

John Barker

Mathematical Beauty

1. Introduction

Is there genuine beauty in mathematics? Or when we speak of beautiful results and elegant proofs, are we merely speaking metaphorically? Among mathematicians, it is a received truth that abstract objects, especially proofs, theorems, and even whole areas of mathematical discourse, can possess aesthetic qualities such as beauty and elegance. Is it possible that they are simply wrong?

Nick Zangwill has offered an answer of “yes” in a number of publications.¹ For Zangwill, genuine aesthetic properties are simply too closely tied to sense perception to properly apply in a completely non-sensory domain such as mathematics. For Zangwill, the properties we ascribe to mathematical objects by calling them beautiful or elegant are simply different in kind from the properties we ascribe to physical objects under the same or similar terms.

It is a philosopher’s job to question platitudes; and the platitude that “Euclid alone has looked on beauty bare” is no exception. That said, I am going to argue that Zangwill has got this one wrong. The beauty we see in mathematics is part of the same overall phenomenon that includes beautiful art, music, literature, and natural formations. For the sake of focus, I will concentrate primarily on one instance of purported mathematical beauty: namely elegance in proofs. I will start by considering Zangwill’s case against beauty in mathematics, which is found primarily in his book, *The Metaphysics of Beauty*. I will then sketch some positive reasons for regarding mathematical elegance as an aesthetic property.

¹ The main argument occurs in Zangwill, *The Metaphysics of Beauty* (Cornell: Cornell University Press, 2001), especially Chapter 8, “Aesthetic/Sensory Dependence.” Parts of the argument also occur in his “Beauty,” Chapter 18 of Jerold Levinson, *The Oxford Handbook of Aesthetics* (Oxford, Oxford University Press, 2003), 325–343, especially section 3, “Relevance.”

2. Aesthetic/Sensory Dependence

In *The Metaphysics of Beauty*, Zangwill offers two reasons to doubt that proofs and other mathematical objects can have genuine aesthetic properties. First, he argues that when we call a proof “elegant,” we are not ascribing an aesthetic property to it at all; instead, we are commenting on its effectiveness as a proof. I will consider this argument in the next section. Second, attributing aesthetic properties to mathematical objects conflicts with his thesis of “partial aesthetic/sensory dependence.” In short, he holds that aesthetic properties depend, in part, on sensory properties. Since mathematical objects have no sensory properties, and since the senses do not seem to be involved in any way in the alleged aesthetic properties of mathematical objects, it follows that mathematical objects lack aesthetic properties.²

Zangwill states the aesthetic/sensory dependence thesis as follows:

Aesthetic properties depend in part on sensory properties, such as colors and sounds.³

And this may appear to rule out aesthetic properties for objects that lack sensory properties, such as mathematical objects. However, it is not immediately clear what the dependence thesis rules in or out. If the qualifier “in part” had been omitted, then the dependence thesis would have been a straightforward supervenience thesis:

Strong Sensory Dependence. Any two objects that are identical with respect to their sensory properties are also identical with respect to their aesthetic properties; equivalently, no two objects can differ with respect to an aesthetic property without also differing with respect to some sensory property.

This strong dependence thesis would just about rule out aesthetic properties for objects that lack sensory properties.⁴ But Zangwill denies this

² There are some who argue that abstract objects are indeed perceptible – see, for example, Jesse J. Prinz, “Beyond Appearances: The Content of Sensation and Perception,” in Tamar Gendler and John P. Hawthorne, *Perceptual Experience* (Oxford: Oxford University Press, 2006), 434–460. If this view is correct, then it would only strengthen my case, I would think; but it is certainly a minority opinion. In any case, it is not clear that such a view, even if correct, would block Zangwill’s argument at this point, since for Zangwill, sensory properties evidently form a more restricted class than perceptual properties. In his discussion of architecture, for example, Zangwill clearly regards spatial properties as non-sensory, even though it seems pretty clear that they are perceptual.

³ Nick Zangwill, *The Metaphysics of Beauty*, op. cit. 127.

