CHAPTER

1

Domain Specificity: Introduction and Overview

Summary: This chapter outlines the issue of the domain specificity/domain generality of creativity, comparing it to similar controversies in intelligence. It argues that although people generally think of creativity in a domain-general way, our intuitions, when guided by the right questions, actually suggest a much more domain-specific view. The goal of this chapter is not to convince the reader that domain specificity is the correct theory but simply to introduce the controversy, break down some seemingly commonsensical (but incorrect) biases in the ways we tend to think about creativity, and introduce the kinds of tools needed in creativity research to make reasonable judgments about domain specificity and generality.

Can one use the same set skills, the same aptitudes, and the same abilities to do creative things in very different domains? Can one apply one's creativity in writing poetry, playing the piano, or glazing pottery to cooking, chemistry, or chess in ways that will result in more interesting and delicious recipes, more original theories and experimental designs, or more innovative ways to checkmate one's opponents? Is there a way of thinking or approaching problems that will lead to creative outcomes no matter the field in which one chooses to apply them? Is there a personality type that results in creativity in the arts, sciences, human relations, or anywhere else that creativity matters? These are the questions this book attempts to answer. Then, having answered those questions as far as current research can take us (which is rather far – the answers are surprisingly clear), later chapters explain what those answers mean for creativity research, creativity theory, creativity testing, and creativity training. The things one needs to know to be a competent

Domain Specificity of Creativity. http://dx.doi.org/10.1016/B978-0-12-799962-3.00001-X

poet, musician, florist, chef, chemist, or chess master are, of course, very different. No one would suggest that knowing what a *haiku* is will be of much use when cooking, designing chemistry experiments, or playing chess. But given reasonable levels of domain-specific knowledge in several domains, is there some broadly applicable way of conceptualizing or approaching problems, some general tendency to think in unusual or offbeat ways, some individual personality trait or quirk, or some comprehensive kind of thinking skill or thinking style that will generally lead to more creative outcomes no matter which of those domains one happens to be working in?

If one were asking similar questions about intelligence, the likely answer would be yes. Although there are some notable and even famous dissenters (such as Gardner, 1983), the consensus among those who study intelligence is that it is a domain-general set of abilities that are associated with performance across a very wide range of domains (Neisser et al., 1996). If someone shows intelligence in one area, it is likely that person will exhibit intelligence in the many other areas in which intelligence is thought to matter. As a result, people with more intelligence are likely to be better at chemistry, cooking, chess, writing poetry, composing music, and flower arranging than people with less intelligence (other things – such as domain-specific knowledge – being equal). Intelligence is fungible, like money: it can be used profitably in many very different kinds of endeavors. That doesn't mean that intelligence (or *g*, as psychometricians often call it) can be used, or that it will be useful, everywhere — just as money can be useful in many different, but not *all*, situations. (As the Beatles and others observers have warned us, money can't buy one love, among other things.)

A domain-general theory of intelligence therefore has limits; it argues that the skills that make up *g* can be widely useful in many diverse and seemingly unrelated contexts, but not *all* contexts. But even with this limitation, a domain-general view of intelligence is very broad and far-reaching, claiming almost (but not quite) universal applicability. In doing so it does not, however, insist that the skills that make up *g* are the *only* kinds of cognitive abilities that matter or deny the importance of many domain-specific cognitive abilities that operate primarily in a single domain. The sole claim of *g* is that there are many very significant domain-general cognitive abilities, not that *g* matters in *every* domain or that it is the only thing that matters in *any* domain (other than the "domain" of taking IQ tests).

Expertise, in contrast, works in an entirely different way. No one assumes that someone who is an expert in modern art will also know a great deal about Heian literature, auto mechanics, or dentistry.

Expertise is domain-specific, and to my knowledge no one has ever seriously made a case for expertise being domain-general other than noting, perhaps, that people with higher intelligence are more likely to have multiple areas of expertise because they can acquire knowledge in most domains more easily. If one factors out intelligence and opportunity to learn, no domain generality is likely to be left when it comes to expertise. (I know of no studies that have attempted to test this claim directly, which is perhaps evidence just how obvious it seems to most psychologists.)

For many years psychologists assumed that creativity was, like intelligence, domain-general. If someone was creative in one area, then that person was more likely than chance to be creative in many other areas; all that would be needed was the acquisition of the necessary skills and knowledge in the new domains. Creative thinking skills could, like the cognitive skills that we call intelligence, be deployed in any field or endeavor. And understanding creativity did not require domain-by-domain investigations because if one understood creativity in one domain, the same general understanding would apply equally in other domains.

