

ABSTRACT

Title of dissertation: VISUAL INSIGHT IN GEOMETRY
Logan Fletcher, Doctor of Philosophy, 2016

Dissertation directed by: Professor Peter Carruthers
Department of Philosophy

According to a traditional rationalist proposal, it is possible to attain knowledge of certain necessary truths by means of *insight*—an epistemic mental act that combines the 'presentational' character of perception with the *a priori* status usually reserved for discursive reasoning. In this dissertation, I defend the insight proposal in relation to a specific subject matter: elementary Euclidean plane geometry, as set out in Book I of Euclid's *Elements*. In particular, I argue that visualizations and visual experiences of diagrams allow human subjects to grasp truths of geometry by means of visual insight.

In the first two chapters, I provide an initial defense of the geometrical insight proposal, drawing on a novel interpretation of Plato's *Meno* to motivate the view and to reply to some objections. In the remaining three chapters, I provide an account of the psychological underpinnings of geometrical insight, a task that requires considering the psychology of visual imagery alongside the details of Euclid's

geometrical system. One important challenge is to explain how basic features of human visual representations can serve to ground our intuitive grasp of Euclid's postulates and other initial assumptions. A second challenge is to explain how we are able to grasp general theorems by considering diagrams that depict only special cases. I argue that both of these challenges can be met by an account that regards geometrical insight as based in visual experiences involving the combined deployment of two varieties of 'dynamic' visual imagery: one that allows the subject to visually rehearse spatial transformations of a figure's parts, and another that allows the subject to entertain alternative ways of structurally integrating the figure as a whole. It is the interplay between these two forms of dynamic imagery that enables a visual experience of a diagram, suitably animated in visual imagination, to justify belief in the propositions of Euclid's geometry. The upshot is a novel *dynamic imagery account* that explains how intuitive knowledge of elementary Euclidean plane geometry can be understood as grounded in visual insight.

Visual Insight in Geometry

by

Logan Fletcher

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2016

Advisory Committee:

Professor Peter Carruthers, Chair
Professor Robert Friedel
Professor Jerrold Levinson
Professor Aidan Lyon
Professor Eric Pacuit

©Copyright by

Logan Fletcher

2016

TABLE OF CONTENTS

List of Figures	iv
Chapter 1: Geometrical Insight	1
1 Introduction	1
2 Reasons for skepticism	2
2.1 First objection: lack of motivation	4
2.2 Second objection: unimpressive phenomenology	8
2.3 Third objection: mere discursive reasoning	12
2.4 Fourth objection: mysterious access to abstract objects	19
2.5 Fifth objection: sui generis psychology	20
3 A phenomenological perspective	21
4 Plato: geometrical insight as recollection	30
4.1 Context	32
4.2 The geometrical demonstration	34
4.3 Plato's notion of a priori justification	42
4.4 Recollection as recognition	48
4.5 Integrative understanding as a priori	51
5 Conclusion and prospect	55
Chapter 2: Reliability	57
1 Introduction	57
2 An idea from Plato	61
3 Attribute substitution	63
4 Approximation to perfection	65
5 Failures of imagination	68
6 Conclusion	71
Chapter 3: Towards a Theory	72
1 Introduction	72
2 Giaquinto's project	75
3 Spatial structure and symmetry	78
4 A perceptual concept for squares	90
5 Basic geometrical knowledge	95
6 Criticism of Giaquinto's view on generality	104
Chapter 4: Basic Knowledge of Geometry	115
1 Introduction	115
2 An assessment of Giaquinto's account	115
3 Applying dynamic imagery to Euclid's starting points	121
4 The common notions	128
5 Straight lines	143

6	Linearity in general	151
7	Straightness	163
8	From straight lines to construction postulates	170
9	Indefinite extensions of straight lines	180
10	Circles	189
11	The structure of Euclidean space	193
12	Conclusion	206
Chapter 5: The Generality Problem		207
1	Introduction	207
2	The problem	207
3	Historical solutions	213
4	Contemporary solutions	220
5	Review of the dynamic imagery account	225
6	A novel solution	230
References		241

LIST OF FIGURES

Figure 1.1: Step one of Plato's proof	35
Figure 1.2: Step two of Plato's proof	40
Figure 2.1: Fischbein's example	64
Figure 2.2: Missing piece puzzle	67
Figure 2.3: Pythagorean proof by rearrangement	67
Figure 2.4: The von Koch snowflake	70
Figure 3.1: Square with diagonal	76
Figure 3.2: Square and diamond	79
Figure 3.3: Effect of framing on orientation	82
Figure 3.4: Intrinsic orientation	83
Figure 3.5: Sequence of triangles	84
Figure 3.6: Deviant orientation	85
Figure 3.7: Square and diamond with main symmetry axes	87
Figure 4.1: Diagram for Euclid's Proposition I.1	126
Figure 4.2: Circle with diameter	158
Figure 4.3: Pair of points	175
Figure 4.4: Straight line drawn between points	176
Figure 4.5: Step one of uniqueness demonstration	177
Figure 4.6: Step two of uniqueness demonstration	178
Figure 4.7: Circle with radius	192
Figure 4.8: Right angles	194
Figure 4.9: Diagram for fifth postulate	202
Figure 4.10: Animated diagram for fifth postulate	203
Figure 5.1: Diagram for Euclid's Proposition I.32	209
Figure 5.2: Demonstration of Euclid's Proposition I.13	231
Figure 5.3: Demonstration of Euclid's Proposition I.15	233
Figure 5.4: Animating the diagram for Euclid's Proposition I.32	236

Chapter 1: Geometrical Insight

1 Introduction

A familiar proposal, one that lies at the core of traditional rationalism, is that human beings are able to attain knowledge of necessary truths by means of the mental act of *insight*.¹ The basic picture of insight is roughly as follows: One becomes consciously aware of the subject matter involved in a proposition, in some sense ‘bringing it before one’s mind’, and by attending to that subject matter in the right way, one is able to *see* or *grasp* or *apprehend* or *recognize* that the proposition concerning it must be true. Insight is distinguished from both empirical observation, on the one hand, and from discursive reasoning, on the other. It is held to differ from empirical observation in respect of the *a priori* character of the justification it provides. It is held to differ from discursive reasoning in virtue of its ‘directness’ or ‘immediacy’—the perception-like characteristic whereby one enjoys a presentational² awareness of the subject matter of one’s judgment. Insight is further held to enjoy a special kind of epistemic security or certainty, presumably in connection to its *a priori* character.