⁴ The qualifier “just about” is needed here, because strictly speaking, strong sensory dependence does not entail that objects without sensory properties lack aesthetic properties; it simply

strong thesis, preferring instead to say that aesthetic properties depend “in part” on sensory properties. Unfortunately, it is not clear what this “in part” amounts to. The strong dependence thesis was naturally construed as a supervenience thesis; however, there is no such thing as partial supervenience. Properties of one type either supervene on properties of a second type, or they do not. Now in motivating the partial dependence thesis, Zangwill points out that other properties besides sensory properties can play a role in determining an object’s aesthetic properties, as long as sensory properties play a role as well. We might try to capture this thought as follows:

Weak Sensory Dependence. There is some class *P* of properties such that aesthetic properties supervene on the combination of sensory properties and *P*-properties.

Or in other words: two objects cannot differ with respect to their aesthetic properties without differing with respect to their sensory properties, or their *P*-properties, or both. Unfortunately, the weak sensory dependence thesis is now *too weak*. One way for objects to differ with respect to the combination of sensory properties and *P*-properties is to differ with respect to *P*-properties alone, not differing at all with respect to sensory properties. And one way for this to come about is for the objects in question to be mathematical objects that lack sensory properties, and for *P* to include the sorts of properties that mathematical objects have. Thus, aesthetic properties for mathematical objects are not ruled out. We may be tempted at this point to strengthen the thesis and require objects with different aesthetic properties to differ with respect to their sensory properties *and* their *P*-properties. However, the resulting dependence condition would now be too strong again: in fact, it would entail the strong sensory dependence thesis, since objects that differ with respect to both sensory and *P*-properties differ, *a fortiori*, with respect to sensory properties.

Zangwill does offer some hints about what he means here. He writes:

The [aesthetic/sensory dependence] thesis is that sensory properties are *necessary* for aesthetic properties, not that they are *sufficient*. Accepting a weak dependence thesis is compatible with admitting that other factors are also necessary.⁵

A supervenience thesis is a sufficiency thesis: if a domain *A* supervenes on a domain *B*, then for any *A*-fact, there is some *B*-fact (or some conjunction

implies that all such objects have the *same* aesthetic properties. Realistically, however, anyone who ascribes aesthetic properties to mathematical objects will also hold that different mathematical objects can have different aesthetic properties.

⁵ Ibid.

of *B*-facts) that entails, i.e., is sufficient for, that *A*-fact. Zangwill's remark therefore suggests that aesthetic/sensory dependence should be construed as follows:

Sensory Necessity. For every aesthetic property, there is some sensory property that is necessary for, i.e., is entailed by, that aesthetic property.

For example, let the aesthetic property in question be the property of being beautiful. According to sensory necessity, this aesthetic property entails some sensory property: i.e., there is some sensory property *P* such that necessarily, all beautiful objects have property *P*. However, this is quite obviously too strong: there is no one sensory property that all beautiful objects have in common.

Thus, it is not clear that the dependence thesis should be construed in terms of entailment relations between aesthetic and sensory properties. Given the work that Zangwill wants the dependence thesis to do, the following seems closer to the mark:

Sensory/Aesthetic Explanation. For any object *X* with an aesthetic property *A*, there are sensory properties that play a role in explaining why *X* has property *A*, although non-sensory properties may also play such a role.

And this probably rules out the sorts of aesthetic properties that are sometimes ascribed to mathematical objects, though it is hard to say for certain without knowing more about what constitutes an explanatory role.

In any case, let us assume that we have before us some version of the aesthetic/sensory dependence thesis that does effectively rule out aesthetic properties for mathematical objects. Why should we believe such a thesis? What argument could be given in its support? Zangwill argues for the thesis by defending it against a series of purported counterexamples. In so doing, he apparently regards it as the default view, the view we should accept unless a good reason can be found for rejecting it. However, he never actually gives an argument for regarding it as the default view, and it is not at all clear why we should so regard it. Could we not, with equal justice, take it as our default position that proofs can be elegant, that theorems can be beautiful, etc., and then challenge Zangwill to refute that position?

If we suppose, for the sake of argument, that the dependence thesis holds for the various traditional art forms, then this in itself might be taken as evidence that the thesis holds more broadly. I am not sure that Zangwill himself makes this argument, however; and this is just as well, because the argument is a questionable one. Just because a generalization holds for one domain, it by no means follows that it holds for other domains. Consider, for example, the following thesis, which Zangwill does *not* hold:

Artifacts. Only an artifact can have aesthetic properties.

