundane instances of creativity like Norm's learning quantum mechanics as it is of the innovations of Bach, Gettier, and Kekulé. Cognition that *is* creative requires more than rote memorization or mere acquisition of knowledge or concepts. Considerations of cross-categorical comparisons (in each of the Karmiloff-Smith, Caciarri, and Ward studies on the drawings of both children and adults, discussed in 5.1.1) and development of figurative and literal linguistic competence (in the Caciarri studies, among others,

discussed in 5.1.1) encouraged this thesis, since even these very mundane cognitive capacities plausibly require cognitive manipulation. If these accomplishments are creative or, less committally, if creative cognition requires *at least* this much of its agents, then the cognitive manipulation thesis is well-motivated. We thus identify an important feature of P_4 for creative instances of (T_R): cognitive manipulation (partly) occupies P_4 in cases where the occupant of P_3 is a creative thought. In other words, a sub-role of the heuristic set that leads to creative solutions is the cognitive manipulator role. In fact, the cognitive manipulation thesis coupled with our definition *MC* tells us that this *must* be so. If an agent's thoughts are about some conceptual space, then some of those mental tokens must *not* immediately derive from the space; they must be manipulated in some non-truth bound way, else they not be p-novel and thus not minimally creative. Until cognitive manipulation becomes part of the heuristic method, the agent cannot token a p-novel thought. This provides a clear way of understanding the claim, discussed in 1.3.3, that *MC* implies: modal cognitive profiles depend on actual cognitive profiles, so that what an agent *can* think (and thus whether a thought satisfies the modal condition) depends upon what and how she *has* thought.

The analysis then identifies imagination as possessing the cognitive freedom requisite for occupying the cognitive manipulator role. This identification was founded on the conceptual and empirical features of imagination distinguished in chapter 4. Imagination varies in *character*. These variations can be identified with the phenomenologies of different imaginative states—of, for example, make-beliefs versus make-desires versus visual images. Or they can be identified with the functional roles of those states. Or both. (See 4.1.2, 4.1.3, and 4.2.1.) Imagination may vary in *content*; it

depends who you ask. Some maintain that mental images possess a unique, pictorial content. Others maintain that images are no less propositionally contentful than beliefs or desires. (See 4.2.2.) Imagination varies in *richness*. We may imagine in ways that are quantitatively rich, imagining more propositions or forming more images that surround the target imagining. This is richness in content. We may imagine in ways that are qualitatively rich. Here again the issue can be parsed into questions about phenomenology and questions about functional role. We may engage in qualitatively rich imaginings in the sense that we really "get ourselves in the mood" of our imaginative project, forming states that are phenomenologically robust like their non-imaginative counterparts—feelings of conviction for make-belief, drive for make-desire. Alternatively, we may engage in qualitatively rich imaginings in the sense that we form imaginative states—like make-beliefs and visual images as contrasted with mere suppositions—that connect with other cognitive states in a variety of important ways. (See 4.1.3 and 5.2.2.)

These various features reveal a complex and diverse set of mental phenomena that fall under the rubric 'imagination.' In spite of all of these differences, there is at least one common thread: any state that is an imaginative state lacks, qua imaginative state, commitment to truth or action. This fact is sufficient to support the assumption that imagination is a generic cognitive capacity; it is a capacity that enjoys cognitive freedom. It is a playful cognitive capacity (5.2.1). At the same time, it is workful (5.2.2). It is this very combination of playfulness and workfulness that makes imagination the ideal cognitive manipulator. It possesses the freedom to allow for non-truth bound manipulation; it possesses the right cognitive and inferential connections to encourage the

cognitive purchase of those manipulations. Imagination is thus ideally suited to actively generate novel cognitive states (5.2.3, 5.2.4, 5.3).

This only gets us part of the story about the cognitive manipulator role and the heuristic methods for creative cognition. Introspection and volumes of literature tell us that creative thoughts are sometimes tokened non-actively or unconsciously. Creative thoughts do not always come during conscious attention to the problem or task. *Incubation*, it seems, is sometimes an occupant of P_4 in creative instances of our (T_R). Chapter 6 took on the task of explaining, naturalistically, incubated cognition and how it should be understood within the framework on offer.

After narrowing our understanding of unconscious cognitive processing to processes or states that lack phenomenal consciousness (6.1.2), and then motivating the scientific plausibility of these kinds of processes (6.1.3), our explanation of incubated cognition turned in 6.2 to neuropsychology and behavioral psychology. As a first pass, incubation was understood as a stage of cognitive processing where conscious attention is diverted from a task and during that time conceptual associations are activated more freely and perhaps more broadly. Adopting the framework of the late Donald Hebb, we identified synaptic plasticity of the brain as the neural basis for these incubated processes. Cell assemblies activate and strengthen in response to current stimuli, and this activation continues after conscious attention to the initiating stimuli (6.2.2). A number of researchers have studied how attention to and performance of cognitive tasks affects these very networking features of the brain. Sustained attention to a task will eventually result in the internalization of parts of the task: the more we learn or do something, the more automatic it becomes. This coincides with a reduction in the area of cortical arousal

and results in a lessened potential for cognitive interference (6.2.3). However, these functional changes in automaticity and localization are consistent with the synaptic activation that Hebbian learning involves, since proximal, simultaneous neural activation becomes more likely with the increased concentration of activity.

All of this provides a neuropsychological foundation for incubated cognition, while allowing for some flexibility in describing the concept. Consider the subjective experience of a problem solver and the neural correlates of that experience. As we consciously attend to a stimulating problem, the number and density of connections between certain neural cells and assemblies increases. As we continue to attend to that stimulus, cortical arousal decreases and some of the information becomes subject to automatic processing. Attention is then diverted, defocused or halted altogether, while the relevant cell assemblies and connections between assemblies continue in activation. Our subjective experience at this stage is less important (we might be working away at something else or nothing else). What is important is that this stage is where *incubated cognitive processing* takes place. Returned (conscious) attention to the problem might result in an *incubation effect*. This could come in the form of newly created or strengthened associative connections that solve the problem (*incubatory solution*) or new associations that are crucial to securing a solution (*incubatory preparation*). The model provided is consistent with either way of thinking about incubation, and it is plausible that both are needed to describe a variety of unconscious cognitive processing (6.2.4). Moreover, the model is general in the sense that it explains incubated cognition whether that cognition yields creative results or not. The other components of our analysis do this

job of demarcation in a particular way, but the model of incubated cognition is available for use by alternative theories of creativity.

Finally, incubated cognition can, at the appropriate level of description, be described as involving imagination. We see this by considering the operating constraints on and output of incubated cognitive processing. The association building and strengthening that take place during incubated processing is unconstrained by truth or action: it enjoys cognitive freedom. Moreover, the output of these processes do not necessarily track or mirror the input information: wild conceptual associations can result from tame and predictable input stimuli. Incubated processing thus involves *doing something with* the input information; it involves non-truth bound cognitive manipulation. Imaginative states, as states that may function at both an occurrent and non-occurrent level, are good candidates for components of this kind of processing (though perhaps not the only components). And this is for the same reasons that imagination was identified as ideal active cognitive manipulator: it combines the right amount of cognitive freedom with cognitive purchase. A conclusion? Imagination is part of the heuristic method that leads to creative thoughts, both when those methods are active and when they are incubated. Thus in creative instances of (T_R) imaginative processes will (at least partly) occupy P_4.