In this dissertation, I will be concerned only with *geometrical* insight—that is, insight into geometrical truth. In particular, I will be concerned with basic Euclidean

¹ Terminology here varies. What I will call ‘insight’ is also often called ‘intuition’. Both terms, when used to denote the phenomenon under consideration, are often preceded by a qualifier such as ‘rational’ or ‘*a priori*’. When the truths in question are those of mathematics, it is common to speak of ‘mathematical intuition’.

² Cf. Chudnoff (2012), whose use of the term ‘presentational’ corresponds closely to the sense I have in mind here.

plane geometry, roughly corresponding to the subject matter of Book I of Euclid's *Elements*. In this chapter, I will argue that we have good reason to take seriously the proposal that there is indeed a real phenomenon of geometrical insight. In Section 2, I begin by setting out some of the main reasons for skepticism about the very idea of geometrical insight, and offering some partial responses to those objection. The full responses to these objections will not be apparent until later in the chapter, and in some cases, until later chapters. In Section 3, I provide a phenomenological perspective on geometrical insight, as part of an overall case that our phenomenology, in relation to the relevant geometrical examples, gives us at least strong *prima facie* reason for taking the geometrical insight proposal seriously. Finally, in Section 4, I will consider Plato's view of geometrical insight as a kind of 'recollection'. My aim here will be to show that, while *prima facie* Plato's view seems to be in tension with a naturalistic picture of the world, there is a plausible reading of Plato's view which points the way towards an appealing account of geometrical insight.

2 Reasons for skepticism

The very idea of geometrical insight is likely to be met with a high degree of skepticism on the part of contemporary philosophers. Since my primary aim in this chapter is to argue that geometrical insight is indeed quite real, it is important to be clear up front about the primary reasons that underlie this contemporary skepticism. In my view, the case against geometrical insight can be usefully separated into five distinct lines of criticism. First, we lack sufficient motivation for believing in

geometrical insight in the first place. Second, our first-personal experience in the relevant cases is suggestive not of insight, but rather of uncertain empirical inference. Third, our first-personal experience in these cases does suggest *a priori* justification, but due not to insight but rather to discursive reasoning. Fourth, geometrical insight is at odds with naturalism, because it depends upon a mysterious, *sui generis* form of cognitive access to a realm of abstract objects. Fifth, geometrical insight appears to be mysterious from a purely psychological perspective.

I believe that each of these challenges can be met. To be sure, doing so adequately will require embracing a picture of geometrical insight that is in certain crucial ways more modest than the traditional view. Some of the stronger claims traditionally made about geometrical insight will need to be qualified in significant ways, and in some cases, rejected outright. But I believe that this will still leave us with a view of geometrical insight that is substantially true to spirit of the proposal put forward, for instance, by Plato.

In this section, I will set out the main lines of criticism against the very idea of geometrical insight, and explain how they will be handled by the view to be put forward in this chapter. The arguments in this section, therefore, are largely prospective, anticipating the account to be developed below, which aims to answer these objections. In relation to certain points, the full response to these objections will not be evident until the discussion undertaken in later chapters.

2.1 *First objection: lack of motivation*

An initial criticism of the idea of geometrical insight is just the straightforward objection that there seems to be little to motivate the proposal in the first place. Contemporary philosophers might understandably take the view that the notion of ‘geometrical insight’ (along with ‘insight’ more generally) is merely a historically entrenched philosophical dogma, one motivated by nothing more than an unwarranted optimism about human epistemic abilities. In order to confront this criticism, it is important to consider what motivation Plato offers for taking the proposal seriously. The main motivation Plato offers is provided by the famous geometrical demonstration in *Meno*, which we will consider in detail in Section 4.

In the passage, Socrates leads his student, an uneducated slave who has had no prior training in geometry, eventually to arrive at the correct answer to a geometrical problem. Since Socrates brings this about by doing nothing aside from posing questions in relation to drawn figures, he claims that the correct answer must have arisen ‘from within’ the student’s soul—in particular, it was not reached by receiving testimony from a teacher already in possession of the relevant knowledge. A frequent objection to Socrates’ inference here—one that is invariably raised by my undergraduate students, when I teach the passage—is that the ‘experiment’ has not been conducted in a methodologically sound manner. For Socrates has posed *leading* questions, ones that might well allow the student to infer, with high reliability, *which* responses Socrates *wants* to receive. In many cases, the correct response is

conspicuously encoded within the question itself, which from a pragmatic perspective, often seems only to ask for affirmation: “Why yes, Socrates!”

This is, to be sure, a reasonable objection to the methodological soundness of the demonstration, regarded *qua* experiment. But to take this as an objection to Plato’s case for geometrical insight is to suppose that this case rests on our taking the passage to accurately report the results of a well-conducted experiment, and it is far from clear that this is the best reading available. As one commentator observes:

There is an alternative approach to the text which undercuts this reaction. For by following the text supplemented by diagrams, one can discover for oneself the geometrical theorem as it might have been discovered by the slave if he had complied with Socrates’ request to give as answers only what he genuinely believed (83d2) rather than what he guessed Socrates believed; or, if one already knows the theorem, one can see how it *could* be discovered that way by someone not already in the know. We do not have to be convinced that people in the slave’s position, answering as the slave does in the text, would not be picking up latent information conveyed by the manner of questioning, in order to become convinced of the possibility of discovery for oneself. We can become directly acquainted with this possibility merely by following the exchange. (Giaquinto 1993, 82)

On this reading, Plato intends for his readers to rehearse the steps of the geometrical exercise for themselves, and thereby, to experience geometrical insight firsthand.

One reason to favor this reading is that it harmonizes with the very phenomenon Plato is drawing to our attention: that of grasping a truth *for oneself*—by one's own lights—as opposed to relying on testimony from others. In Section 4, I will argue that Plato's central concern in the dialogue is in fact to investigate precisely this contrast, and to argue for the epistemic superiority of the former, internalist, variety of justification over the latter, externalist kind. It would be deeply ironic, even thematically incoherent, if Plato were to argue for the importance of recognizing the truth for oneself, by asking the reader to trust that the passage in *Meno* is an accurate recounting of an experiment actually performed. The principle of charity suggests that we ought to interpret Plato so that his argumentative methods are understood as aligning with his argumentative aims—that is, as appealing to what his readers experience for themselves when they follow along with the demonstration, rather than on the accuracy of his testimonial reports.

This impression, that Plato really is inviting his readers to experience geometrical insight for themselves, is further reinforced by reflecting on how *well suited* the passage is to serving this very purpose. Here is a quote from another commentator:

In a sense, it doesn't matter that the slave sees it; what matters is that we do.