Under a domain-general conception of creativity, neither creativity testing nor creativity training needs to be concerned with domains. Consider the task facing those who wished to measure creativity — domain-general creativity, that is, which was the only kind of creativity in which most creativity test designers were interested. Domain-general creativity was, by definition, independent of domains, and so the wisest thing for creativity testers to do would be to make every possible effort to avoid any potentially contaminating effects of differences in domain-relevant skills or knowledge. For this reason, creativity test items have typically been designed to require as little domain knowledge as possible (such as listing possible uses for some common object with which every test taker would be expected to be familiar), because creative-thinking skills were believed to be universal and to exist independent of any specific content on which those skills might be applied.

Similarly, if creativity could be in some way increased through creativity training, it would be increased across the board (under a domain-general understanding of creativity), so the specific content of any creativity-training exercises designed to increase domain-general creativity was inconsequential. Whatever any increase in domain-general creative-thinking skills might produce in one domain it would also produce in most other domains, and brainstorming uses for a brick to increase one's divergent-thinking skill (theorized to be a key creativity-relevant thinking skill) would therefore lead to more creativity when writing poems, solving

puzzles, choreographing dances, designing experiments, or developing theories.[1]

All of these beliefs about the nature of creativity (and about how to test and train it) were grounded in the untested and generally unstated assumption that creativity is a domain-general entity that attaches to domains rather than something that forms part of the essential fabric of each separate domain (and cannot therefore be detached from its respective domain and applied wherever one might wish), as domain specificity theorists claim.

In the past quarter century the idea that creativity is domain general has been seriously challenged. To give a sense of the significance of this issue in the world of creativity research and theory, the *Creativity Research Journal* has published just one invited debate (in the form of a pair of Point-Counterpoint articles) in its history. The two articles that constituted that debate (Baer, 1998b; Plucker, 1998) addressed this crucial domain specificity/generality question, a hugely significant one for creativity research and theory. Even the author of the paper arguing for domain generality acknowledged that the tide had turned in favor of a domain-specific view:

> Recent observers of the theoretical (Csikszentmihalyi, 1988) and empirical (Gardner, 1993; Runco, 1989; Sternberg & Lubart, 1995) creativity literature could reasonably assume that the debate is settled in favor of content specificity. In fact, Baer (1994a, 1994b, 1994c) provided convincing evidence that creativity is not only content specific but is also task specific within content areas. (Plucker, 1998, p. 179)

[1]This parallels expectations for the effects of programs designed to increase intelligence (conceptualized as a domain-general set of skills), based on the fact that those with higher IQs, on average, tend to perform at a higher level of competence across a wide range of tasks and task domains. If *g* – domain-general intelligence – could be increased (as can happen with better nutrition; e.g., Deary, 2008; Glewwe & King, 2001; Katzen-Luchenta, 2007), then performances on a wide range of tasks across a wide range of domains should all improve. It is, in fact, the failure of many supposed intelligence-raising programs to improve skill in more than a few domains that has caused many such efforts to be deemed failures. ("The central question is not whether performance on cognitive tests can be improved by training, but rather, whether those benefits transfer to other untrained tasks or lead to any general improvement in the level of cognitive functioning"; Owen et al., 2010, p. 775). In one comprehensive six-week training study, 11,430 participants practiced and trained several times each week on cognitive tasks designed to improve reasoning, memory, planning, visuospatial skills, or attention. "Although improvements were observed in every one of the cognitive tasks that were trained, no evidence was found for transfer effects to untrained tasks, even when those tasks were cognitively closely related" (Owen et al., 2010, p. 775). Increasing skills in one task seemingly related to intelligence has generally not transferred to improved performance on other intelligence-related tasks, but the domain-generality of intelligence – unlike creativity – has been supported by strong correlations with performance on a wide range of tasks in diverse domains with IQ test scores and similarly strong correlations among performances by individuals across a similarly wide range of tasks (Neisser et al., 1996).

This change represented a nearly 180° turn from just a decade earlier (when domain generality was simply assumed, often implicitly), and as will be shown in Chapter 2, the evidence favoring a more domain-specific view has continued to accumulate.