This brings us full circle. *MC* identified three conditions on the psychology of minimal creativity. Through functional analysis we identified the various roles operative in problem solving and creative thinking. Identification of the relations and constraints between roles afforded many lessons. With respect to creative process, the relation between the heuristic set role and the solution set role is especially instructive. *MC*

requires of creative thoughts that they be psychologically novel and that they possess a certain modal status. What occupies the heuristic set role will depend on all other elements in the formula, and not until this role is adequately occupied can the solution set role be occupied. In other words, we see why an agent *cannot* secure a solution or token a certain thought until a particular heuristic path is in place. We thus see how, for some particular thought before some time, the modal condition is satisfied. We also see, given our discussions of cognitive manipulation, imagination, and incubation, how the agent *does* secure the solution. Creative thoughts require cognitive manipulation, and imagination and incubation serve this role. It is thus that agents token p-novel thoughts, and thus that the novelty condition is satisfied. We have both defined a psychological phenomenon and modeled how it works.

Some observations and consequences

This model splits many differences. We have acknowledged extreme views on creativity throughout our analysis: the mystic versus the close-minded scientist, the mechanism profligate versus the creativity eliminativist, the incubation phobic versus the Freudian incubation essentialist, inspirationalism versus blue-collar-just-keep-grinding-away-at-it-learning. Our model has consistently learned from these extremes while maintaining a balance between them. Here are three important examples.

Cognitivist extremes: We avoided endorsing both the special mechanism views and the nothin' special views canvassed in chapter 2. Our theory of creativity maintains the uniqueness of creativity, without invoking special creativity-specific mechanisms or modules. The mechanisms invoked—propositional and

procedural knowledge (plus other truth-bound processes), imagination of various modes, incubation—are none of them specific to creative cognition. However, we have not reduced creativity to a mere conjunction of these capacities. Instead, we have identified the roles that each play and suggested that creativity emerges only when particular circumstances obtain. We thus avoid the nothin' special extreme as well.

Incubation extremes: The phenomenon of incubation, like the other spooks we discuss below, invites two opposite reactions. Some buy the whole hog: creativity always and only occurs when agents enter an incubated stage of cognition. Others avoid the hog altogether: incubation is unfit for scientific analysis and thus should be excluded from a scientific theory of creativity (should one choose to give one). Our model of incubation shows how both the incubation essentialist and the incubation phobic are guilty of *over*reaction. Incubated cognition can be understood and explained naturalistically, and is *but one* important but inessential stage in the creative process.

Spooky extremes: The problems encouraged by spooky features like flash phenomenology and novelty, in addition to incubation, were both prefaced and discussed in various parts of the thesis. Our model accommodates these features as part of the phenomenon, but does so within the confines of naturalism. Flash phenomenology is just that, *phenomenology*. It is a feature of the way creative thoughts seem to us as agents. This acknowledgement puts creativity on an explanatory par with more theorized cognitive states like beliefs, desires, and memories: if we can accommodate the phenomenology of the latter states without

spookiness then we can do the same for the former. Moreover, the flash is often symptomatic of incubated cognitive processing, and we have explained that as well. Robust novelty is possible given the plasticity of the brain and our ability to manipulate information in a variety of rich ways.

The model is general. As prefaced, avoiding an analysis of creativity as found in some particular domain is preferable. This is so for at least two reasons. One avoids issues and problems that are symptomatic of the domain and not the creativity in that domain. And, second, a comprehensive theory of creativity will simply find more applications, even if it misses some of the nuances of creativity in this or that domain. Our method was to narrow the focus to creative cognitive processes. At the general schematic level of processes, domain-specific issues are ironed out since they come in at the level of content and thus concern particular instances of creativity (rather than schemas). The focus on creative processes nonetheless enjoys wide application, since all of creativity, be it in pure mathematics, crossword puzzles, or freestyle rap, involves creative cognitive thinking. This is at least part of what makes them all instances of the same phenomenon, of creativity.

The model is entirely descriptive. The conditions of *MC*, and all features of the functional analysis and the cognitive processes that serve creative thought have been described without invoking considerations of value. In the preface, this methodological desideratum was first suggested as a way of simplifying a conceptually rich and often confused set of issues: too many theories of creativity attempt to do too much at once. As a result, they make little headway on the subject they initially target. Our tack was then

to provide a descriptive foundation for the evaluative issues by setting those very issues to one side.

The model is informative and useful. And this is so in spite of its lack of a value component. As we just suggested, it avoids the problems that plague evaluative theories. This makes it an effective and flexible research tool. Once the descriptive story about the psychological conditions is in place, one can build in whatever evaluative issues one likes—socio-historical, art-historical, political, aesthetic, and so on. The descriptive account is open to and provides a stable basis for such addendums.

So much of the mystery of creativity is how we are able to do it. Here are but a few of the puzzling questions that surround the phenomenon: how are we able to think in novel ways, solve new problems, invent surprising things? What processes led to this or that creative act? What kinds of skills, knowledge, or capacities are typically needed for creative accomplishment? What goes on in the mind when a creative insight flashes to consciousness? Our model answers these questions, providing a detailed conceptual and empirical model of the creative cognitive process. This is useful both for general theories of mind, as we will discuss below, and for any theory of creativity.

The model makes some evaluative issues more tractable. Making substantial good on this claim requires more treatment than can be offered here, but here are a few brief sketches.

> *Best first step*: Assume that creativity does turn out to be an ineliminably evaluative concept, either generally or as situated in particular domains. Even if this is so, the first step in an account of value is always to identify what we can at a descriptive level. Even a value theorist who insists that evaluative facts cannot

be entirely demarcated from descriptive ones can acknowledge the merit of proceeding as if such a cut can be made. That is, one might insist that phenomena like aesthetic value, moral value, or creativity, cannot be accounted for if evaluative facts are excluded, and this is partly *because* those facts are not distinguishable from the descriptive facts. Nonetheless, such a theorist can maintain that a good first step is to proceed as if the distinction can be made, such that the non-controversial descriptive facts, however few they should be, can be identified and put aside. Consider a few analogies with mundane cases of evaluative judgment. If we want to ask whether someone's actions were praiseworthy or blameworthy, we first get straight on the descriptive facts of his action. If a court is deciding if a defendant is guilty of some crime, they first debate and identify the physical evidence: what he did, what he left behind, where, how, and so on. If you and I are arguing over the aesthetic or artistic value of some artwork, we had better be sure that we agree what the artwork is: is it just that object affixed to the wall, or that plus some artistic intention, or that plus some social context, and so on. Individuating these descriptive facts precedes consideration of the evaluative ones in the order of analysis. To be clear, these cases are different since they concern individual instances of evaluative judgment and our concern is with general accounts of values or inherently evaluative phenomena. Nevertheless, the progression from descriptive to evaluative concerns is common to both; and this is instructive for theories of creativity. Our model best provides this first step: it identifies the descriptive features of creative cognition.

Narrowing: This point emerges directly from the last one. If our theory of creativity is accurate, and the descriptive features identified are in fact features of any creative property instance, then the theory effectively narrows the decision space for a value theory of creativity. Descriptive issues are often entangled with evaluative issues. Our analysis gives a model for identifying the first, so that one can then set them aside and move on to the second. In other words, whatever evaluative concerns attach to creativity, and whatever features they typically target, we can initially set aside as irrelevant issues of agency, p-novelty, modality, and the various components of creative cognitive processing as modeled. These features are descriptive and, at least to start, are to be separated from the evaluative ones. This gives the value theorist of creativity less to think about.

Marking: At the same time, there is no reason that these descriptive features cannot, from within some particular domain or social context, be evaluated. That is, although the features can be identified without appeal to values, they can be assessed through appeal to values. Certain technological innovations, made possible given conceptual or social change, might be judged destructive from a standpoint of world peace or health; musical works might be, though minimally creative, considered hackneyed or symptomatic of a sexist or racist worldview; creative properties as defined might be included as constitutive of a broader theory of artistic or scientific value. Our model does the work of marking these features. The value theorist of creativity can then use or judge these features from within whatever framework she chooses.