Arguably, this thesis has no counterexamples in the art world, because works of art are also artifacts. Yet as most people acknowledge (including Zangwill), other objects besides artifacts can have aesthetic properties. In particular, natural objects can be beautiful. Thus, the fact that the artifact thesis holds for works of art is at best weak evidence that it holds in general, and likewise for the aesthetic/sensory dependence thesis. In the case of the artifact thesis, the inference fails because the evidence base is too narrow: it is too narrow precisely because it only includes artifacts. And in the case of aesthetic/sensory dependence, the evidence base may also be too narrow precisely because it excludes proofs, theorems, and other mathematical objects.

Thus, I am arguing, Zangwill's *general* case for aesthetic/sensory dependence does not provide a compelling reason to accept that thesis in the *specific* instance of mathematical objects. To make that case compellingly, one would have to make a specific argument that mathematical objects lack aesthetic properties. Fortunately, Zangwill provides such an argument, to which we will now turn.

3. Elegance or Effectiveness?

In arguing that proofs lack genuine aesthetic properties such as elegance, Zangwill makes two closely related points. First, our basis for attributing such properties to proofs is too closely tied to the *function* of proofs to count as aesthetic judgments, or for the properties thus attributed to count as aesthetic properties. And second, when we do attribute elegance to proofs, we are commenting not on the proof's aesthetic properties, but on its *effectiveness* as a proof. I think both points make a mistake about what we are commenting on when we describe a proof as elegant.

What is the function, purpose or end of a proof? The answer is simple: a proof purports to establish a given result, and it succeeds if and only if it really does establish that result with mathematical rigor. For Zangwill, the so-called elegance of a proof is too closely tied to its purpose to count as beauty. He writes: "Our admiration of a good proof, theory, or chess move turns solely on its effectiveness in attaining these ends, or else on its having properties which make attaining these ends likely."⁶ It may seem that a proof could nonetheless have *dependent* aesthetic properties, which are the aesthetic properties an object has as an object with a given purpose. Zangwill disagrees: "[W]hat we are appreciating in these cases is not dependent beauty or elegance but the mere technical achievement of finding a very effective means to an end."⁷

⁶ Ibid., 141.

⁷ Ibid., 142.

However, I think it is simply wrong to say that attributing elegance to a proof is a comment on the proof's effectiveness. First of all, effectiveness is an all-or-nothing affair when it comes to proofs. Either a proof establishes its result, or it does not. One proof simply cannot establish a result more effectively than another proof, assuming both proofs are successful: a successful proof is *entirely* effective. There therefore seems to be no room for comparative judgments about the effectiveness of successful proofs. However, if Zangwill is right then that is exactly what judgments of elegance would have to be. It is commonplace among mathematicians that one (successful) proof of a given theorem can be considerably more elegant than another; yet both, being successful proofs, are equally effective.

We may be tempted to say that an elegant⁸ proof does its job *better* than an inelegant proof. But an elegant proof of a theorem does not make that theorem more true, or more likely to be true, than an inelegant proof would. Thus, in saying that an elegant proof does its job better, we are not saying it does its job more effectively. Instead, the fact that we have preferences among equally effective proofs shows that in addition to judging proofs for their effectiveness, we also judge them by some standard other than effectiveness.

Could it be, as Zangwill suggests, that elegant proofs have more features that are *conducive* to success than do inelegant proofs, or perhaps features that are more *strongly* conducive to success? There are several problems with this move. First of all, it is not entirely clear what it means to say that one proof has features that are more conducive to its success than another proof of the same theorem. After all, a proof is, by definition, *fully* successful. Now perhaps the success-conducive features Zangwill has in mind are methods or patterns of reasoning that are fruitful, in that they lead to, or can be found in, many other successful proofs. Mathematicians take methods and strategies of proof very seriously, and for good reason: they can be re-used in other proofs. However, it is not at all clear that the fruitfulness of a strategy or method makes any contribution to the *elegance* of the proofs in which they occur, and it is especially unclear that they make the *only* such contribution.

Consider elementary proofs. An elementary proof is one that can be grasped without much advanced or specialized mathematical knowledge. An elementary proof might, for example, use nothing more than high school algebra. Thus, elementary proofs are relatively lacking in fruitful methods and strategies: the methods and strategies they employ will tend to be fairly basic and not particularly interesting. Yet such proofs can be very elegant, and indeed, they can be elegant *because* they are elementary, not despite that fact.

⁸ Here and throughout, I use the term "elegance" to denote whatever it is that we ascribe to proofs by calling them elegant. In so doing, I am remaining neutral on whether the property so ascribed is an aesthetic property.