The domain specificity/generality debate was also at the heart of the first debate ever sponsored by the American Psychological Association's Division 10 (Psychology of Aesthetics, Creativity, and the Arts). The topic of that APA debate was the validity of divergent-thinking tests like the Torrance Tests of Creative Thinking, which are generally assumed to be domain-general assessments (even though Torrance himself found that his two versions of the test, verbal and figural, measured essentially orthogonal variables that were uncorrelated with each other; Cramond, Matthews-Morgan, Bandalos, & Zuo, 2005). Although the APA debate was nominally about the validity of the Torrance Tests, the underlying issue and the central question that animated the debate was the question of domain specificity (Baer, 2009; Kim, 2009; see also Baer, 2011b, 2011c; and Kim, 2011a, 2011b for a follow-up written version of the same debate that was solicited by the APA journal *Psychology of Aesthetics, Creativity, and the Arts*).

Domain specificity/generality is no longer an issue just for creativity specialists. It is the focus of one of the six chapters of *Creativity 101* (Kaufman, 2009), a textbook that is widely used in undergraduate Introduction to Creativity courses, and it will be a featured topic in a forthcoming Oxford University Press *Handbook of Educational Psychology* (O'Donnell, in press), a volume addressed to the field of educational psychology more broadly. Sawyer's *Explaining Creativity* is probably the most comprehensive creativity textbook on the market, and in the preface to the second edition of this text (2012) he noted that the issue of domain specificity had become one of the most controversial topics in the field. After discussing the issue in several chapters and weighing the various research findings, Sawyer concluded that "[t]he consensus among creativity researchers is that although there are domain-general creative strategies, creativity is primarily domain-specific" (p. 395). These examples are evidence of how broadly significant the issue of domain specificity has become for creativity theory, even among nonspecialists.

What difference does it make whether creativity is domain-general or domain-specific? How would this distinction play out in how people outside the field think about and understand what it means to be creative? What might be the educational implications? To answer these questions (which this chapter will preview and which will be developed more fully in later chapters), it is helpful first to address the intuitive understandings most people have about how creativity works. So think for a moment about your own creativity. How creative are you? If you were to give yourself a "creativity score" on a scale of 1–100, where would you place yourself?

The answer for most people is something on the order of, "Well, it depends." There are many things on which it might depend, such as the time of day, one's motivation, how much ethanol or other drugs one might have ingested, and the social and physical environment. But the big "it depends" issue is what one is asked to be creative *in*. Are you equally creative in everything you do, whether writing poetry, solving math equations, woodworking, dancing, solving interpersonal problems, designing science experiments, composing music, developing sports strategies, sculpting, gardening, teaching children how to do something, solving puzzles like the Rubik's cube, or arranging complex schedules? Of course, one needs training to do many of those things, but is that the only thing that causes you to be less creative in some areas than others? Are there areas in which you have had some experience and yet find yourself far less creative than you are in other areas? Is one reason that you are more creative in some areas than others that it just seems easier for you to be creative in those areas?

It is not my goal to convince you that creativity is domain-specific based on your intuitions. How you and others might answer these questions is not the kind of evidence that counts in psychology: intuitions can be wrong, and common sense is often a poor guide. My goal in asking them is simply to get you thinking about what it would mean if creativity *were* truly domain-general (the way it is claimed intelligence is), and what it would mean, on the other hand, if it were domain-specific (the way expertise seems to be). Here's another such question: Think of an area in which you are especially creative. It doesn't matter whether it is an academic field or a field far from academic pursuits, an artistic or a scientific field, a practical field or a theoretical one. Now think of a field of a different kind, one that you have in some way engaged but in which you are not especially creative. (Perhaps you are a creative woodworker but an uncreative poet, or vice versa, or a creative writer but rather uncreative when it comes to drawing or solving math puzzles. Pick a contrast of that sort.) Now think about this: Could you apply whatever it is that makes you creative in the area in which you are especially creative in ways that would produce much more creative work in the other field? Could your creativity when fixing mechanical things be put to good use in writing sonnets or one-act plays, or vice versa?

Throughout the book, I present evidence of a much more scientific nature that should have more weight than any intuitions you might have, based either on your answers to these questions or other hunches you might have about how creativity works. The goal here with these questions is to contemplate the possibility that creativity may be more like expertise than intelligence, that it may be much more domain-specific than domain-general, and to highlight what that would mean.