The model thus accomplishes the goals set in the preface. We have accommodated the spooks with a cold stare and flat feet, as naturalism should. We have avoided confusion by offering a conceptual analysis and then maintaining the results of that analysis throughout. We have avoided getting stuck in the mud by further narrowing the explanandum to the descriptive but general features of creative cognition. Our diagnosis and prognosis have thus repaid us well, motivating from the shadow of mystery, an informative naturalistic model of creative thought processes.

Future research

A topic like creativity, it should now be obvious, involves a complex of issues relevant to a variety of theoretical domains. We thus mention only three sets of issues for further research, with emphasis on how these issues are either informed or treated by our model and/or, conversely, how our model stands to benefit from considering them.

Issues in the philosophy of science: We have given a model of creativity understood as a psychological phenomenon. This model sheds interesting light upon and itself stands to be enlightened by consideration of the following two issues.

(a) *Scientific method as creative*: Traditionally, the story goes that artists are in the business of creating and scientists are in the business of discovering. Perhaps this is a wedge diminishing in purchase, but our model gives a fairly simple way of saying why. We simply look at the cognitive processes modeled and ask, do scientists think like that? Do their methods require cognitive manipulation; do they use imagination in the varied ways proposed; do they enjoy incubatory solutions or preparation? We need not

carefully answer these questions here. But on the face of it, the quick answers are 'yes' across the board. This gives us a reason to proceed as if science is a creative activity and then to analyze the features specific to that activity qua scientific activity.

(b) *Rationality of science*: There is a suite of complicated issues that concern the rational status of scientific method and theorizing. Since the publication of Thomas Kuhn's famous work on scientific revolutions, philosophers have debated the rationality of paradigm shifts, say from Newtonian to Einsteinian mechanics (Kuhn [1962] 1970). A related debate concerns what Michael Friedman calls *communicative rationality*. Even if Kuhn's defenses are correct and basic principles of scientific theorizing are sufficient to maintain some level of rationality across paradigm shifts, Friedman argues that any two paradigms will involve incommensurably different concepts and terminology, thus preventing rational communication between the two parties involved (Friedman 2001, 2002; see also Habermas 1984). The most pronounced form of scientific irrationalism is embodied in Paul Feyerabend's *Against Method* (1975). Feyerabend opposes rationalists like Popper and Lakatos, arguing that science, as a pursuit of knowledge, lacks any clear methodological rules or constraints that would render it a rational enterprise. If it is bound by any rule at all, it is the obviously trivial rule "anything goes."

These debates are like old trees. It is thus unlikely that our model suffices to sever any one branch or traces any one root. But here is a hunch: creative cognition as we have modeled it is rational. If scientific method involves, as a significant part, this kind of cognitive processing, then we have identified one portion of the scientific process that is rational. We then continue our analysis by determining how this fact speaks to other points of the rationality debates.

Issues in philosophy of mind and cognitive science: A theory of the psychology of creativity informs a broader theory of cognition. Here are two ways that our theory stands to do so, both of them abstractable from concerns specific to creativity.

(c) General importance of imagination: We argued in chapter 5 for the role of imagination in creative cognition. Creative thoughts require for their tokening some kind of non-truth bound manipulation. Given its inherent cognitive freedom, imagination serves this role uniquely well. But we can extract from this a more general lesson: much of cognition requires cognitive manipulation and thus much of cognition plausibly involves imagination. Traditionally, imagination is reserved for discussions of fantasy, fiction and creativity, and this is no less true of theoretical analysis than of folk idiom. But through identifying the importance of imagination to creativity, we see that imagination enables and/or constitutes much of mundane cognition, some of it creative some of it not (5.4). Novel cognitive states, be they beliefs, desires, concepts, or whatever, require for their formation the kind of cognitive manipulation that imagination affords. Sometimes these states satisfy the other conditions requisite for creativity, other times not. Either way, we see that imagination plays a much broader and more important role in human cognition than we might have thought. This is surely a fact that warrants further analysis by philosophers of mind and cognitive scientists.

(d) General importance of incubation: In chapter 6, we took seriously the tradition of including an incubated stage in creative processing. We succeeded in explaining this stage in naturalistic terms. But, once again, we can extract a more general point: incubated cognition is important not just for creative problem solving, but for much of mundane cognition. It enables us to perform a multiplicity of tasks, to learn new

information, and to recall relevant associations and memories. This, it was suggested, is amenable to a simple evolutionary explanation: we need some degree of automatized unconscious processing to learn to navigate a world that provides us with an overwhelming array of stimuli. This is a dramatically important lesson for the philosophy of mind, which has largely shrugged the very possibility of unconscious cognition. If our model is accurate, such cognition is not merely a possibility, it is a necessity.

Issues in evolutionary theory: The ability to think creatively possesses obvious importance for the success and survival of the human species. Evolutionary psychologists, anthropologists, and archeologists have begun to theorize this fact, with heavy emphasis on explanations of the *creativity explosion* (see 5.3). Here are two ways this set of concerns might benefit from a model like the one on offer.

(e) Creativity as psychological phenomenon: Much of the emphasis of evolutionary theory goes on social context and culture (Atran 1998; Dunbar 1996; Miller 2000; Mithen 1996; Tomasello 1999). This is likely as it should be, both for explanations of creativity and of other cognitive capacities and social phenomena. However, for evolutionary explanations of creativity, we might do well to set cultural issues aside and start with creativity understood as a psychological phenomenon. This suggestion is analogous to the one already made about starting with descriptive concerns and then moving on to evaluative concerns. The second might, as a matter of fact, be necessary to giving a complete theory of creativity, but that does not preclude us from starting with an analysis that is exclusive to the first. Our model encourages the same strategy in the

evolutionary domain. Start with a precisely defined and carefully modeled set of psychological phenomena and then situate them within one's evolutionary theory.

(f) Imagination as explanandum for selection pressures: This point immediately follows from the last. We have identified the importance of imagination to creativity and gestured at the importance of creativity to species survival and culture. We have also identified imagination as a cognitive capacity with more general import, serving the role of cognitive manipulator in all manner of cognition. This frames an explanandum that some have already taken to explaining, namely, the selection pressures responsible for the evolution of imaginative capacity (Carruthers 2001, and in preparation; Harris 2000). This is a very difficult story to tell, since it requires an explanation both of why we would, as a species, start imagining and why or how we would begin doing it in seriously minded or cognitively valuable ways. Our model lends some conceptual clarity to this project, since it identifies a functional role for cognitive manipulation and then shows why imagination, by its nature, is best suited to serve this role (see chapter 4 and 5).

We have left some things out, only hinted at some others, and probably spent too much time on others still. Nonetheless, these six sets of issues ((a) through (f)), parsed into three broad theoretical domains, are each of them sufficient for a book-length manuscript. Conjoining them with those we have neglected provides for entire volumes and several academic careers. The hope is that this thesis encourages philosophers and scientists to take on some of this explanatory work. Here are ways it might. The thesis identifies or at least provides a reminder of the import of creativity. It also provides an example of the amenability of the phenomenon to naturalistic analysis. Finally, the thesis

provides the best first step for taking on and continuing the heavy explanatory burden that surrounds creativity.

Bibliography

Alston, W. (1989) 'The Deontological Conception of Epistemic Justification,' in *Epistemic Justification*, Ithaca: Cornell University Press.

Amit, D. (1995) 'The Hebbian paradigm reintegrated: Local reverberations as internal representations,' *Behavioral and Brain Sciences* 18: 617-57.

Anderson, J.R. (1983) 'A spreading activation theory of memory' *Journal of Verbal Learning and Verbal Behavior* 22: 261-95.

Atran, S. (1998) 'Folk Biology and the Anthropology of Science,' *Behavioral and Brain Sciences* 21: 547-609.

Barbur, J.L., Ruddock K., and Waterfield, V. (1980) 'Human visual responses in the absence of the geniculo-calcarine projection,' *Brain* 102: 905-928.