Being elementary is one way of being simple, and simplicity is widely acknowledged as conducive to elegance. All else being equal, a simpler proof of a given theorem is likely to be more elegant than a complicated one. Yet simplicity does not seem to be in any way success-conducive. Here a distinction must be made. There is a debate about whether the simplicity of a *theory* makes that theory more likely to be true. On the one hand, when we decide what theory to believe, we will tend to pick a theory that is simpler than its competitors, provided the theories are otherwise equally well supported. On the other hand, it is hard to justify this preference, since there is no a priori reason why the truth should be simple. That is an interesting question, but it is a separate question. Scientists may judge a theory to be likely true partly on the basis of the theory's simplicity; but mathematicians *never* judge a proof to be successful on the basis of its simplicity. Nor is there any reason I know of to believe that simplicity makes a proof more likely to work.

Zangwill might argue that even if the features of proofs that form the basis of our judgments of elegance are not conducive to the success of individual proofs, they are still best explained in terms of the overall truth-seeking goals of the discourse in which those proofs occur. Some proofs, for example, are more explanatory than others, in that they leave the reader with a better understanding of why the theorem in question is true. A proof can be perfectly adequate even if it is not explanatory; nonetheless, it seems fair to say that mathematicians' preference for explanatory proofs is a direct result of their concern for mathematical truth in general. Thus, a case can be made that explanatoriness is in some broad sense success-conducive, and at any rate that it is not an aesthetic property.

Unfortunately, an explanatory proof can fail to be elegant, and vice versa. In fact, there is often a tradeoff between these two properties. The most elegant proofs tend to be short, taut and clever: they get their work accomplished as efficiently and directly as possible. However, while these features may contribute to elegance, they often come at the expense of explanation: such a proof can leave the reader rather mystified about *why* the theorem in question is true, even though the reader does not doubt *that* it is true. To redress this deficiency, a longer and more discursive proof may be required, achieving explanatoriness at the expense of elegance. The Second Recursion Theorem from mathematical logic provides a well known example.⁹ This theorem has a short, elementary proof that invariably leaves people completely mystified: everyone who can read a proof in the first place acknowledges that the proof successfully establishes the result,

⁹ The Second Recursion Theorem, usually known simply as the "Recursion Theorem," makes essentially the following assertion. Suppose we have an effective enumeration of all Turing machines, say M_1, M_2, \dots . Now let F be any recursive function, i.e., a function that can be computed by a Turing machine. The Recursion Theorem states that there is some integer i such that M_i and $M_{F(i)}$ compute exactly the same function, i.e., have the same output for any given input.

but most people are left with a very strong feeling that they still do not understand why the result holds. Indeed, there is a small but serious body of literature that seeks to explain why this theorem is true. Most people would consider the proof to be elegant, but hardly anyone considers it to be explanatory.

Thus, it seems that elegance in proofs cannot simply be identified either with the success of the proof itself, or with any feature that is conducive to the success of the overall enterprise in which the proof is situated, since inelegant proofs can possess the very same success-conducive features. That said, elegance in proofs is not wholly separate from the success of the proof, as Zangwill rightly notes. It seems strange to call a proof elegant if it is unsuccessful, and all the more so if it is completely unsuccessful. Zangwill argues that for this reason, elegance in proofs is too closely tied to the success of the proof to count as an aesthetic property.

Now the elegance of a proof is no doubt tied to the manner in which the proof achieves its purpose, even if it cannot simply be identified with the proof's success in achieving its purpose. But in this respect, elegance resembles dependent beauty. A beautiful object is dependently beautiful if it is beautiful in a way that is tied to its function: specifically, an object is dependently beautiful, for Zangwill, because of the manner in which it expresses and articulates its function. Thus, we might be tempted to regard elegance in proofs as an instance of dependent beauty. Zangwill anticipates this move, and counters it by pointing out that a dependently beautiful object can be dependently beautiful even if it wholly fails to fulfill its purpose. For example, imagine a building whose purpose essentially involves sturdiness, and which appears to be quite sturdy, but is not: its apparent sturdiness is due entirely to a façade. We may imagine that the building wholly fails to achieve its purpose for this reason. Yet the façade, while not contributing to the fulfillment of the building's purpose, nonetheless contributes to its dependent beauty, because it expresses and articulates sturdiness. By contrast, we cannot imagine a proof that wholly fails in its purpose but which is nonetheless elegant.