As we will see, the implications of domain specificity in creativity are both interesting and important. In fact, the impact of domain specificity

for much of the work that has been done in creativity research and testing is potentially devastating. Domain specificity calls into question the assumption that a general theory of creativity is even possible. In contrast to the one-theory-fits-all approach of domain generality, domain specificity calls for one theory to explain creativity in poetry, a different theory to explain creativity in chemistry, yet another theory to explain creativity in film-making, and so on. Similarly, domain specificity argues that one cannot simply apply one's creativity as a poet to help solve problems in chemistry, or vice versa. It also suggests that much of what researchers may think they know about creativity may not be true because the so-called creativity tests used in much of the research could not possibly be valid, or at least not valid outside a particular domain. In addition, generic creativity training – learning how to think outside of just about any kind of box – is seen to be impossible when understood through the lenses of domain specificity (although domain specificity also shows how to make creativity training much *more* effective in a given domain). Domain specificity points research in creativity in an entirely different (and frankly more difficult) direction – many directions, actually.

How is Creativity Measured?

Research in creativity has been hampered by the lack of good measures of creativity. Chapter 5 contains a discussion regarding research about domain specificity means for creativity testing. What I need to explain now is (sort of) the opposite: what creativity testing means for research about domain specificity.

Unfortunately, the ways creativity has most often been tested and the assumptions made by many creativity tests make those tests unsuitable for use in determining whether creativity is domain-specific or domain-general. Even if the most widely used tests were valid, which for the most part they are not, they would still not be useful for judging questions about domain generality and domain specificity because they are simply the wrong kinds of tests. The situation is rather like being forced to use a spelling test to determine whether musical, mathematical, artistic, athletic, and verbal abilities are related. Those five kinds of abilities may or may not be related, and there are research designs that might help probe what, if any, those relationships might be. But even a well-supported, valid test of spelling, used by itself, would be of little use in answering questions about possible connections among these different kinds of abilities.

The situation is difficult, but it's not hopeless. Just as one might use scores on separate tests of musical, mathematical, artistic, athletic, and verbal abilities to probe what inter-relationships there may be among these abilities, there are ways to assess creativity in different domains that can be used to

answer questions about domain generality/specificity. But the kinds of tests needed are not the inexpensive, easy-to-administer, and objectively scorable domain-general tests that have long dominated creativity assessment.

Tests of divergent thinking were for many years the most commonly used measure of creativity. In a 1984 review of all published creativity research, the Torrance Tests, which are not the only tests of divergent thinking in use, but certainly the most widely used, accounted for three-quarters of all creativity research involving students and 40% of the smaller subset of all creativity research involving adults as subjects (Torrance & Presbury, 1984). The Torrance Tests and other divergent-thinking tests are based on Guilford's (1956) Structure of the Intellect model, in which he argued that "divergent production" — thinking of a wide variety of ideas in response to an open-ended question or prompt — was a significant contributor to creativity. In defining divergent production (which means the same thing as divergent thinking, a term Guilford also used; *divergent thinking* is the term more commonly used today), Guilford clearly distinguished between divergent and convergent thinking:

> In convergent-thinking tests, the examinee must arrive at one right answer. The information given generally is sufficiently structured so that there is only one right answer. . . . [A]n example with verbal material would be: "What is the opposite of hard?" In divergent thinking, the thinker must do much searching around, and often a number of answers will do or are wanted. If you ask the examinee to name all the things he can think of that are hard, also edible, also white, he has a whole class of things that might do. It is in the divergent-thinking category that we find the abilities that are most significant in creative thinking and invention. (Guilford, 1968, p. 8)

Torrance, whose eponymous Torrance Tests of Creative Thinking are actually tests of divergent thinking (these tests will be discussed in some detail later), made a similar point:

> Learning by authority appears primarily to involve such abilities as recognition, memory, and logical reasoning – which are, incidentally, the abilities most frequently assessed by traditional intelligence tests and measures of scholastic aptitude. In contrast, learning creatively through creative and problem-solving activities, in addition to recognition, memory, and logical reasoning, requires . . . evaluation . . ., divergent production . . ., and redefinition. (Torrance, 1970, p. 2)

Four aspects of divergent thinking are frequently mentioned in the literature:

- Fluency is the total number of responses to a given stimuli, "the total number of ideas given on any one divergent thinking exercise." (Runco, 1999a, p. 577)
- Originality is the distinctiveness of responses to a given stimuli, "the unusualness . . . of an examinee's or respondent's ideas." (Runco, 1999a, p. 577)

- Flexibility is the number of different categories or kinds of responses to a given stimuli, or more broadly, "a change in the meaning, use, or interpretation of something." (Guilford, 1968, p. 99)
- Elaboration is the extension or broadening of ideas in one's responses to a given stimuli, "the richness of detail in the ideas one produces." (Baer, 1997a, p. 22)

A recent book on creativity assessment illustrated these with the following scenario:

> [I]f a person were planning a social occasion at a restaurant to celebrate a special occasion, she may want to produce a list of possible locations. She may produce a list of 50 potential restaurants (high fluency), a list that includes restaurants her friends would be unlikely to think about (high originality), a list with a wide range of types of restaurants (high flexibility), or a list that includes only Indian restaurants but lists every possible such establishment in the area (high elaboration). (Kaufman, Plucker, & Baer, 2008a, p. 18)

Most early tests of creativity were essentially divergent-thinking tests, which had very little competition for many years except for one another. Their seniority is probably one reason why the tests have been used so widely, but they also had other advantages. They provided a convenient parallel to single-number IQ testing (even though its proponents, including Torrance himself, often argued against such a conceptualization; Kim, Cramond, & Bandalos, 2006); the tests are simple to administer, even to young children; and the idea of divergent thinking on which they are based is easy to understand and has a strong intuitive appeal (Baer, 1993; Kaufman et al., 2008a; Kim, 2008; Runco, 1999a; Torrance, 1993; Torrance & Presbury, 1984; Wallach & Wing, 1969).

Divergent thinking, conceptualized as one component of creative thinking, remains an important concept among creativity researchers and is the basis of some of the most common creativity-training activities (such as brainstorming, even though brainstorming predated Guilford's discovery of divergent production by a few years; Guilford, 1956; Osborn, 1953). Divergent thinking can be thought of as either domain-general or domain-specific, but its conception as a domain-general skill is far more common. To the extent that creativity is domain-specific, however, domain-general theories of divergent thinking cannot be valid and must be replaced by domain-specific versions.

Domain-specific divergent thinking works exactly the same as domain-general divergent thinking in producing a range of possible responses to an open-ended prompt (and fluency, flexibility, originality, and elaboration remain key components of divergent thinking under domain specificity). The difference is simply that the divergent-thinking skills that promote creativity in one domain differ from the divergent-thinking skills that lead to creativity in other domains (e.g., being able to think of many different

and unusual ways to explain division by fractions might lead to creativity in teaching mathematics but have little value in other domains, such as sculpting, composing music, or teaching history).

Because divergent thinking can be conceptualized as a wide variety of domain-specific skills rather than a single, domain-general skill, acceptance of domain specificity does not require creativity researchers to abandon divergent thinking as an important contributor to creativity. Domain specificity argues that the use of domain-general tests of divergent thinking cannot be valid, but domain-specific tests could still be devised and used if one needed them for some special purpose, as one might in creativity research. Domain specificity also changes the ways one must go about teaching people to be more creative, but even when teaching such creative-thinking skills directly, as will be shown in Chapter 6 on creativity training, divergent thinking may be equally important under a domain-specific interpretation; it just needs to be applied in a somewhat different manner, which will influence the kinds of prompts and training activities one might choose. It is probably true that some form of divergent thinking (in either its generic domain-general version or its more recent domain-specific conceptualization) is probably part of creative thinking – that is an empirical question the answer to which has been somewhat clouded by the use of supposedly (but in fact not actually) domain-general divergent-thinking training and testing – although it no longer seems likely that it is the sole or primary ingredient as has sometimes been assumed in the past (Amabile, 1996; Kaufman, 2009; Kaufman & Baer, 2005a, 2006; Simonton, 2010a; Sternberg, 1999).

Unfortunately, it is the domain-general version of divergent thinking that serves as the basis for some of the most widely used creativity tests. Domain specificity calls into question the use of such tests and challenges the validity of research results that have been based on those tests. In selecting measures to use when conducting research about domain generality/specificity, divergent-thinking tests present a special problem because domain generality is a built-in assumption of the tests. There are, for example, two distinct versions of the Torrance Tests, one figural, one verbal, but both are routinely used as domain-general tests.

Each of the two Torrance Tests reports various subscores. There have been many changes in these subscores over the years, but as an example, the figural test currently claims to "assess five mental characteristics" and 13 "creative strengths" (Scholastic Testing Service, 2013). There is also an overall "creativity index," but Torrance himself cautioned against any single-number interpretations of his tests:

> Torrance has discouraged the use of composite scores for the TTCT. He warned that using a single score like a composite score may be misleading because each subscale score has an independent meaning. (Kim et al., 2006, p. 461)

Torrance also found that his two domain-specific tests of divergent thinking were essentially uncorrelated:

> Reponses to the verbal and figural forms of the TTCT are not only expressed in two different modalities . . . but they are also measures of different cognitive abilities. In fact, Torrance (1990) found very little correlation (r = .06) between performance on the verbal and figural tests. (Cramond et al., 2005, pp. 283–284)

Torrance's cautions have fallen on deaf ears, however. Subscale scores that measure different aspects of divergent thinking are routinely ignored in favor of overall creativity index scores, especially by gifted/talented programs, which are the most active users of the Torrance Tests (Scholastic Testing Service, 2013), and researchers now often argue that the overall Creativity Index is the best predictor of creative ability (e.g., Plucker, 1999; Yamada & Tam, 1996).