Barsalou, L.W. (1999) 'Perceptual Symbol Systems,' *Behavioral and Brain Sciences* 22: 577-660.

_____and Prinz, J. (1997) 'Mundane Creativity in Perceptual Symbol Systems,' in T. Ward, S. Smith, and J. Vaid (eds) *Creative Thought: An Investigation of Conceptual Structures and Processes*, Washington, D.C., American Psychological Association.

_____(2002) 'Acquisition and Productivity in Perceptual Symbol Systems: An Account of Mundane Creativity,' in T. Dartnall (ed.) *Creativity, Cognition and Knowledge*, Westport, CT: Praeger.

Beardsley, M. (1958) *Aesthetics: Problems in the Philosophy of Criticism*, New York: Harcourt, Brace & World.

Bennett, J. (1984) "Deciding what to believe," in *A Study of Spinoza's Ethics*, Indianapolis: Hackett.

_____(1990) "Why is belief involuntary?" *Analysis* 50: 87-107.

Beringer, L. and Stein, K. (1930) 'Analyse eines Falles von "Reiner" Alexie,' *Zeitschrift fuer die Gesamte Neurologie und Psychiatrie* 123: 472-8.

Bisiach, E. and Luzzatti, C. (1978) "Unilateral Neglect or Representational Space, *Cortex* 14: 129-33.

Block, N. (1980a) 'Introduction: What is Functionalism?' in N. Block (ed.) *Readings in Philosophy of Psychology*, Vol. 1, Cambridge: Harvard University Press.

_____(1980b) 'Problems with Functionalism,' in N. Block (ed.) *Readings in Philosophy of Psychology*, Vol. 1, Cambridge: Harvard University Press.

_____(1995) 'On a Confusion About a Function of Consciousness,' *Behavioral and Brain Sciences* 18: 227-87.

Boden, M. (1999) 'Computer Models of Creativity,' in R. Sternberg (ed.) *Handbook of Creativity*, Cambridge: Cambridge University Press.

_____(2004) *The Creative Mind*, 2nd Edn., London: Routledge.

Bowers, K., Farvolden, P. and Mermigis, L. (1995) 'Intuitive Antecedents of Insight,' in S.M. Smith, T.B. Ward, and R.A. Finke (eds), The Creative Cognition Approach, Cambridge: MIT Press.

Boyer, P. (1994) *The Naturalness of Religious Ideas: A Cognitive Theory of Religion*, Berkeley: University of California Press.

Braitenberg, V. (1989) 'Some arguments for a theory of cell assemblies in the cerebral cortex,' in L. Nadel, L. Cooper, P. Culicover and R. Harnish (eds) *Neural connections, mental computation*, Cambridge: MIT Press.

Brent, P., Kennard, C., and Ruddock, K. (1994) 'Residual colour vision in a human hemianope: spectral responses and colour discrimination,' *Proceedings of the Royal Society of London: Series B: Biological Sciences*, 256: 219-25.

Brewer, W. (1999) *Perception and Reason*, Oxford: Oxford University Press.

Bridgeman, B. and Staggs, D. (1982) 'Plasticity in human blindsight,' *Vision Research* 22: 1199-1203.

Brooks, L. (1968) 'Spatial and Verbal Components of the Act of Recall,' *Canadian Journal of Psychology* 22: 349-68.

Buchanan, B., Smith, D., White, W., Gritter, R., Feigenbaum, E., Ledgerberg, J., Djerassi, C. (1976) 'Applications of Artificial Intelligence for Chemical Inference: XXII Automatic Rule Formation in Mass Spectrometry by Means of the Meta-Dendral Program,' *Journal of the American Chemistry Society* 98: 6168-78.

Buckner, R., Peterson, S., Ojemann, J., Miezin, F., Squire, L., Raichle, M. (1995) 'Functional anatomical studies of explicit and implicit memory retrieval tasks,' *Journal of Neuroscience* 15: 12-29.

Cacciari, C. and Levorato, M. (1989) 'How children understand idioms in discourse,' *Journal of Child Language* 16: 387-405.

Cacciari, C., Levorato, M., and Cicogna, P. (1997) 'Imagination at Work: Conceptual and Linguistic Creativity in Children,' in T. Ward, S. Smith, and J. Vaid (eds) *Creative Thought: An Investigation of Conceptual Structures and Processes*, ed. Washington, D.C., American Psychological Association.

Calvin, W. (1989) *The Cerebral Symphony: Seashore Reflections on the Structure of Consciousness*, New York: Bantam.

Campbell, D. (1960) 'Blind variation and selective retention in creative thoughts as in other knowledge processes,' *Psychological Review* 67: 380-400.

_____(1965) 'Variation and selective retention in socio-cultural evolution,' in H. Barringer, G. Blanksten, R. Mack (eds) *Social Change in Developing Areas*, Cambridge: Schenkman.

Carruthers, P. (in preparation) *The Architecture of the Mind.*

_____(2002) 'Human Creativity: Its Cognitive Basis, its Evolution, and its Connections with Childhood Pretence,' *British Journal for the Philosophy of Science* 53: 225-49.

Chalmers, David. (2002) 'Does Conceivability Entail Possibility?' in T. Gendler and J. Hawthorne (eds) *Conceivability and Possibility*, Oxford: Oxford University Press.

Changeux, J.P. and Danchin, A. (1976) 'Selective stabilisation of developing synapses as a mechanism for the specification of neuronal networks,' *Nature* 264: 705-12.

Changeux, J.P. and Dehaene, S. (1989) 'Neuronal models of cognitive function,' *Cognition* 33: 63-109.

Churchland, P. (1999) *Matter and Consciousness*, 9th Edn., Cambridge: MIT Press.

Cohen, L.J. (1989) 'Belief and Acceptance,' *Mind* 98: 367-369.

_____(1992) *An Essay on Belief and Acceptance*, Oxford: Oxford University Press.

Collins, A., and Loftus, E. (1975) 'A Spreading Activation Theory of Semantic Processing,' *Psychological Review* 82: 407-28.
Coslett, H., and Saffran, E. (1989) 'Evidence for preserved reading in "pure alexia,"' *Brain* 112: 327-59.
_____(1992) 'Optic Aphasia and the Right Hemisphere: A Replication and Extension,' *Brain and Language* 43: 148-61.
Craver-Lemley, C. and Reeves, A. (1992) 'How Visual Imagery Interferes With Vision,' *Psychological Review* 99: 633-49.
Crimmins, M. (1992) *Talk About Beliefs*, Cambridge: M.I.T. Press.

Cummins, R. (1975) 'Functional Analysis,' *Journal of Philosophy* 72: 741-65.

Currie, G. (1989) *An Ontology of Art*, London: Macmillan.
_____(1990) *The Nature of Fiction*, Cambridge: Cambridge University Press.
_____(1995) 'Imagination and Simulation: Aesthetics Meets Cognitive Science,' in M. Davies and T. Stone (eds) *Folk Psychology*, Oxford: Blackwell.
_____(1996) 'Simulation-theory, Theory-theory, and the Evidence From Autism,' in P. Carruthers and S. Smith (eds) *Theories of Theories of Mind*, Cambridge: Cambridge University Press.
_____(1997) 'The Paradox of Caring: Fiction and the Philosophy of Mind', in M. Hjort and S. Laver (eds) *Emotion and the Arts*, Oxford: Oxford University Press.
_____(2002a) 'Desire in Imagination,' in T. Gendler and J. Hawthorne (eds) *Conceivability and Possibility*, Oxford: Oxford University Press.
_____(2002b) 'Imagination as Motivation,' *Proceedings of the Aristotelian Society* 102.3: 201-216.
_____and I. Ravenscroft (2002) *Recreative Minds*, Oxford: Oxford University Press.