Thus, Zangwill is making the following argument:

1. A proof cannot be elegant if it is unsuccessful, or at least if it lacks any features that are conducive to success.
2. Therefore, what we call elegance in proofs is not a dependent aesthetic property.

To justify the inference from (1) to (2), he needs something like the following principle:

Non-Instrumentality. When a dependent aesthetic property P is connected to a function F , objects must be capable of having property P while completely failing to fulfill the function F .

This is a rather strong claim, and while it is not without some plausibility, Zangwill provides no real argument for it: he simply provides a few examples that conform to the principle, such as the example of false sturdiness cited above. Thus, I suppose one could simply deny the principle, citing elegant proofs as a counterexample. After all, the set of special cases offered in support of the principle may be too narrow *precisely because* it fails to include elegant proofs.

However, I think we can make a more satisfying reply to Zangwill here. Consider the case of a building with a misleading façade. The building manages to express and articulate sturdiness because it *looks* sturdy. When it comes to sturdiness, there is a difference between appearance and reality. The two can come apart; and if the building's aesthetic properties are tied specifically to its appearance, and not (or not just) to its function, then its aesthetic properties can come apart from its function, or from the fulfillment of its function, as well. In proofs, however, there is little or no distinction between appearance and reality. The correctness of a proof is a *manifest* property of the proof. An incorrect proof simply cannot appear correct in the way that a flimsy building can appear sturdy. And a failed proof certainly cannot appear to be correct while at the same time being *completely* devoid of success-conducive properties. At most, a failed proof might fail because of a few subtle flaws in an otherwise sound argument, and might therefore appear correct to many readers. And in that case, the proof may well have a good deal of elegance in it, notwithstanding that it is not entirely correct.

Thus, what is special about proofs here is that being successful is a manifest property of proofs, while being sturdy is not a manifest property of buildings. We might even maintain that an object's aesthetic properties depend on its manifest properties, or at least partly depend on them, thus generalizing Zangwill's aesthetic/sensory dependence thesis. In doing so, we could easily explain why an object can be dependently beautiful while failing to fulfill its purpose: it is dependently beautiful because its beauty derives primarily from its manifest properties, which in turn can come apart from how well or badly it fulfills its function. And at the same time we can accommodate elegant proofs, whose function cannot come apart from their manifest properties in the same way. I am not arguing that we should adopt this manifest property thesis; I am simply arguing that since it explains all the cases Zangwill presents, it is the most he is entitled to assert.

I suspect that Zangwill is right to insist on some separation between an object's aesthetic properties and its ability to fulfill a function. In short, the ability to fulfill a function is typically not an aesthetic property, or so I would think. Being a good hammer is not an aesthetic property of hammers, for example. One might also argue, though this is more controversial, that sturdiness is not an aesthetic property of buildings, even if the *appearance* of sturdiness is. (If you disagree with either example, no worries: you probably already disagree with Non-Instrumentality, which is needed for Zangwill's argument.) Likewise, being a correct proof is not an

aesthetic property, and a proof cannot have an aesthetic property simply because it is correct. But as I have argued at length, proofs are not elegant simply because they are correct. It may be impossible to remove a proof's correctness while leaving its elegance intact, but the reverse is not true: it *is* possible to remove a proof's elegance while leaving its correctness intact.

4. Elegance as an Aesthetic Property

In this final section, I want to present some positive reasons for regarding beauty and elegance, as applied to mathematical proofs, theorems and objects, as aesthetic properties, and not just as misleadingly named non-aesthetic properties. In so doing, I am not making an argument about the boundaries of a set of ordinary concepts. I am not terribly interested in whether the term "elegant," for example, as used by ordinary language users, includes some proofs in its extension. I am more interested in how we *should* use terms like "beautiful" or "elegant" if we are to carve up the world at its joints. Unfortunately, this is a difficult question to answer, and nothing I have to say will be conclusive.

In what follows, I will focus on aesthetic judgments, and the felt responses that underlie those judgments, and deal only indirectly with aesthetic properties. I will argue that judgments of mathematical elegance ought to be counted among aesthetic judgments, and that these judgments are formed on the basis of felt responses that ought to be counted among aesthetic responses. Admittedly, as we move from aesthetic judgments to the aesthetic properties referenced in those judgments, we get into a number of difficult issues about the objectivity (or lack thereof) of aesthetic statements and properties; but these issues are everybody's problem, not just mine. All I will assume in this regard is that when we make a judgment to the effect that a given proof is elegant, we are often thereby saying something true. If we further suppose that the judgments in question are aesthetic judgments, then it is hard to avoid the conclusion that mathematical elegance is a genuine aesthetic property, and that some proofs have it.