The fact that Torrance created two different domain-specific tests of divergent thinking and found that they were essentially orthogonal and therefore measuring two very different abilities (Cramond et al., 2005) has naturally caused problems for those who have used both tests in the same study and interpreted them both as measures of domain-general creativity. For example, a recent study conducted with the aim of validating the Torrance Tests found that one of the tests correlated with key outcome measures but the other did not. In that case, verbal divergent-thinking scores predicted many of the kinds of things the study had used as evidence of creative performance (things subjects had self-reported as personal accomplishments from a checklist of creative achievements), but figural divergent thinking scores did not. As the author explained:

> The importance of verbal DT relative to figural DT may be due to a linguistic bias in the adult creative achievement checklists. For example, if a majority of the creative achievements required a high degree of linguistic talent, as opposed to spatial talent or problem solving talents, the verbal DT tests would be expected to have a significantly higher correlation to these types of achievement than other forms of DT. (Plucker, 1999, p. 110)

This outcome is exactly what domain specificity theory would predict. Different measures of creativity rooted in different domains will predict creative performance only in their respective domains. Unfortunately, these kinds of findings (including those of the tests' creator) have not caused those who market the Torrance Tests to scale back their claims. Both the figural and verbal forms of the test purport to be general tests of creativity (Scholastic Testing Service, 2013).

Because divergent-thinking tests like the Torrance Tests assume domain generality, they can hardly be used in studies whose aim it is to test whether creativity is domain-general or domain-specific (even though results

from tests do, despite their claims of domain generality, provide evidence, such as that offered by Torrance himself, that creativity is domain-specific; Cramond et al., 2005). For that kind of research, an assessment technique that is agnostic about domain generality/specificity is needed.

Fortunately, divergent thinking no longer has either the kind of monopoly it once enjoyed in creativity theory and testing nor the wide respect that it once engendered among creativity theorists and researchers. Other creativity tests have been developed, and although none has gained the kind of near-universal acceptance that the Torrance Tests (and other divergent-thinking tests) once had, they do provide other possibilities as research tools.

A recent book about creativity assessment (Kaufman et al., 2008a) contained chapters about four kinds of creativity measures: divergent thinking, assessments by others (teachers, peers, parents), self-assessments, and a procedure called the Consensual Assessment Technique (CAT), which uses experts to judge the creativity of things people have created (poems, artwork, theories, puzzles, soufflés, advertisements, performances of any kind; it can be used for just about anything).

Assessments by others and self-assessments do not require an assumption of either domain generality or domain specificity, but the way the assessments are structured can (and often does) add such an assumption. If one asks about X's creativity as an architect, no assumption of either domain generality or domain specificity is made. (If creativity is in fact domain general, then X's creativity in architecture would, of course, speak to X's creativity in general. But asking about X's creativity in architecture only *allows* domain generality – it does not assume it – and it equally allows for creativity to be domain-specific.) But if one simply asks how creative X is, then the question assumes the answer will apply to X more generally. A general "How creative is X?" question therefore assumes domain generality, and as such cannot help a researcher trying to determine just how domain-general or domain-specific creativity might be. Unfortunately, in most instances in which researchers have asked about the creativity of others, the questions have been framed in ways that require domain-general answers and as such are of little use in settling any disputes about domain generality/specificity.

Self-assessments of creativity are similar to assessments by others in that they can ask about creativity in general (thereby assuming domain generality) or about creativity in specific domains (which makes no assumptions about generality/specificity). Studies of this sort have tended to show a great deal of domain specificity (as discussed in Chapter 2), but this technique has two great weaknesses:

1. Self-assessments in general, and self-assessments of creativity in particular, tend to have limited validity. (Some might go even further to argue that they have no validity whatsoever, but either way – with

either very limited validity or no validity at all – they are significantly less-than-ideal research tools).

2. Although when asked to judge their own creativity in different areas people tend to rate themselves differently in different domains, one could argue that although such questions do not actually assume domain specificity, they might tend to push responses in that direction. After all, if creativity were domain-general, why would one ask about creativity like in many different areas? It might seem to respondents that those asking the questions do, in fact, assume domain specificity.