Damasio, A. (1999) *The Feeling of what Happens: Body and Emotion in the Making of Consciousness*,New York: Harcourt Brace and Co.
Davies, M. (1995) 'Reply to Searle: Consciousness and the Varieties of Aboutness,' in C. Macdonald and G. Macdonald (eds) *Philosophy of Psychology*, Oxford: Blackwell.
_____and Stone, T. (eds) (1995a) *Folk Psychology*, Oxford: Blackwell.
_____(eds)(1995b) *Mental Simulation*, Oxford: Blackwell.

De Haan, E. (2001) 'Face perception and recognition,' in B. Rapp (ed.) *The Handbook of Cognitive Neuropsychology*, Philadelphia: Psychology Press-Taylor & Francis.
_____and Young, A., and Newcombe, F. (1987a) 'Face recognition without awareness,' *Cognitive Neuropsychology* 4: 385-415.
_____(1987b) 'Faces interfere with name classification in a prosopagnosic patient,' *Cortex* 23: 309-16.
De Sousa, R. (1987) *The Rationality of Emotion*, Cambridge: MIT Press.

Dennett, D. (1988) 'Quining Qualia,' in Marcel, A. and Bisiach, E. (eds) *Consciousness in Modern Science*, Oxford: Oxford University Press.
_____(1991) *Consciousness Explained*, Boston: Little, Brown & Company.

Desimone, R. (1996) 'Neural mechanisms for visual memory and their role in attention,' *Proceedings of the National Academy of Sciences of the United States of America* 93: 13494-99.
Donald, M. (2001) *A Mind So Rare: The Evolution of Human Consciousness*, London: Norton.
Dretske, F. (1969) *Seeing and Knowing*, Chicago: University of Chicago Press.

Dunbar, R. (1996) *Grooming, Gossip, and the Evolution of Language*, Cambridge: Harvard University Press.
Durup, G., & Fessard, A. (1935) 'L'ectroencephalogramme de l'homme,' *Annual Psychology* 36: 1-32.
Edelman, G. (1987) *Neural Darwinism*, New York: Basic Books.
_____and Tononi, G. (2001) *A Universe of Consciousness*, New York: Basic Books.

Farah, M. (1990) *Visual Agnosia*, Cambridge: MIT Press.
_____(2001) 'Consciousness' in B. Rapp (ed.) *The Handbook of Cognitive Neuropsychology*, Philadelphia: Psychology Press-Taylor & Francis.
Fentress, J. (1999) 'The Organization of Behaviour Revisited,' *Canadian Journal of Experimental Psychology* 53: 8-19.
Feyerabend, P. (1975) *Against Method*, London: New Left Books.

Findlay, A. (1965) *A Hundred Years of Chemistry*, 3rd Edn., ed. T.I. Williams, London: Duckworth.
Finke, R., and Pinker, S., (1982) 'Spontaneous imagery scanning in mental extrapolation' *Journal of Experimental Psychology: Learning, Memory, and Cognition* 8: 142-7
_____(1983) 'Directional scanning of remembered visual patterns,' *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9: 398-410.
Finke, R., Ward, T., and Smith, S. (1992) *Creative Cognition*, Cambridge: MIT Press.

Flanagan, O. (1991) *The Science of the Mind*, 2nd Edn., Cambridge: MIT Press.

Freud, S. (1940) *An Outline of Psychoanalysis*, trans. by J. Strachey, New York: Norton.

Friedman, M. (2001) *Dynamics of Reason: The 1999 Kant Lectures at Stanford University*, Stanford: The Center for the Study of Language and Information.
_____(2002) 'Kant, Kuhn, and the Rationality of Science,' *Philosophy of Science* 69: 171-90.

Funkhouser, E. and Spaulding, S. (in preparation) 'Imagination and Other Scripts'

Gabora, L. (2000) 'Toward a theory of creative inklings' in R. Ascott (ed.) *Art, Technology, and Consciousness*, Oxford: Intellect Press.

_____(2002) 'Cognitive mechanisms underlying the creative process' in T. Hewett and T. Kavanagh (eds) *Proceedings of the Fourth International Conference on Creativity and Cognition*, 126-133.
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.
Gendler (2000) 'The Puzzle of Imaginative Resistance' *Journal of Philosophy* 2: 55-81.
_____(2003) 'On the Relation Between Pretense and Belief,' in D. Lopes and M. Kieran (eds) *Imagination, Philosophy and the Arts*, London: Routledge.
Gettier, E. (1963) 'Is Justified True Belief Knowledge?' *Analysis* 23: 121-3.

Gibbs, W. (1987) 'Linguistic Factors in Children's Understanding of Idioms,' *Journal of Child Language* 14: 569-586.
_____(1991) 'Semantic Analyzability in Children's Understanding of Idioms,' *Journal of Speech and Hearing Research* 34: 613-620.
_____(1994) *The Poetics of the Mind: Figurative Thought, Language, and Understanding*, Cambridge: Cambridge University Press.
Goldman, A. (1989) 'Interpretation psychologized,' *Mind and Language* 4: 161-85.
_____(1992) 'In defense of simulation theory,' *Mind and Language* 7: 104-19.

Goodale, M. , Milner, A., Jakobson, L, and Carey, D. (1991) 'A neurological dissociation between perceiving objects and grasping them,' *Nature* 349: 154-6.
Gordon, R. (1986) 'Folk psychology as simulation,' *Mind and Language* 1: 158-70.

Habermas, J. (1984) *The Theory of Communicative Action*, Boston: Beacon.

Hadamard, J. (1954) *An Essay on the Psychology of Invention in the Mathematical Field*, New York: Dover.
Haier, R. Siegel, B.V. Jr., MacLachlan, E., Soderling, S., Lottenberg, S. and Buchsbaum, M. (1992) 'Regional glucose metabolic changes after learning a complex visuospatial/motor task: A positron emission tomographic study," *Brain Research* 570: 134-43.
Hampson, P.J. and Duffy, C. (1984) 'Verbal and Spatial Interference Effects in Congenitally Blind and Sighted Subjects,' *Canadian Journal of Psychology* 38: 411-20.
Harris, P. (2000) *The Work of the Imagination*, Oxford: Blackwell.

Heal, J. (1995) 'How to Think About Thinking,' in M. Davies and T. Stone (eds) *Mental Simulation*, Oxford: Blackwell.
_____(1996a) 'Simulation and Cognitive Penetrability,' *Mind and Language* 11: 44-67.
_____(1996b) 'Simulation, Theory and Content,' in P. Carruthers and P. Smith (eds) *Theories of Theories of Mind*, Cambridge: Cambridge University Press.
Hebb, D. (1949) *The Organization of Behavior*, New York: Wiley.

Helmholtz, H. von (1896) *Vorträge und Reden*, Brunswick: Friedrich Viewig und Sohn.

Hinton, G., and Anderson J. (eds)(1981) , *Parallel Models of Associative Memory*, Hillsdale, NJ: Lawrence Earlbaum Associates.

Hinton, G., McClelland, J.L. & Rummelhart, D.E. (1986) 'Distributed representations' in D.E. Rummelhart and J. L. McClelland (eds) *Parallel Distributed Processing: Explorations in the Microstructure Cognition*, Cambridge: MIT Press.

Hofstadter, D. and the Fluid Analogies Research Group (FARG) (1995) *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*, New York: Basic Books.

Hopfield, J. and Tank D. (1986) 'Computing with neural circuits: A model,' *Science* 233: 625-33.

Horgan, T. and Tienson, J. (2002) 'The Intentionality of Phenomenology and the Phenomenology of Intentionality,' in D. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings*, Oxford: Oxford University Press.

Hull, C. (1943) *Principles of Behavior*, New York: Appelton-Century-Crofts.

Humberstone, I.L. (1992) 'Direction of Fit,' *Mind* 101: 59-83

Hume, David. (1987) *A Treatise of Human Nature*, 2nd Edn., ed. P.H. Nidditch, Oxford: Oxford University Press.