In the remainder of this section, I will consider several respects in which judgments of mathematical beauty and elegance are similar to ordinary aesthetic judgments. Hopefully, this will lend some credence to the idea that both types of judgment should be classified together as aesthetic.

4.1. Subjective Universality

Zangwill adopts the following Kantian framework for aesthetic judgments. On the one hand, aesthetic judgments purport to describe objective features of objects, and not simply report a subject's mental state. On the other hand, these judgments are made on the basis of felt responses to the objects they purport to describe; we judge objects to be beautiful or ugly, for example, on the basis of the pleasure or displeasure we take in them. How well do judgments of mathematical elegance fit this framework?

Judgments of elegance clearly satisfy the “universal” side of subjective universality. In judging a proof to be elegant, we purport to describe the proof itself, not just our own feelings in contemplating the proof. We expect sufficiently discerning people to come to similar judgments; and if there is disagreement about how elegant a given proof is, we are very willing to count some people’s judgments as better than others’. What may be less clear is that such judgments are subjective, i.e., based on felt responses.

First of all, judgments of elegance are surely *associated* with felt responses. Not everyone knows enough about mathematics to appreciate the difference between elegant and inelegant proofs, just as not everyone knows enough about art to see the beauty in a given artwork. But those who do have the relevant sensitivity take great pleasure in elegant proofs, beautiful theorems, etc. Moreover, this pleasure seems, subjectively, to be similar in character to the pleasure one takes in appreciating beautiful objects, at least to me – and presumably to many other people, as well, considering the widespread use of terms like “beautiful” and “elegant” in mathematical contexts. Thus, our next question is: do these felt responses form the basis of judgments of elegance, or do they merely accompany such judgments?

There are at least two reasons to believe that judgments of mathematical elegance are based on subjective, felt responses. First, the connection between judgments of mathematical elegance and the corresponding felt responses seems to be a fairly tight one. Just as it is hard to imagine judging a painting beautiful without having any inclination to take pleasure in the painting, it is hard to imagine judging a proof elegant without having any inclination to find the proof pleasing. In other words, judgments of elegance seem closely tied to, and not fully separable from, the *appreciation* of elegance in much the way that judgments of beauty are tied to the appreciation of beauty.

Second, it is hard to see what judgments of mathematical elegance could be based on if they are not based on felt responses. The reason for this is that such judgments are, or tend to be, non-inferential. When we judge a proof to be correct, we are explicitly inferring one fact about the proof (its correctness) from another (the proof’s contents), and we have a fairly clear idea of what standards we are employing when we make such inferences. By contrast, when it comes to elegance in proofs, we simply know it when we see it. And much the same thing can be said about aesthetic judgments about artworks. Granted, we may have various insights into what makes a given proof elegant, or a given painting (say) beautiful. We may feel that something is beautiful or elegant in part because it has features *X*, *Y* and *Z*. But we do not simply *infer* that it is beautiful or elegant from the fact that it has features *X*, *Y* and *Z*. In any case, the non-inferentiality of elegance judgments seems to me to be an important feature of aesthetic judgment and response generally, and I will now say a few words about it.

4.2. Non-Inferential Responses

Aesthetic judgments and responses seem to me to involve non-inferential and non-conceptual knowledge in an important way. Consider a fairly typical case of aesthetic response. A subject looks at a painting. She notices several features of the painting: both low-level features like the placement of individual colors on the canvas, and high-level features such as the painting's overall balance and composition. Her perception of these features prompts an aesthetic response: she likes them, and more specifically, she likes them in a way that supports, for her, a judgment that the painting is beautiful. Non-inferential knowledge enters this picture in a number of important ways.

First, the features of the painting that she is responding to – in other words, the features she appreciates aesthetically – are directly present in the representational content of her visual experience. They are not features that she consciously infers from her visual experience. This is obvious in the case of low-level features such as color placement, but it is true of high-level features also. When we recognize a balanced composition, we do not count up individual low-level features of the scene and make a conscious inference that the scene is balanced. We simply see the scene as balanced. The balance that the subject appreciates in the scene is part of the non-conceptual content of her visual experience, not part of the conceptual knowledge that she arrives at on the basis of this experience.