Repeatedly, when I have taught the passage, someone gasps or even cries out.

The impact of the proof is unquestionable. We see that it *has* to be so—that it is not a matter of convention, or custom, or even an empirical fact. It is seeing *this*—that it *has* to be so—that is at the heart of the passage, and the dialogue, and, I believe, Plato’s lifework.... When I reflect on my own experience, it is clear that the perception of necessary truth involves a kind of intellectual *phenomenology*—that necessary truth has a distinct *feel*, especially when it is given elegant and economical expression. This is what prompts the gasps, or the involuntarily raised eyebrows, in the classroom. (Zwicky 2009, 47)

It is noteworthy how closely this *phenomenological* description of ‘getting’ the visual proof in *Meno* corresponds to the *theoretical* description of geometrical insight briefly sketched in the introductory section—as presentational, *a priori*, and epistemically certain. If we were to accept both descriptions, it would be difficult not to conclude that the defining features of geometrical insight are in some way directly manifest in the very phenomenology of the experience. What is suggested in the above quotation is threefold: that what has so impressed Plato about geometrical discovery is its distinctive phenomenology, that this phenomenology is indeed impressive, providing at least *prima facie* subjective evidence in favor of the geometrical insight proposal, and that Plato’s primary aim in going through the demonstration is therefore to prompt *us* to experience this phenomenology for ourselves. In Section 4, when we consider the *Meno* demonstration directly, we will

revisit the matter of its associated phenomenology, and reflect on the lessons Plato draws from the latter.

I do not mean to suggest, of course, that the entire case for geometrical insight can be rested on first-personal experience. Indeed, this is far from the case; theoretical considerations will matter a great deal. But it does seem that first-personal experience is what initially motivates the proposal of geometrical insight, by providing *prima facie* evidence for its reality. Accordingly, before considering Plato's own example in Section 4, Section 3 will be devoted to first-personal reflections on the phenomenon.

2.2 *Second objection: unimpressive phenomenology*

Of course, appeals to first-personal experience are notoriously difficult to adjudicate. The skeptic about geometrical insight might attempt to defuse the apparent motivation from first-personal experience by taking one of two tacks. First, one might deny that one's own first-personal experience, when one considers examples of the sort that have impressed Plato, is really suggestive of geometrical insight in the way that has been supposed. This would effectively result in a standoff, since of course neither party is in a position to evaluate the phenomenology of the other. Second, one might accept that one's first-personal experience is indeed suggestive of geometrical insight, but claim that there is no reason to presume that one's phenomenology in these cases is reflective of the actual epistemological qualities of the experience.

I will be focused primarily on responding to the first of these two options, but I will briefly comment on the second one here at the outset. I think that in cases where

one does genuinely experience the phenomenon in question, it is difficult to maintain that one's phenomenology does not provide strong *prima facie* reason for believing that one is having an experience that genuinely possesses the qualities of geometrical insight. For in these cases, one does not merely experience a *brute sensation* of certainty or aprioricity or necessity, one that can at most be said to *accompany* belief in the relevant geometrical proposition. Instead, as suggested above, the experience is one in which the subject *sees why* the geometrical result *must* be true. That is, one has a presentational awareness of the features and relationships that secure and underlie the sense of necessity, and which suffice to make one certain about one's judgment. The experience, that is, is a transparent one. We do not merely note the presence of a *feeling* of certainty about our judgment, whose origins and basis are opaque to us. Rather, the experience is one in which we enjoy an intimate, presentational awareness of the spatial relationships that serve to *ground* our sense of certainty. I think, therefore, that when critics object that there is no reason to take the phenomenology of geometrical insight as reflective of the actual epistemological qualities of the experience, it is because they are simply failing to experience this phenomenology.

The real obstacle to granting the first-personal motivation that Plato claims, then, is that some subjects, when confronted with the relevant examples, fail to recognize anything in their own phenomenology that is suggestive of geometrical insight. Anecdotally, I can report that this is a fairly common initial reaction to the diagrammatic 'proofs' I present in undergraduate lectures and at conference presentations. Typically, these individuals agree that consideration of the diagram

does make the target proposition seem ‘plausible’, and does perhaps incline them in the direction of believing it, but that they lack any phenomenology of necessity or certainty, of the sort described above in relation to the *Meno* demonstration. As such, their first-personal experience seems compatible with the view that the geometrical beliefs are reached in these cases by means of ordinary, highly fallible, empirical inference. This may lead them to suspect that proponents of geometrical insight are merely being overly eager in interpreting their own phenomenology, self-ascribing a sense of ‘necessity’ and ‘certainty’ that is simply not there to be found in the actual experience, soberly appraised.

I think this discrepancy can be explained by the straightforward proposal that there are multiple ways in which an experience of a diagram or a visual image is apt to incline one towards a geometrical belief, and that only some of these qualify as geometrical insight. Geometrical insight does not arise in any robust way without a significant degree of careful, directed attention by the subject. Moreover, the very same images that are held to support geometrical insight, when one attends to them carefully, are also capable of yielding less impressive epistemic phenomena, such as visually-based hunches. This should not be at all surprising to a proponent of geometrical insight. By way of example, consider a diagram of the sort that might allow a subject to grasp, through geometrical insight, that an angle constructed in a certain way must *necessarily* be a right angle. If the diagram is drawn with any reasonable degree of accuracy, it will surely display an angle that can be seen at a glance to be *approximately* right. Since the subject will arrive much more readily at

this latter kind of observation, and since it will often be sufficient at least to incline the subject towards belief in the target proposition, there is an understandable temptation to presume that the mere sense of plausibility or suggestiveness that the experience of the diagram delivers at first glance comprises all that it has, epistemically, to offer. In this way, I think there is a systematic tendency for subjects to underestimate the power of visual images to yield geometrical insight, because they are apt to conclude too hastily that they have already discovered what is there to be found. The contrast drawn here, between what initially *seems plausible* when confronted with a diagram, and what one can genuinely grasp *via* geometrical insight, is a central concern of Plato's, as we will see in Section 4.

Anecdotally, I can report that in at least some cases in which individuals initially find a diagrammatic example to deliver only a sense of plausibility, they do eventually come to appreciate the example in a way that strikes them as more phenomenologically suggestive of geometrical insight. In other cases, they fail to achieve this result in relation to the original example, but succeed in doing so with simpler examples. I don't regard these anecdotal reports as having much in the way of probative value, however. As I've already suggested, there is no adequate substitute for experiencing the phenomenon for oneself. That is why Section 3 will be devoted to first-personal reflection on a simple example.