Jackendoff, R. (1987) *Consciousness and the Computational Mind*, Cambridge: MIT Press.

Jacobs, D., Shuren, J., Bowers, D., and Heilman, K. (1995) 'Emotional Facial Imagery, Perception, and Expression in Parkinson's Disease,' *Neurology* 45: 1696-702.

James, W. (1890) *The Principles of Psychology*, Vol. 1, New York: Holt.

_____(1892) *Psychology: The briefer course*. New York: Holt.

John, E. and Killiam, K. (1959) 'Electrophsyiological Correlates of Avoidance Conditioning in the Cat,' *Journal of Pharmacological and Experimental Therapeutics* 125: 252.

Kahneman, D., & Treisman, A. (1984) 'Changing views of attention and automaticity' in R. Parasuraman and D. R. Davies (eds), *Varieties of attention*, New York: Academic Press.

Kanerva, P. (1988) *Sparse Distributed Memory*, Cambridge: MIT Press.

Kant, I. (1987) *The Critique of Judgement*, trans. W.S. Pluhar, Indianapolis: Hackett.

Kaplan, D. (1989) 'Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals,' in J. Almog, J. Perry, H. Wettstein (eds) *Themes From Kaplan*, Oxford: Oxford University Press.

Kaplan, S., Sonntag, M., and Chown, E. (1991) 'Tracing recurrent activity in cognitive elements (TRACE): A model of temporal dynamics in a cell assembly,' *Connection Science* 3: 179-206.

Karmiloff-Smith, A. (1990) 'Constraints on representational change,' *Cognition* 34: 57-83.

_____(1992) *Beyond Modularity*, Cambridge: MIT Press.

Karni, A., Meyer, G., Jezzard, P., Adams, M. M., Turner, R., & Ungerleider, L. G. (1995) 'Functional MRI evidence for adult motor cortex plasticity during motor skill learning,' *Nature,* 377: 155–158.

Kim, J. (2005) *Physicalism, Or Something Near Enough*, Princeton: Princeton University Press.

Kivy, P. (2001) *The Possessor and the Possessed*, New Haven: Yale University Press.

Klein, R. (1999) 'Donald O. Hebb,' in R. Wilson and F. Keil (eds) *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge: MIT Press.

Koestler, A. (1964) *The Act of Creation*, New York: Macmillan.

Kosslyn, S. (1994) *Image and Brain: The Resolution of the Imagery Debate*, Cambridge: MIT Press.

_____(1981) 'The medium and the message in mental imagery: A theory,' in N. Block (ed.) *Imagery*, Cambridge: MIT Press.

_____(1980) *Image and Mind*, Cambridge: Harvard University Press.

_____and Pomerantz, J.R. (1978) 'Imagery, propositions, and the form of internal representations,' in N. Block (ed.) *Readings in Philosophy of Psychology*, Vol. 2, Cambridge: Harvard University Press.

_____and Ball, T., and Reiser, B. (1978) 'Visual images preserve metric spatial information,' *Journal of Experimental Psychology: Perception and Performance* 4: 47-60

Kripke, S. (1980) *Naming and Necessity*, Cambridge: Harvard University Press.

Kris, E. (1952) *Psychoanalytic Explorations in Art*, New York: International Universities Press.

Kuhn, T. ([1962]1970) *The Structure of Scientific Revolutions*, 2nd Ed., Chicago: University of Chicago Press.

Lamarque, P. and Olsen, S.H. (1994) *Truth, Fiction, and Literature: A Philosophical Perspective*, Oxford: Oxford University Press.

Langley, P., Simon, H., Bradshaw, G., Zytkow, J. (1987) *Scientific Discovery: Computational Explorations of the Creative Process*, Cambridge: MIT Press.

Lashley, K. (1950) 'In search of the engram,' *Symposia of the Society for Experimental Biology* 4: 454-82.

Levorato, M. and Cacciari, C. (1992) 'Children's Comprehension and Production of Idioms: The Role of Context and Familiarity,' *Journal of Child Language* 19: 415-33.

_____(1995) 'The Effects of Different Tasks on the Comprehension and Production of Idioms in Children,' *Journal of Experimental Child Psychology* 60: 261-83.

Lewis, D. (1970) 'How to Define Theoretical Terms,' *Journal of Philosophy* 67: 427-446.

_____(1972) 'Psychophysical and Theoretical Identifications,' *Australasian Journal of Philosophy* 50: 249-58.

_____(1979) 'Attitudes De Dicto and De Se,' *The Philosophical Review* 88: 513-43.
_____(1989) 'Dispositional Theories of Value,' *Proceedings of the Aristotelian Society* 63: 113-37.
_____(1996) 'Elusive Knowledge,' *The Australasian Journal of Philosophy* 74: 549-67.

Lindsay, R. Buchanan, B., Feigenbaum, E., and Lederberg, J. (1980) *Application of Artificial Intelligence for Chemistry: The DENDRAL Project*, New York: McGraw-Hill.
Loar, B. (2003a) 'Phenomenal Intentionality as the Basis of Mental Content,' in M. Hahn and B. Ramberg (eds) *Reflections and Replies: Essays on Tyler Burge*, Cambridge: MIT Press.
_____(2003b) 'Transparent Experience and the Availability of Qualia,' in Q. Smith and A. Jokic (eds) *Consciousness: New Philosophical Perspectives*, Oxford: Oxford University Press.
Lopes, D. (2003) 'Out of sight, Out of Mind,' in D. Lopes and M. Kieran (eds) *Imagination, Philosophy and the Arts*, London: Routledge.
Marcel, A. (1998) 'Blindsight and shape perception: Deficit of visual consciousness of visual function?' *Brain* 121: 1565-88.
Marr, D. (1969), 'A theory of cerebral cortex', *Journal of Physiology* 202: 437-470.
_____(1976) 'Early processing of visual information' *Philosophical Transactions of the Royal Society* 275: 483-524.
_____(1982). *Vision*, San Francisco: W.H. Freeman.

Martindale, C. (1977) 'Creativity, consciousness, and cortical arousal,' *Journal of Altered States of Consciousness* 3: 69-87.
_____(1981) *Cognition and Consciousness*, Homewood, IL: Dorsey Press.
_____(1989) 'Personality, Situation, and Creativity,' in J. Glover, R. Ronning, and C. Reynolds (eds) *Handbook of Creativity*, New York: Plenum.
_____(1990) 'Creative Imagination and Neural Activity,' in R. Kunzendorf and A. Sheikh (eds) *The Psychophysiology of Mental Imagery*, New York: Baywood.
_____(1995) 'Creativity and Connectionism,' in S. Smith, T.B. Ward, and R. Finke (eds), *The Creative Cognition Approach*, Cambridge: MIT Press.
_____(1999) 'Biological bases of creativity,' n R. Sternberg (ed.) *Handbook of creativity*, Cambridge: Cambridge University Press.
_____and Armstrong, J. (1974) 'The relationship of creativity to cortical activation and its operant control,' *Journal of Genetic Psychology* 124: 311-320.
_____and Hines, D. (1975) 'Creativity and cortical activation during creative, intellectual and EEG feedback tasks,' *Biological Psychology* 3: 71-80.
Maudsley, H. (1876) *The Physiology of the Mind: Being the First Part of a 3rd Edition, Revised, Enlarged, and in Great Part Rewritten, of "The Physiology of Mind"*, London: Macmillan.
Matravers, D. (2003) 'Fictional assent and the (so-called) puzzle of imaginative resistance,' in D. Lopes and M. Kieran (eds) *Imagination, Philosophy and the Arts*, London: Routledge.
Matthen, M. (2005) *Seeing, Doing, and Knowing*, Oxford: Oxford University Press.

McDowell, J. (1994) *Mind and World*, Cambridge: Harvard University Press.