Self-assessments and assessments by others, then, have not been especially useful in answering questions about domain generality/specificity. Divergent-thinking tests typically assume domain generality; in fact, even though they need not do so, all commonly used divergent-thinking tests – even ones like Torrance's that bear the domain-based labels *figural* and *verbal* – make this assumption and encourage domain-general interpretation. Divergent-thinking tests also face the problem that the evidence for the validity of divergent-thinking tests as measures of creativity is, at best, somewhat weak. As already noted, the first debate ever sponsored by the American Psychological Association's Division 10 (*Psychology of Aesthetics, Creativity and the Arts*) was about the validity of divergent-thinking tests like the Torrance Tests (Baer, 2009; Kim, 2009), suggesting it is an open question. (The title of the debate was "Are the Torrance Tests of Creative Thinking Still Relevant in the 21st Century?") So even if one could find a way to use divergent-thinking tests to measure the domain generality/specificity of creativity (such as by giving subjects tests of divergent-thinking in different domains and comparing the results, something Torrance himself did, with results that clearly pointed to domain specificity; Cramond et al., 2005), questions about the validity of divergent-thinking tests (even domain-specific divergent-thinking tests) would undermine confidence in the results obtained.

That leaves the other primary method of creativity assessment, the CAT (Amabile, 1982, 1983, 1996). The CAT assesses creativity at all levels (whether the garden-variety little-c creativity that even children demonstrate or the paradigm-shifting Big-C creativity of the most original and influential thinkers in their fields) in the same way that creativity is most often assessed in the real world – by the opinions of experts in the relevant domain. Just as Nobel Prize winners are selected by panels of experts in each field who judge the creativity of contributions to their respective fields, the CAT employs experts in a domain to judge the creativity of actual products in that domain. The judgments of experts can, of course, change over time; the standards in any field, whether artistic, scientific, or practical, are not immutable, and what might be viewed as creative in one

era might be thought of less highly in another (and vice versa), as might the qualifications of experts in any field. But the best possible estimate of the creativity of any product at a given point in time is the collective assessment of the acknowledged experts in that field. There is simply no better gauge available (Baer & McKool, 2009, 2014).

The experts doing creativity ratings in a CAT assessment make their judgments independently — there is no opportunity for them to influence one another's opinion — which allows for a check on inter-rater reliability, which is generally quite good. Different experts are of course needed depending on the artifacts to be judged. Poets, poetry critics, and poetry teachers might serve as judges if the artifacts in question were *haiku* poems, whereas artists, art critics, and art teachers would be appropriate if the artifacts were collages. Each expert is asked to assess individually the creativity of every product in the study in relation to all the others in the sample and not in comparison to any external standard. All judgments are relative to the creativity of the other artifacts in the group being judged. A Likert-type scale is used so that the scores are spread across a range of possible ratings, and judges are encouraged to use the full scale, rating the most creative artifacts in the group at the highest score and the least creative at the lowest score, with the goal of differentiating the comparative creativity among the artifacts in the group. Mean ratings of all the judges (who might number 10–15 in a typical study) are used as the creativity scores of each of the artifacts (Amabile, 1996; Baer, Kaufman, & Gentile, 2004; Kaufman et al., 2008a).

The process is both simple and straightforward, although simple unfortunately does not translate to being either easy or inexpensive. Unlike the college students who make up the subject pool for so much of psychological research, experts are not so freely available. Different kinds of experts are needed, depending on the kinds of artifacts to be judged, and the experts are typically paid for their work. Some studies have shown that quasiexperts in some domains (such as students in a field who might not yet qualify as experts) produce ratings rather similar to those of experts, which can reduce the costs somewhat. Use of novice raters (such as college students), however, rarely produces the same kinds of ratings as experts, so the CAT generally requires judges with at least a modest level of expertise in the field in question and is therefore more expensive than many other methods of creativity assessment (Kaufman, Baer, & Cole, 2009b; Kaufman, Baer, Cole, & Sexton, 2008b; Kaufman, Baer, Cropley, Reiter-Palmon, & Sinnett, 2013a).