This feature of our subject's aesthetic response – its sensitivity to non-conceptual knowledge arrived at non-inferentially – seems to me to be the norm in aesthetic response generally. Of course, purely intellectual knowledge can strongly influence our aesthetic responses as well; but this is at least partly due to the fact that intellectual knowledge can affect the contents of our perceptual experience. It is fairly well established that observation is theory-laden. Someone who knows the difference between an elm and an oak will see an elm as an elm; others will simply see it as a tree. In this case, one's knowledge of trees actually has an effect on the content of one's visual experience, and not merely on one's judgment. Likewise, intellectual knowledge can help one more readily detect aesthetically relevant features of a scene. Nonetheless, the perceived feature of the scene is genuinely *perceived*. To aesthetically appreciate a balanced scene, one must see it as balanced, and not just become convinced intellectually that it is balanced.

In short, for a property of an object to have an effect on one's aesthetic appreciation of that object, the property in question must normally be perceived, and not merely known about. Moreover, the subject's response to this property – her aesthetic appreciation of the property, and the judgment that this appreciation supports – also tends to be largely non-inferential. A subject appreciates the beauty in a scene by seeing the scene as beautiful, not by inferring that it is beautiful, and this appreciation directly supports her judgment that the scene is beautiful. Again, explicit knowledge and rational inference can play a role, but as in the case of perception, it tends to play a role by influencing how things are perceived and how we feel about them,

not by lending direct inferential support to aesthetic judgments themselves. It is rare, for example, to decide that a given painting is beautiful *simply* because one has a prior belief that all of that artist's paintings are beautiful, though such a background belief could certainly influence such a judgment indirectly, by influencing how the painting is seen.

Now all of this may seem to argue against my claim that judgments of mathematical elegance are aesthetic judgments. After all, we do not literally see proofs, nor do we perceive mathematical objects and results through the senses.¹⁰ However, non-inferential and even non-conceptual knowledge play a strikingly similar role in the understanding and appreciation of mathematical proofs. First, simply *understanding* a proof requires more non-conceptual knowledge than many people realize. As any aspiring mathematician soon realizes, there is a difference between understanding a proof as a whole, and understanding each individual inference in the proof. Grasping a proof, understanding its gist, seeing why it works, is an important further step, and an essential step if one is to become a competent mathematician. However, by simply following each move in a proof, one has learned everything that is explicitly stated in the proof. Therefore, in really understanding a proof, one must be learning something that is not explicitly stated in it.

Moreover, I think it is pretty clear that this extra something constitutes non-conceptual knowledge. That is, it constitutes something that cannot be stated in language, or at least, that can be grasped independently of one's ability to state it in language. If this were not the case, then the extra knowledge could simply be written down as a further line of the proof, or as a remark following the proof, saving the reader much trouble and effort.

Likewise, in judging a proof to be elegant, we rely on insight, not inference. We simply see the proof as elegant. I actually suspect that the parallel to perception here is strong, though how strong is a psychological question outside the scope of this paper. In either case, we detect higher-order features, either of a scene or of a proof, in a non-inferential and largely unconscious process of analysis and integration.

4.3. Features of Elegant Proofs

No one can explain exactly what features a proof must have in order to be elegant, just as no one can explain what features an object must have in order to be beautiful. Nonetheless, there are certain properties of proofs that tend to contribute to their elegance, and I would argue that the same properties can also make a contribution to aesthetic properties in other domains.

The most obvious feature of elegant proofs is *simplicity*. All else being equal, simple proofs are usually considered more elegant than complicated ones. Simplicity itself is hard to define, but we tend to know it when we

¹⁰ Here I am ignoring the arguments in Prinz, *op. cit.* If mathematical elegance can be genuinely perceived, then the analogy between mathematical elegance and sensory beauty is simply that much more direct, and my argument is, if anything, strengthened.

see it. A proof that proceeds by enumerating seventeen special cases is probably less simple, and almost certainly less elegant, than a proof without cases. Relatedly, elegant proofs tend to be *economical*. That is, they represent a large payoff for a small investment: e.g., a simple but well-placed move creates a large effect in terms of advancing the proof. A third relevant feature is *directness*. A direct proof avoids detours, in the form of unproductive moves and extraneous constructions. Here it should be noted that extraneousness is not the same thing as logical irrelevance. A proof is a chain of inferences, and usually no one inference can be omitted without invalidating the proof. But sometimes this chain of inferences will carry the proof into territory that seems off-topic, and when it does, this usually detracts from the perceived elegance of the proof.