2.3 *Third objection: mere discursive reasoning*

One sometimes encounters a different sort of skeptical reaction to diagrammatic examples of the relevant kind. The skeptical reaction considered just above agrees with the proponent of geometrical insight that the mental act or episode that produces belief in these cases is presentational, but denies that it is either *a priori* or certain. Alternatively, it is sometimes claimed by subjects that when they consider these examples, they enjoy subjective justification that is perhaps both *a priori* and certain, but which is not presentational. Here the claim is generally that, although a diagram or visual image may stimulate or provoke or suggest a certain line of reasoning, it is in fact this discursive reasoning alone that carries the burden of justifying belief in the geometrical proposition, yielding the conclusion in a way that is justificationaly independent of any contributions from visual imagery, and indeed from any form of ‘presentational’ awareness of the geometrical subject matter. In short, the claim is that what one really experiences is not geometrical *insight*, but rather ordinary discursive reasoning about geometry.

This is not a novel proposal. Indeed, Leibniz famously advanced precisely this claim regarding Euclid-style geometrical practice:

The force of the demonstration is independent of the figure drawn, which is drawn only to facilitate the knowledge of our meaning, and to fix the attention; it is the universal propositions, i.e., the definitions, axioms, and

theorems already demonstrated, which make the reasoning, and which would sustain it though the figure were not there. (1704, 403)

Leibniz, however, is simply mistaken on this point. As is clearly established by Manders' (2008) analysis, the deductions found in Euclid's proofs will not go through without the contribution of the spatial relationships seen to obtain in the accompanying diagrams, which serve to justify what would otherwise be glaring inferential gaps. Moreover, this point generalizes to ordinary geometrical reasoning of the sort here under consideration. This is not to deny that the content of Euclidean geometry can be captured by a thoroughly deductive axiomatic system; indeed, both Hilbert and Tarski have achieved this, respectively, in second-order and first-order axiomatizations. Nor is it to deny that it may be possible to provide a formalized rendering of Euclid's own proof procedures, by carefully selecting axioms to replace the justificational contribution that would ordinarily be made by visual diagrams; indeed, recent work in proof theory promises to achieve precisely this result (e.g., Avigad et al., 2009). It remains the case, however, that such purely formal-deductive treatments of Euclidean geometry do not provide a route for ordinary reasoners to grasp geometrical propositions in the cases here at issue. If an ordinary reasoner reflects on a geometrical problem like the one posed in *Meno*, and comes to grasp the solution in a way that is certain and *a priori*, they must either be relying on a presentational awareness of spatial relationships evident from the diagram, or else

their conclusion will be conditional on assumptions which are themselves nontrivial, and which fall within the purview of geometrical insight.

The upshot is that discursive reasoning alone cannot provide certain and *a priori* justification for geometrical belief in the cases at issue. In that case, why do some subjects believe this to be true of their own experience? One possibility is that they are reaching the result through a combination of insight and discursive reasoning. Indeed, this is perhaps the most natural way to follow Euclid's diagram-based proofs: The overarching frame is a discursive, deductive one, and diagram-based insight is invoked, tacitly, at various stages, in order to justify unstated premises on which the argument would be seen to depend, were it to be spelled out in fully explicit detail. If one approaches the Euclidean proof in this manner, primarily attending to the deductive track set out in the text, and looking to the diagram ostensibly just to remind oneself of the concrete meaning of the textually encoded inferences, it is easy to overlook the fact that one's visual understanding of the diagram is in fact making essential contributions to the deduction. For in this case, the contribution of insight becomes fragmented and piecemeal, and is thereby reduced to judgments that are so visually obvious (for instance, that a line drawn to certain specifications will have to lie *inside* a given angle, rather than outside it) that it is easy to miss the fact that one is relying upon them at all. I presume this is what accounts for Leibniz's mistaken impression that the deductions set out in the text suffice to establish Euclid's conclusions, independently of any contribution from the diagram itself.

What is indicated by the example of Euclid's proofs is that when it comes to the propositions of Euclidean geometry, there is generally available a *partial* discursive argument for those propositions. While this partial discursive argument will in almost every case contain gaps in its deductive structure, these gaps are readily overlooked because they can be bridged by premises whose truth is visually *obvious*, and is therefore easily taken for granted. As such, so long as one has available to one's awareness a diagram drawn on paper, or one entertained in visual imagination, it is relatively easy for one to rely on visual understanding in one's reasoning, without realizing that one is doing so.

The response to the objection, then, is that subjects may well be correct in thinking that they have reached the result in a way that is *partially* discursive. In this case, however, they have not grasped the result itself through geometrical insight—as such, they are simply experiencing a different epistemic phenomenon, the existence of which does not count against that of geometrical insight. It is worth taking a brief digression in order to reflect further on the nature of this contrast between insight and (partially) discursive reasoning in geometry.

When one reasons in the manner I have described above, such that the thematic focus is on the discursive text, with only occasional, isolated appeals to visual understanding, one's grasp of the conclusion *itself* does not count as an instance of geometrical insight, merely because this conclusion was reached in a way that *depends* upon geometrical insight, in relation to certain (perhaps implicit) premises in the discursive argument that supports the conclusion. This is because, in

order for one's grasp of a given proposition to qualify as geometrical insight, it is necessary that one have a presentational awareness of the geometrical relationships that constitute the truth-makers for that very proposition.

There is, I think, a way of approaching Euclid's proofs that does yield geometrical insight in relation to the proposition proved. In fact, in my experience this is something that emerges naturally as I become more familiar with a Euclidean proof, and arrive at a more complete understanding of it. When I first encounter a given geometrical proof in Euclid, I find myself constantly shifting attention back and forth between the text and the diagram. I first look to the text to read the current assertion in the deductive sequence, and then attend to the diagram in order to interpret the concrete spatial meaning of what Euclid asserts at that step, decoding Euclid's reference to angle 'ABD', for instance, by noting which letters label which points on the diagram, and apprehending, now in an ostensive way, which angle Euclid has in mind. While, as noted above, some of Euclid's inferences depend critically on the spatial interpretation of his statements, this is by no means always the case. Often, it is unnecessary to consider the spatial meaning at all in order to verify that a given inferential step is deductively valid. For instance, Euclid sometimes reasons by substitution, justifying the replacement of one angle by another in an equation by relying on a prior textual assertion that the angles are equal in magnitude. The validity of such an inferential step can be verified by attending only to the *labels* for the angles, without any concern for which spatial objects these labels represent.