Mednick, S. (1962) 'The associative basis of creative processes,' *Psychological Review* 69: 220-32.
Mellars, P. (1973) 'The character of Middle-Upper transition in southwest France,' in C. Renfrew (ed.) *The Explanation of Culture Change*, London: Duckworth.
_____(1989) 'Technological changes at the Middle-Upper Paleolithic transition: economic, social, and cognitive perspectives,' in P. Mellars and C. Stringer (eds) *The Human Revolution*, Edinburgh: Edinburgh University Press.
Mendolsohn, G. (1976) 'Associative and attentional processes in creative performance,' *Journal of Personality* 44: 341-69.
Metzler, J. & R. N. Shepard (1971) 'Mental rotation of three-dimensional objects,' *Science* 171: 701-3.
Miller, G. (2000) *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*, London: Heinemann.
Miller, R. and Wickens, J. (1991) 'Corticostriatal cell assemblies in selective attention and in representation of predictable and controllable events. A general statement of corticostriatal interplay and the role of striatal dopamine,' *Concepts in Neuroscience* 2: 65-94.
Millgram, E. (1997) 'Deciding to Desire,' in *Practical Induction*, Cambridge: Harvard University Press.
Milner, A.D. et al (1991) 'Perception and action in "visual form agnosia,"' *Brain* 114: 405-28.
_____and Goodale, M. (1995) *The Visual Brain in Action*, Oxford: Oxford Science Publications.
Milner, P. (1957) 'The cell assembly: Mk. II.' *Psychological Review* 64: 242-52.
_____(1999) *The Autonomous Brain*, London: Lawrence Erlbaum.

Mishkin, M. (1993) 'Cerebral memory circuits,' in T. Poggio and D. Glaser (eds) *Exploring Brain Function: Models in Neuroscience*, New York: Wiley.
Mithen, S. (1996) *The Prehistory of the Mind*, New York: Thames and Hudson.

Nagel, T. (1979) 'What is it like to be a bat?' in *Mortal Questions*, Cambridge: Cambridge University Press.
Nichols, S. and Stich, S. (2000) 'A Cognitive Theory of Pretence,' *Cognition* 74: 115-47.
_____(2003) *Mindreading*, Oxford: Oxford University Press.

Norman, D. A., & Shallice, T. (1986) 'Attention to action: Willed and automatic control of behavior,' in R. J. Davidson, G. E. Schwartz, & D. Shapiro (eds), *Consciousness and self-regulation*, Vol. 4, New York: Plenum.
Novitz, D. (1999) 'Creativity and Constraint,' *Australasian Journal of Philosophy* 77: 67-82.
____(2003) 'Explanations of creativity,' in B. Gaut and P. Livingston (eds) *The Creation of Art*, Cambridge: Cambridge University Press.
Nozick, R. (1981) *Philosophical Explanations*, Cambridge: Harvard University Press.

Nudo, R. J., Milliken, G. W., Jenkins, W. M., & Merzenich, M. M. (1996) 'Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys,' *Journal of Neuroscience* 16: 785–807.

Oatley, K. (1992) *Best Laid Schemes: The Psychology of Emotions*, Cambridge: Cambridge University Press.

_____and Johnson-Laird, P. (1987) 'Towards a cognitive theory of the emotions,' *Cognition and Emotion* 1: 29-50.

Olton, R. (1979) 'Experimental studies of incubation: Searching for the elusive,' *Journal of Creative Behavior* 13: 9-22.

Orbach, J. (1998) 'Hebb's *Organization of Behavior* 50 Years After Publication,' in Orbach, J., *The Neuropsychological Theories of Lashley and Hebb*, New York: University Press of America.

Palm, G. (1982) *Neural Assemblies: An Alternative Approach to Artificial Intelligence*, Berlin: Springer-Verlag.

Partridge, D. and Rowe, J. (2002) 'Creativity: A Computational Modeling Approach,' in T. Dartnall (ed.) *Creativity, Cognition and Knowledge*, Westport, CT: Praeger.

Peirce, C.S. (1935 [1898]) *Collected Papers-Vol. 6: Scientific Metaphysics*, ed. C. Harshorne and P. Weiss, Cambridge: Harvard University Press.

Peterson, S., vanMier, H. Fiez, J., and Raichle, M. (1998) 'The effects of practice on the functional anatomy of task performance,' *Proceedings of the National Academy of Sciences of the United States of America* 95: 853-60.

Pettigrew, J. (1974) 'The effect of visual experience on the development of stimulus specificity by kitten cortical neurons,' *Journal of Physiology* 237: 49-74.

Pigarev, I., Northdurft, H., and Kastner, S. (1997) 'Evidence for asynchronous development of sleep in cortical areas,' *Neuroreport* 8: 2557-60.

Plato (1989) *Ion*, *Republic*, both in E. Hamilton and H. Cairns (eds) *Plato: Collected Dialogues*, New Jersey: Princeton University Press.

Poincaré, H. ([1902,1904, 1908] 1982) *Science and Hypothesis*, *The Value of Science*, *Science and Method*, all repr. in *The Foundations of Science*, Washington, D.C.: University Press of America.

Polanyi, M. (1981) 'The Creative Imagination,' in D. Dutton and M. Krausz (eds), *The Concept of Creativity in Science and Art*, Boston: Martinus Nijhoff.

Pollock, J. (2000) 'Epistemic Norms' in E. Sosa and J. Kim (eds) *Epistemology: An Anthology*, Oxford: Blackwell; repr. from J. Pollock (1986) *Contemporary Theories of Knowledge*, Lanham: Rowman and Littlefield.

Poppel, E., Held, R., and Frost, D., (1973) 'Residual visual function after brain wounds involving the central visual pathways in man,' *Nature* 243: 295-6.

Popper, K. ([1934]1959) *The Logic of Scientific Discovery*, New York: Harper and Row.

Posner, M. (1978) *Chronometric Explorations of Mind*, Hillsdale, NJ: Erlbaum.

_____and Raichle, M. (1994) *Images of Mind*, New York: Scientific American Library.

_____and DeGirolamo, G., and Fernandez-Duque, D. (1997) 'Brain Mechanisms of Cognitive Skills,' *Cognition* 6: 267-90.

_____and P. Bourke (2001) 'Attention,' in J. Beaumont, P. Kenaealy, and M. Rogers (eds) *The Blackwell Dictionary of Neuropsychology*, Oxford: Blackwell.

Prinz, J. (2004) *Gut Reactions: A Perceptual Theory of Emotion*, Oxford: Oxford University Press.

Pylyshyn, Z. (1978) 'Imagery and artificial intelligence,' in W. Savage (ed.) *Perception and Cognition, Minnesota Studies in the Philosophy of Science*, Vol. 9, Minneapolis: University of Minnesota Press.

_____(1980) 'Computation and Cognition' *The Behavioral and Brain Sciences* 3: 111-69

_____(1981a) 'Imagery and Artificial Intelligence,' in N. Block (ed.) *Readings in Philosophy of Psychology*, vol. 2, Cambridge: Harvard University Press.

____(1981b) 'The imagery debate: Analog media versus tacit knowledge,' in N. Block (ed.), *Imagery*, Cambridge: MIT Press.

____(2002) 'Mental Imagery: In Search of a Theory,' *Behavioral and Brain Sciences* 25: 157-182.

Putnam, H. (1979) 'The "Corroboration" of Theories,' in T. Honderich and M. Burnyeat (eds) *Philosophy As It Is*, New York: Penguin.

Ramachandran, V.S. (1998) *Phantoms in the Brain*, New York: Quill Press.

Ramsey, F. ([1929]1978) 'Theories', in D. H. Mellor (ed.), *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, London: RKP.

Reisberg, D., Culver, L., Heuer, F., and Fischman, D. (1986) 'Visual Memory: When Imagery Vividness Makes a Difference,' *Journal of Mental Imagery* 10: 51-74.