The CAT is somewhat resource-intensive, but it has much going for it and has been called the "gold standard" of creativity assessment (Carson, 2006). The long-term stability of single CAT assessments is as good as the long-term stability of scores on well-established, multiple-item divergent-thinking tests (e.g., in both cases test-retest correlations after

one year fall in the 0.50 range with elementary-school-age subjects), and if multiple creative products of the same type are created and judged in both pre- and posttesting, CAT long-term stability shows even better results (Baer, 1994c; Kogan, 1983). The CAT can be used to judge the creativity of artifacts in almost any domain, and, unlike divergent-thinking and other tests of subskills theorized to be associated with creativity, the CAT assesses actual creative performance and is therefore not dependent upon acceptance or validity of any particular theory of creativity. The CAT also avoids halo effects and other personal biases that might interfere with assessments of creativity by others or oneself. Although judgments are about the creativity of artifacts in a given domain, no assumption is made when using the CAT about the domain specificity or generality of creativity. It is not linked to *any* theory about the nature of creativity and is completely neutral regarding domain generality/specificity issues, which makes it an ideal measure of creativity in this arena.

Chapter 2 will review research about the domain specificity and generality of creativity in detail, much of it using the CAT. The results have been fairly consistent in showing little domain generality. A key research methodology has been to give subjects a number of different tasks in different domains (e.g., create a collage, write a poem, write a story), have panels of experts in the respective domains independently rate those products for creativity using the CAT, and then look for correlations between the ratings in different domains. The two competing theories – domain generality and domain specificity – make different predictions regarding actual creative performance. Here's how one creativity researcher succinctly summarized how these predictions would differ:

> Domain generality would be supported by high intercorrelations among different creative behaviors . . . while domain specificity would be supported by relatively low correlations among different behaviors. (Ivcevic, 2007, p. 272)

The correlations reported in the many studies that have made exactly this comparison have tended to hover around zero (especially if variance attributable to intelligence is removed; Chapter 2 reviews this research in detail, but see Baer, 2010, 2013 for summaries), and even authors who have claimed to have found some degree of domain generality typically only find it *within* domains. For example, Conti, Coon, and Amabile (1996) reported correlations of creativity ratings on several short story writing tasks ranging from 0.43 to 0.87 and some smaller, but still statistically significant, correlations among several different art tasks (the art tasks were less similar to one another than the writing tasks, so this outcome was expected). These results are all within-domain correlations, however, and they therefore only show that within a domain (such as short story writing or art) there is some generality, as both domain specificity and

domain generality predict. In contrast to the many statistically significant within-domain correlations; however, of the 13 cross-domain (writing-art) correlations reported – the ones that matter for domain specificity – all were tiny and, whether positive or negative, *none* was statistically significant. So despite these authors' claims of evidence for domain generality, all they actually found was *within-domain* generality (i.e., domain specificity).

Feist (2004) commented that it is an "appealing, and ultimately firmly American, notion that a creative person could be creative in any domain he or she chose. All the person would have to do would be to decide where to apply her or his talents and efforts, practice or train a lot, and *voilà*, you have creative achievement. On this view, talent trumps domain, and it really is somewhat arbitrary in which domain the creative achievement is expressed." Although appealing, Feist concluded that "this is a rather naïve and ultimately false position and that creative talent is in fact domain specific . . . creativity and talent are usually not among the domain general skills" (p. 57).[2]

After reviewing the evidence for and against domain specificity in Chapter 2, Chapters 3–6 explore what these research results mean for creativity theory, creativity research, creativity testing, and creativity training, followed by a look at what kinds of creativity theories would be viable under domain specificity. Although readers are invited to read through each of these chapters in order, they have been written with the understanding that many readers will have special interests in just one or a few of the chapters. Skipping chapters should not result in great confusion (although writing them in a way to make this possible has required occasional repetition of some key ideas and research findings). Readers already familiar with (and convinced by) the research evidence supporting domain specificity, for example, may wish to skip the comprehensive review of that evidence in Chapter 2, which covers in greater detail some of the evidence reported more briefly in this chapter.

[2]There is another "firmly American" idea that suggests creativity is easier than it actually is, the belief that "ignorance and lack of experience could actually bestow an advantage, might serve as the wellspring for 'originality and boldness.'" *Lawrence in Arabia* author Scott Anderson described this belief, held by American State Department officers responsible for helping to create the map of the Middle East at the end of World War I, as an "exemplar of the American can-do spirit" (2014, p. 357). Alas, both notions – (a) that one can successfully direct one's creativity at any domain one wishes, and (b) that ignorance and lack of experience are likely to promote creativity – are at odds with what is actually known about creativity. (It hasn't worked out well in the maps that the Western powers created of the Middle East after World War I either.) Expertise matters, although it does matter more in some domains than others (as discussed in detail in Chapter 7).