If one proceeds in this manner, however, considering the spatial meaning of Euclid's statements only when it is strictly necessary in order to verify that the current inferential step is deductively valid, it is difficult to arrive at the proof's conclusion feeling as though one has a genuine understanding of why the proof has succeeded. In practice, coming to understand a Euclidean proof seems to require that one interpret the spatial meaning of each textual statement in succession, and initially, this means shifting attention back and forth between the text and the diagram.

In my own experience, however, when I continue to rehearse a Euclidean proof in this manner, I find myself progressively devoting less attention to the text, and more to the diagram. Having previously decoded each of Euclid's symbolic statements into their spatial meanings, I now find myself understanding these assertions *only* in reference to their spatial interpretations. It is no longer necessary for me to pay attention to the labels, for I am no longer translating the text into its spatial content. After all, I already know what Euclid is saying, in reference to the geometrical situation displayed by the diagram. Moreover, by considering the diagram, I find I can *see why* his inferences succeed. What was at first grasped abstractly as a *textual* substitution of *symbols* now becomes understood more concretely as a *spatial* substitution of *angles*, which I apprehend demonstratively. I find, in fact, that Euclid's entire course of reasoning can be understood in a way that is based on my direction of visual attention to the diagram itself. By in this manner arriving at a concrete, spatial understanding of Euclid's proof, I have effectively transformed it into a thoroughly *visual* proof: I see directly, in the diagram, the spatial

relationships that make the result true. In this way, I am able to attain an integrated, synchronous, synoptic appreciation of the proof *as a whole*. I claim that when I appreciate the proof in this way, I do grasp its conclusion by means of geometrical insight. Interestingly, this results naturally from my arriving at a complete understanding of what is, to begin with, a largely discursive proof.

So that I am not misunderstood, it is worth acknowledging the following point: I do not claim that discursive or linguistic content needs to be altogether absent from my phenomenology, when I grasp a geometrical proposition by means of insight. Indeed, when I visually attend to the diagram in a way that allows me to grasp why the geometrical proposition is true, I may well experience myself rehearsing thoughts in a discursive mode, for example: ‘If I were to substitute *this* angle for *that* one...’. In this case, however, it seems clear that these linguistic or quasi-linguistic contents are playing a supporting, ‘scaffolding’ role, serving to fix visual attention on the relevant relationships in the spatial, geometrical situation presented by the diagram. They do not themselves bear the burden of providing justification, but rather serve to organize and give shape to inferences that have an essentially visuo-spatial character. To put the point in slightly different terms, linguistic thought may serve to keep a running record of judgments made about relevant spatial relationships in the geometrical situation, but the justificational ground for those judgments is visual, based in what can be seen in relation to the diagram itself.

2.4 *Fourth objection: mysterious access to abstract objects*

The objections to the idea of geometrical insight considered above relate to the question of whether first-personal experience provides motivation for taking the phenomenon seriously. I now turn to a pair of related objections of a more theoretical character. These objections question whether the putative phenomenon of geometrical insight can be squared with a naturalistic picture of the world.

The first objection can be traced to arguments advanced by Benacerraf (1973). The objection points out that the propositions of geometry are concerned with abstract mathematical objects. Since these objects, if they are real at all, presumably exist outside of the spatiotemporal realm inhabited by human minds, it would seem that any justification of true beliefs concerning geometrical subject matter will require some mode of cognitive access to the platonistic realm of being that these objects inhabit. If so, the objection goes, the mechanisms through which we are able to enjoy this access to objects in an abstract mathematical realm are utterly mysterious. We don't even know how to begin offering an explanation for their operations.

In brief, my response to this objection is that, indeed, we do not enjoy any special form of cognitive access to a realm of abstract geometrical objects. On the view that I propose to defend, geometrical insight does not provide us with any justification for believing that the subject matter of geometry—consisting of the geometrical objects themselves—actually exists. Strictly speaking, what we grasp through geometrical insight are propositions that are *subjunctive* in character. That is, we apprehend what *would* be true of geometrical objects (and geometrical space) as

we envision them, if those objects (and that space) actually existed. While I will often omit mention of this subjunctive framing, for ease of exposition, it should be borne in mind that at no point in this dissertation should I be interpreted as claiming that we possess any knowledge concerning the metaphysical existence of abstract geometrical objects or of ideal Euclidean space. As such, the purely epistemological proposal that I will be putting forward is intended to be entirely neutral regarding all ontological matters concerning the existence and metaphysical nature of geometrical subject matter. Of course, this means that the version of the geometrical insight proposal I will defend makes a far weaker claim about geometrical knowledge than traditional versions such as Plato's, at least in this particular respect. That this still leaves us with an interesting account of a nontrivial sort of geometrical knowledge—and one that remains substantially true to the spirit of Plato's own account—is something I hope to demonstrate in the remainder of the chapter, and the dissertation as a whole.

2.5 *Fifth objection: sui generis psychology*

Our final objection to consider is that, even if we can set aside the concern about cognitive access to abstract objects in the manner suggested above, the geometrical insight proposal still seems to require us to postulate a mysterious, *sui generis* kind of psychological state. For the very characterization of geometrical insight as both presentational and *a priori* seems to confront us with puzzle. How can an epistemic mental act possibly have both of these features at once? The familiar case of justification based on 'presentational' awareness is that of ordinary perceptual

knowledge, acquired through sensory channels—but this sort of knowledge is of course not *a priori*. Conversely, the familiar case of *a priori* justification is that gained through discursive reasoning, which lacks the requisite presentational character. So the proposal of geometrical insight, which attempts to combine both of these features within a single epistemic act, seems to run against the natural grain of familiar psychology. As I will attempt to show in my interpretation of Plato’s view below, this apparent tension can indeed be reconciled, by regarding geometrical insight as grounded in an experience of *visual understanding* that is plausibly regarded as combining these presentational and *a priori* aspects.

3 A phenomenological perspective

As we noted in the Section 2.1, Plato appeals to first-personal experience in order to offer motivation for believing in geometrical insight. And as we saw in Sections 2.2 and 2.3, subjects often report an absence of the kind of phenomenology that would serve to motivate the proposal. As I suggested there, the proper diagnosis may well be that these subjects are simply not having experiences of the relevant kind. Therefore, in this section I want to take the reader through a very simple example of geometrical insight, exploring it from a first-personal vantage point. This will provide an intuitive familiarity with some of the features of geometrical insight that will take on theoretical importance within the account provided by Plato.