Now this list is obviously rough and incomplete; yet the three features just mentioned seem capable of contributing to (other) aesthetic properties as well. A simple, clean and uncluttered scene can be aesthetically superior to a complicated, busy and cluttered scene at least in part because of its simplicity. A well-placed brush stroke, turn of phrase, or chord sequence – a simple element that creates a significant effect – is an example of economy. And very often, a cluttered scene is so judged because it contains elements that are largely extraneous to the intended overall effect; lack of clutter is therefore an instance of, or at least closely related to, directness. In all of these cases, I am not arguing that simplicity, economy, etc., are necessary conditions for beauty. Far from it: a work of Baroque art, for example, may be beautiful at least in part because of its complexity. I am merely suggesting that *in some instances*, something can be beautiful at least in part because of its simplicity.

Of course, one may argue here that the relevant properties of paintings, musical compositions, etc., are not literally the same properties that one finds in an elegant proof, but merely analogous properties. I disagree, but to explain why it is necessary to distinguish higher-order structural properties from the lower-order properties that they depend on. All of the visual properties of a painting, for example, supervene on the arrangement of colors on a canvas. That is, no two paintings can differ visually without differing in terms of color arrangement. However, the visual properties of a painting – that is, the properties that we can perceive visually – surely include some structural properties that are not identical to color-arrangement properties. Take symmetry, for example. The symmetry displayed in a painting depends, like everything else about the painting, on its specific arrangement of colors. Yet symmetry is a structural property, not a visual one; many different kinds of objects can be symmetrical, even abstract objects. Structural properties are not tied to any one sensory modality or even to sense perception in general, even though they can often be perceived through the senses. I would suggest that the simplicity, economy and directness that we find in proofs are actually highly general structural properties, and that these same properties can also be instantiated in other sorts of objects, including physical objects, and make aesthetic contributions there as well.

4.4. Conclusion

We have now found several points of similarity between aesthetic judgments and judgments of elegance, as well as between the felt responses that underlie these judgments. And I would add one more: there is a strong felt similarity between the two cases. As evidence for this, we need look no further than the near universal tendency to use terms like “elegant” and “beautiful” to describe mathematical proofs and results. Of course, Zangwill regards all such talk as metaphorical. But that simply proves my point, because apt metaphors are based on felt similarities.

Of course, none of these points of similarity, taken either individually or collectively, actually *proves* that mathematical elegance is an aesthetic property. But the more similarities we find between the two cases, the more it seems arbitrary to classify them separately. The best way I know to defeat this line of argument is to find important dissimilarities between the two cases. Can we?

One possible dissimilarity concerns the connection between elegance and correctness. However, we have already seen that this connection is weaker than Zangwill supposes. When we respond aesthetically to a physical object, we are responding mostly to its appearance; so features of the object that are external to its appearance (e.g., sturdiness) should have little or no impact on our aesthetic responses. Proofs do not, strictly speaking have appearances: being abstract objects, they do not affect our sense organs. Thus, when we appreciate the elegance in a proof, we must be responding to something else, and I would argue that we are responding to higher-order structural properties realized in the chain of inferences that constitutes the proof. An argument that does not even come close to being a proof is simply incapable of exhibiting the relevant structural features, and I would suggest that this explains why correctness is relevant to elegance. Moreover, as we have seen already, the appreciation of a proof’s elegance goes significantly beyond the appreciation of its correctness, as evidenced by correct but inelegant proofs.

Beyond this, the only important dissimilarity I can see between the two forms of aesthetic response is that one is based in sense perception and the other is not. But if this is the only basis for excluding mathematical elegance from the aesthetic, then it is surely an arbitrary basis. One could, of course, simply *stipulate* that aesthetic properties and responses are in some sense sensory. Nothing prevents us from using the terms “elegance,” “beauty,” etc., in that way if we so choose. But it is hard for me to see any real benefit in making such a stipulation. Quite the opposite: it will simply blind us to the real and important similarities that exist between mathematical and sensory beauty.¹¹

¹¹ I would like to thank Ewa Bogusz-Boltuc for extremely valuable comments on an earlier draft of this paper, and Roxanne Kurtz, Hei-Chi Chan, and Julia Zhang for much helpful discussion.