Riddoch, G. (1917) 'Dissociation of visual perception due to occipital injuries, with especial reference to the appreciation of movement,' *Brain* 40: 15-57.

Robinson, J. (2005) *Deeper Than Reason*, Oxford: Oxford University Press.

Rosenblatt, R. (1961) *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*, Washington D.C., Spartan Books.

Rumelhart, D., Hinton, G., and McClelland, J. (1986) 'A general framework for parallel distributed processing,' in D. Rumelhart, and J. McClelland (eds) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition,* Vol. 1, Cambridge: MIT Press.

Ryle, G. (1949) *The Concept of Mind*, New York: Barnes and Noble.

Sakurai, Y. (1999) 'How do cell assemblies encode information in the brain?' *Neuroscience and Biobehavioral Reviews* 23: 785-96.

Salmon, N. (1986) *Frege's Puzzle*, Cambridge: M.I.T. Press.

Schopenhauer, A. (1958) *The World as Will and Representation*, trans. E.F.J. Payne, Indian Hills, CO: Falcon Wing's Press.

Scruton, R. (1974) *Art and Imagination: A Study in the Philosophy of Mind*, London: Methuen.

Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*, New York: Cambridge University Press.

Segal, S. (1970) 'Imagery and Reality: Can they be Distinguished?' in W. Keup (ed.) *Origins and Mechanisms of Hallucinations*, New York: Plenum Press.

_____and Fusella, V. (1970) 'Influences of Imaged Pictures and Sounds on Detection of Visual and Auditory Studies,' *Journal of Experimental Psychology* 83: 458-64.

Semon, R. (1921) *The Mneme*, London: Allen and Unwin.

Sejnowski, T. and Rosenberg, C. (1986) *NETtalk: A parallel network that learns to read aloud*, John Hopkins University Technical Report.
Seung, H.S. (2000) 'Half a century of Hebb,' *Nature Neuroscience* 3: 1166.

Shallice, T., and Saffran, E. (1986) 'Lexical processing in the absence of explicit word identification: Evidence from a letter-by-letter reader,' *Cognitive Neuropsychology* 3: 429-58.
Simchen, O. (2004) 'On the Impossibility of Nonactual Epistemic Possibilities,' *The Journal of Philosophy* 101: 527-554.
Simonton, D. (1999) *Origins of Genius*, Oxford: Oxford University Press.

Smith, S. (1995) 'Fixation, Incubation, and Insight,' in S. Smith, T.B. Ward, and R. Finke (eds), *The Creative Cognition Approach*, Cambridge: MIT Press.
_____and Blankenship, S. (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.
_____and T.B. Ward, and R. Finke (eds)(1995), *The Creative Cognition Approach*, Cambridge: MIT Press.
Spatz, H.C. (1996) 'Hebb's concept of synaptic plasticity and neuronal cell assemblies,' *Behavioural Brain Research* 78: 3-7.
Sterelny, K. (1986) 'The Imagery Debate,' *Philosophy of Science* 53: 560-83.

Sternberg, R. and Lubart, T. (1991) 'An investment theory of creativity and its development, *Human Development* 34:1-31.
_____(1992) 'By low and sell high: An investment approach to creativity,' *Current Directions in Psychological Science* 1: 1-5.
_____(1999) 'The Concept of Creativity: Prospects and Paradigms,' in R. Sternberg (ed.) *Handbook of creativity*, Cambridge: Cambridge University Press.
Strawson, G. (1994) *Mental Reality*, Cambridge: MIT Press.

Stringer, C. and Gamble, C. (1993) *In Search of the Neanderthals: Solving the Puzzle of Human Origins*, London: Thames and Hudson.
Taylor, K. (1998) *Truth & Meaning*, Oxford: Blackwell.

Thornton, C. (2002) 'Creativity and Runaway Learning,' in T. Dartnall (ed.) *Creativity, Cognition and Knowledge*, Westport, CT: Praeger.
Tomasello, M. (1999) *The Cultural Origins of Human Cognition*, Cambridge: Harvard University Press.
Tomita, Y. (1998) 'The Sources of J. S. Bach's Well-Tempered Clavier II in Vienna, 1777-1801,' Bach: 29/2: 8-79.
_____(1996) 'The Well-Tempered Clavier, Book 1' http://www.music.qub.ac.uk/~tomita/essay/wtc1.html.

Tye, M. (1991) *The Imagery Debate*, Cambridge: MIT Press.

Umiltà, C. (2001) 'Mechanisms of Attention,' in B. Rapp (ed.) *The Handbook of Cognitive Neuropsychology*, Philadelphia: Psychology Press-Taylor & Francis.

Van Inwagen, P. (1998) 'Modal Epistemology,' *Philosophical Studies* 92: 67-84.

Velleman, D. (2000) *The Possibility of Practical Reason*, Oxford: Oxford University Press.

Von der Malsburg, C. (1973) 'Self-organization of orientation sensitive cells in striate cortex,' *Kybernetick* 14: 85-100.

Wallas, G. (1926) *The Art of Thought*, New York: Harcourt, Brace, and World.

Walton, Kendall (1994) 'Morals in Fiction and Fictional Morality/I,' *Proceedings of the Aristotelian Society* 68: 27-50

_____(1990) *Mimesis as Make-Believe*, Cambridge: Harvard University Press.

Ward, T.B. (1994) 'Structured Imagination: The role of category structure in exemplar generation,' *Cognitive Psychology* 27: 1-40.

_____(1995) 'What's old about new ideas?' in S.M. Smith, T.B. Ward, and R.A. Finke (eds), *The Creative Cognition Approach*, Cambridge: MIT Press.

Weisberg, R. (1986) *Creativity: Genius and Other Myths*, New York: W.H. Freeman.

_____(1995) 'Case Studies in Creative Thinking: Reproduction versus Restructuring in the Real World,' in S. Smith, T.B. Ward, and R. Finke (eds), *The Creative Cognition Approach*, Cambridge: MIT Press.

_____(1999) 'Creativity and Knowledge: A Challenge to Theories,' in R. Sternberg (ed.) *Handbook of Creativity*, Cambridge: Cambridge University Press.

Weiskrantz, (1986) *Blindsight: A Case Study and Implications*, Oxford: Oxford University Press.

_____and Warrington, E., Sanders, M., and Marshall, J. (1974) 'Visual capacity in the hemianopic visual field following a restricted occipital ablation,' *Brain* 97: 709-28.

White, R. (1982) 'Rethinking the Middle/Upper Paleolithic transition,' *Current Anthropology* 23: 169-92.

_____(1993) 'Technological and social dimensions of "Aurignacian-Age" body ornaments across Europe,' in H. Knecht, A. Pike-Tay, and R. White (eds) *Before Lascaux: The Complete Record of the Early Upper Paleolithic*, Boca Raton: CRC Press.

Williams, B. (1973) 'Deciding to believe,' in *The Problems of the Self*, Cambridge: Cambridge University Press.

Willshaw, D.J. (1981) 'Holography, associative memory, and inductive generalization' in G.E. Hinton and J.A. Anderson (eds), *Parallel Models of Associative Memory*, Hillsdale, NJ: Lawrence Earlbaum Associates.

Woodworth, R. and Schlosberg, H. (1954) *Experimental Psychology*, New York: Holt, Rinehart, and Winston.

Yablo, S. (2002) 'Coulda, Woulda, Shoulda,' in T. Gendler and J. Hawthorne (eds) *Conceivability and Possibility*, Oxford: Oxford University Press.

_____(1993) 'Is Conceivability a Guide to Possibility?' *Philosophy and Phenomenological Research* 53:1-42.

Young, J.Z. (1951) Ferrier Lectures, *Proceedings of the Royal Society of London, B,* 139: 18.