Suppose I imagine a circle, with a straight line drawn through its center. It seems clear to me, bringing the relevant image to mind, that the line will divide the

circle into exactly two parts. If I now pose to myself the question of whether these two parts will be congruent to one another, I find myself judging very confidently that they will *have* to be congruent—indeed, *precisely* so. And if I consider whether this would still be the case if the circle were of some different size, or if the line were drawn through the center at a different orientation, it seems clear to me that these variations would not interfere with the congruence of the circle's parts. I now possess a belief about geometry that I hold with a high degree of confidence: *In general, a straight line drawn through the center of a circle will divide the circle into two congruent parts.*

When I consider the question of what justifies my holding this belief with such confidence, it seems to me that my confidence is justified by the experience I have when I entertain the relevant geometrical situation in visual imagination. When I *visualize* a circle with a straight line drawn through its center, and I attend to what I am visualizing in a certain way—specifically, in the context of wondering whether the parts on either side of the line are congruent—I have an experience of things falling into place, with a sense of inevitability. I seem to find, within that experience, immediate justification for believing that the parts on either side of the line will indeed be congruent—and for believing that the size of the circle, and the orientation of the line, will not make any difference when it comes to this relationship. I am tempted to say that when I bring to mind a visual image of the geometrical situation, and I consider it in the right way, I can simply *see* that the two parts of the circle could not fail to be congruent.

If I am now asked to justify to someone *else* that a straight line through the center of a circle will divide it into two congruent parts, I may be unable to articulate, in a fully explicit manner, my reasons for holding this belief. After all, my reasons are private ones. Their basis lies not in any deductive argument that I could hope to communicate, but rather in a visual-imaginative experience I am undergoing. When I pose to myself the question of what justifies my holding this geometrical belief, it seems a sufficient answer to simply point to my own experience. Of course, that is not possible in the interpersonal case—that is, not unless my friend is willing to settle for secondhand justification, by testimony, and simply take my word for it that *I* am in possession of justification of the firsthand variety. It seems that the best I can do, then, is to point to the geometrical situation I am having the experience *of*, in the hopes that my friend will experience it in a similar way. I can do this by drawing a suitable diagram, and presenting it along with the necessary stipulations: that the diagram is supposed to depict a *circle* and a *straight line*, and the straight line is supposed to pass precisely through the circle's *center*.

When I now look at the diagram I have drawn, I have a visual experience that seems to be, in all relevant respects, similar to the visual-*imaginative* experience I was having previously. It is true that I am now also having a visual experience *of* a physical object—the drawn diagram itself—which I take to be veridical. But *that* experience—or that *aspect* of my experience, anyway—strikes me as quite beside the point. What seems important is not my experience *of* the diagram, but my experience of what I see *in* the diagram: the same geometrical situation that I had previously

been imagining. When I now attend to what I see *in* the diagram, in the same way I was previously attending to the visualized situation, I experience the same sense of things falling into place, in a way that seems to justify my belief in the geometrical proposition. Considering both as experiences *of* the geometrical situation, the only significant difference I notice is that the diagram-aided experience seems more vivid than the imaginative one, and the justification it provides, accordingly stronger.

Of course, there is no guarantee that my friend, confronted by the same diagram, will have the same experience of ‘things falling into place’ in a way that seems to provide justification for believing the geometrical proposition. If not, there is little that I can do, except to try to gesture at the features and relationships that seem, within my own experience, to justify this belief. Figuring out how to do this is not a trivial matter. At first, I might simply point to the diagram, repeat the target judgment, and assert that an experience of the former justifies the latter: “Don’t you see that the arrangement *just fits together* so that this *has* to be the case?” After reflecting a bit on my experience—on *why* the arrangement seems to ‘fit together’ in a way that justifies the judgment—I might be able to say something a bit more helpful: “Suppose the circle were *folded* along the straight line—wouldn’t the parts of the circle on either side of the line now have to end up getting folded exactly onto one another? Alternatively, suppose the circle were *rotated* by precisely a half-turn about its center—wouldn’t each of the parts now have to occupy exactly the same space that the other one did before the circle was rotated?” In defense of the generality of the proposition, I might add: “Suppose the circle were larger or smaller, or the line were

drawn through its center at a different orientation—couldn't we compensate for these changes simply by adjusting the viewing distance or the orientation of the diagram itself? Isn't this very diagram, then, adequate to represent all of the possible variations that satisfy the initial characterization *circle with a straight line drawn through its center?*"

As these examples already suggest, my attempting to *point out* the relevant features and relationships might well result in my formulating something partially resembling an *argument*, one with the target judgment as its conclusion. Might this not imply that it is, after all, possible for me to *articulate* my justification for the judgment—or at least some nontrivial part of it—by providing an argument that captures that justification? Perhaps, then, my justification is not inherently bound to a private, 'presentational' experience? Suppose I decide that one of the key features to appreciate, in order to grasp the geometrical proposition, is the *reflection symmetry* of the circle about the straight line. Indeed, by attending to this symmetry, I seem immediately to grasp that the target proposition itself must be true, by appreciating the way that this symmetry seems to *force* congruence upon the circle's parts, with a kind of necessity. (Making this relationship salient had been the point of my earlier suggestion that my friend consider the circle being *folded* along the line.) I might now attempt to explicitly formulate the relationship I have appreciated, in the form of a deductive argument. Taking 'C' to name the circle and 'L' to name the straight line drawn through its center, I could argue as follows:

- (P1) C is symmetrical about L.
- (P2) In general, if x is symmetrical about y, then y divides x into two congruent parts.
- (P3) If C is symmetrical about L, then L divides C into two congruent parts.
(from P2)
- (C) L divides C into two congruent parts. (from P1 & P3)

Of course, a deductive argument can only be as convincing as its basic premises. Here (P1) is a fact about the geometrical situation that seems *so* basic that I can, apparently, *only* point to it—if my friend does not accept that the circle will have to be symmetrical about the line through its center, there seems to be little more that I can say. And my visual experience does not seem to *assume* but rather seems to *show me* that the circle indeed has this symmetry. This already suggests that an argument will not be able to *replace* the justification provided by my visual experience. Still, it might be able to replace *some* of it.

The universal statement (P2), in contrast to (P1), might reasonably be taken to be analytic, following from the concepts of *symmetry* and *congruence*. So let's suppose that there is no difficulty accounting for knowledge of (P2). (P3) follows from (P2), by instantiating the variables x and y to the constants that appear in (P1), and (C) then follows by applying *modus ponens* to (P1) and (P3). Suppose I present this argument to my friend, and my friend grants that (P1) is visually obvious and that (P2) is analytic, and then validly reasons from these premises to the conclusion (C), in

the way indicated. Is my friend now justified in believing (C), in the same way I was justified in believing it, in the first place?

When I consider this question, while rehearsing for myself the steps of reasoning through the argument, it seems to me that while this procedure *does* yield justification for believing (C), it is not the same *kind* of justification that I had initially. In particular, there seems to be an important phenomenological difference. In following the argument, I start by separately confirming the truth of (P1) and (P2)—in the first case, by looking to what I see in the diagram, in the second (let's suppose) by reflecting on the concepts involved. I then instantiate the variables in (P2) to the constants in (P1), 'plugging in' C for x, and L for y, in order to derive (P3). What strikes me about my phenomenology while I am performing this step is a feeling of uncertainty about the conclusion I reach. It is not that I feel uncertain about the step *itself*, while I am carrying it out. Indeed, the operation that takes me from (P2) to (P3) is purely syntactic, and I feel confident that I can perform it reliably, given the clarity with which I apprehend (P3)'s syntactic form. But precisely because I am only attending to the syntax of (P2) while carrying out this operation, it occurs to me in this moment to wonder whether the *formulation* of (P2) that I now hold before my attention was correctly stated in the first place. Even if the judgment I previously attempted to formulate as premise (P2) had been perfectly justified, am I certain that I did not made a mistake in symbolically 'transcribing' this judgment, say by mixing up the variables 'x' and 'y', and placing them in the wrong order?

I find myself confronted by a similar doubt when I perform the final step, deriving (C) from (P1) and (P3) by *modus ponens*. While I am performing this step, I feel very confident that I am correctly carrying out the appropriate *syntactic operation* on premises (P1) and (P3), *as formulated*. From this ‘syntactic’ point of view, however, I find that I have lost sight of my reasons for having accepted (P1) and (P3) as expressions of true statements in the first place, those reasons being essentially bound up with the spatial *meanings* of the premises. Again, perhaps I have made an error in transcribing the thought.

When I rehearse the argument in this way, then, I find myself grasping the truth of the premises independently, then operating syntactically on the forms of their statements, and eventually deriving a formula that encodes the conclusion. When I finally interpret the meaning of this formula, I find myself looking back to the diagram, and thinking: “I guess those parts either side of the circle must be congruent, then. After all, that is the meaning of the statement I seem to have derived.” This gets me to belief in the proposition, to be sure, but I do not thereby replicate my original experience of *seeing* that it must be the case.

To be sure, this is an implausibly ‘mechanical’ way of following the argument, by carrying out inferences without paying attention to the meaning of what is said (with the exception of the logical vocabulary, whose content must still be interpreted in order to perform the inferences). But in this case, the ‘meaning’ is a spatial, geometrical one—to attend to this in the course of following the argument, would be to help it along with the aid of visual understanding. In order to pinpoint the

relevant contrast, it seems appropriate to consider precisely this experience of following the argument ‘mechanically’. And it is clear that when I do so, I seem to miss out on the distinctive sort of justification that I am able to enjoy by means of visual understanding. What is lacking is something that seems both phenomenological as well as epistemological: The different steps of the argument are not sufficiently integrated in my experience, and correlatively, I am unable to achieve a synoptic understanding of the demonstration as a whole. I feel I have no appreciation of *why* the two parts of the circle must be congruent. Instead, I am simply faced with the fact that my reasoning process—in some way or other—has produced this result as output.

A closely related point bears specific emphasis: When I follow this argument in order to justify belief in the geometrical proposition, I experience a sense of doubt. This seems to result from the fact that my appreciation of the truth of the basic premises, by appeal to their semantics, occupies a distinct cognitive moment from the purely syntactic operations by which I carry out inferences on those premises. I wonder, in particular, if I might have made an error in transcription, for instance, by mixing up the symbols in some way. This sort of error is commonplace in such ‘mechanical’ procedures—one may easily forget to ‘carry the one’ in performing the algorithm for addition of large numbers, for instance—which is why it is good practice to perform a ‘sanity check’ to ensure that the result produced by such a mechanical process is a plausible one.

In the present context, however, to perform a sanity check would just amount to my comparing the output of my reasoning process with what is intuitively apparent to me on the basis of my visual understanding, to make sure that the two align. If this is necessary in order for me to attain certainty about my conclusion, then it is clear that the process of reasoning, taken on its own, does not provide me with the same kind of certainty that I seem to have when I reply just on my visual understanding. Practically speaking, the doubt I feel may seem rather silly—I can, after all, simply rehearse the argument several times carefully, in order to increase my confidence that I have committed no error. But the important point is that, silly or not, the doubt does arise, in a way it does not when I simply attend to what I am able to visually appreciate about the way the spatial parts of the geometrical situation hang together.

4 Plato: geometrical insight as recollection

In this section I consider Plato's account of geometrical insight as 'recollection', which is considered most directly in reference to the geometrical demonstration in *Meno* but also discussed in *Phaedo*.³ On the face of it, Plato's recollection account is essentially in conflict with a naturalistic picture of the world. The proposal Plato seems to advance is that prior to its mortal existence, the disembodied soul was able to enjoy a direct contact with ideal geometrical objects, and now retains a buried memory of this prior knowledge by acquaintance, which needs only to be brought to the surface by the right stimulus or trigger. This view asks us to embrace not only

³ Citations for both dialogues are to Plato (2002).

(unsurprisingly) a platonistic conception of geometrical objects, but also the existence of immortal souls. Since this extreme conflict with naturalism is too high a cost to bear for many contemporary theorists, there is a temptation to dismiss Plato's view of geometrical insight outright. But I submit that to do so would be a mistake.

Whether or not Plato genuinely held the view just outlined is a difficult interpretive question, especially given Plato's frequent appeal to myths in order to present his philosophical views. This interpretive question, however, need not detain us much here. For present purposes, we are interested in whether Plato's recollection view contains ideas that place us in a better position to make sense of geometrical insight as a naturalistically respectable phenomenon. I will attempt to show that this is indeed the case. If we are willing to forego the contentious ontological baggage associated with the 'recollection' idea, what remains of Plato's view provides a very appealing account of geometrical insight as a kind of *recognition*. In particular, Plato helps us to resolve the puzzle raised in Section 2.5: how to make sense of a mental act of insight that has the 'presentational' character associated with sensory perception, as well as the *a priori* character that is usually defined in *opposition* to sensory perception. On the view that emerges out of consideration of Plato's ideas, geometrical insight can be seen to rest on a form of *understanding* that is both presentational and *a priori* in nature.

4.1 Context

In order to appreciate the significance of the geometrical demonstration presented below, it is important to consider it in context of the earlier discussion in the dialogue. Socrates and Meno have been inquiring into the nature of virtue. Meno repeatedly shows impatience with the Socratic method of investigation: He is unwilling to earnestly consider matters by his own lights, and instead simply wants to be told the answers to his questions.⁴ Thus when Socrates urges him to answer for himself what he thinks virtue is (71d), Meno responds by parroting a view acquired secondhand from Gorgias (71e). Again, when Socrates requests that Meno propose a definition of *shape* as a preliminary exercise to defining virtue (75a), Meno flatly refuses, and insists that Socrates just tell him (75b). Meno's attitude here reflects a view of learning as the receipt of *information*, by testimony from those who are somehow already informed.

Given Meno's view of learning, it is not surprising that he reacts in the way he does when Socrates brings the discussion to its moment of *aporia* at 80c-d. Socrates has by this point revealed the inadequacies of each of the definitions of virtue that Meno has acquired secondhand from various sources. Meno is now thoroughly perplexed; the ground has been cleared for philosophical investigation to proceed, unencumbered by the false presumption of knowledge. Socrates, noting again that he himself is as perplexed as anyone regarding the nature of virtue, proposes they begin their inquiry afresh. Meno, who had regarded Socrates' initial profession of ignorance

⁴ This aspect of Meno's character is brought out clearly in Scott's (2005) book on the dialogue, particularly in Chapter 5.

dubiously (71b-c), now realizes that Socrates is intent on standing his ground; he is not going to provide Meno with ready answers that he can simply accept on authority. For Meno, this is enough to establish that the inquiry into virtue is bound to be fruitless: “How will you look for it, Socrates, when you do not know at all what it is? ... If you should meet with it, how will you know that this is the thing that you did not know?” (80d)

The question is posed rhetorically. In effect, Meno is insisting that it is impossible that he should ever come to recognize, by his own lights, that something must be the case. This brings us to the crux of the dialogue. The point of the geometrical demonstration will be to show that, on the contrary, it *is* possible to recognize truth for oneself. First, however, Socrates answers Meno’s question directly, by giving an account of *how* this could be possible (81b-e). This is Plato’s doctrine of learning as *anamnesis*, or ‘recollection’. On this view, the human soul already contains latent knowledge of all things, from its prenatal existence. What we call learning is in fact just the conscious recovery of this latent knowledge, cued by some appropriate trigger. From the discussion in *Phaedo* (73a-75b), it is clear that on Plato’s view, the soul’s prenatal knowledge includes knowledge of abstract geometrical properties, such as *equality*. Thus when we perceive two objects that are roughly equal, even though they can only approximate perfect equality, we are reminded of equality itself—a perfect geometrical property with which our soul has been previously acquainted. Similarly, an imperfect geometrical diagram can trigger the recovery of latent knowledge of geometry, as Socrates will now show.

4.2 *The geometrical demonstration*

The geometrical demonstration takes place in *Meno* at 82b-85b. Socrates will teach one of Meno's slaves a simple proposition of geometry (in effect, the isosceles case of the Pythagorean theorem, reformulated as a construction problem), but will do so without imparting any *information* about geometry. Rather, Socrates will present the slave with a geometrical diagram, and by posing a series of questions in relation to the figure, will bring the student to recognize for himself that the theorem must be true. This will be learning through recollection.


Socrates begins by drawing a diagram of a square (82b). The student immediately appreciates that the figure represented, being square, has its four sides equal, and also that it could be larger or smaller while remaining the same shape (82c). Clearly, the area of any particular square will be a function of the length of the side on which it is constructed, and the longer the side is, the greater in area will be the square constructed from it. Socrates now poses the question: Supposing we had a square twice the area of the given one, how long would *its* side be? The student answers immediately: "Obviously, Socrates, it will be twice the length" (82e)—that is, twice the length of the side of our original square.

The student's initial answer, despite its presumed obviousness, is incorrect. What prompts this incorrect response? The swiftness with which the judgment is issued, together with the apparent lack of any conscious *reason* for so judging (the slave says it is simply "obvious"), suggests as a plausible hypothesis that the student

has reached this judgment by employing an unconscious *heuristic* of the sort widely studied within the ‘heuristics and biases’ research program in cognitive psychology.⁵ In this case, the heuristic assumes that one spatial magnitude (the area of the square) will increase in proportion to another spatial magnitude (the length of the square’s sides). It thereby issues in a swift intuitive judgment, with roughly the content: *twice the length, twice the area*.

Socrates now leads the student to recognize that doubling the length of the square’s side will yield a square with four (rather than two) times the area of the original square (83a-c). He does so by augmenting the original diagram, first taking the square’s base and extending it so it is twice the length of the original base. He then constructs a square on this extended base, which contains the original square inside itself. Drawing lines between the midpoints of the larger square yields the diagram in Figure 1.1.

Figure 1.1: Step one of Plato’s proof



⁵ Tversky and Kahneman (1974) is the classic source; see also Gilovich et al. (2002).

Observing the augmented diagram, the student recognizes that the larger square can be seen as composed out of four smaller squares—the four ‘quadrants’ of the large square—each of which is equal to the original. The student now judges, correctly, that doubling the length of a square’s side yields a square four times the area of the original (83c).

On what basis has the student arrived at this second, correct judgment? Here we can appeal to our own phenomenology, as we follow the demonstration through for ourselves. For when we look at Figure 1.1, *we too* can recognize what the student has now grasped. Taking an arbitrary one of the ‘quadrants’ as our original square, we can readily see the figure as one in which the whole, large square has sides twice as long as those of the original, smaller one. Perceiving the figure in this way, we see that the large square is made up out of four smaller squares, all of which are equal, and one of which is just our original square—so the area of the large square will have to be four times that of the original.

Notice that we arrive at this judgment without scrutinizing the metric properties of the drawn figure, for instance by measuring the four smaller squares carefully (or, less reliably, by ‘eyeballing’ them) to satisfy ourselves they have been drawn precisely equal. Proceeding in this manner would be sensible only provided that we were in the first place confident that the diagram had been drawn to metric precision, with the base of the original square having been extended to *exactly* twice its original length, and a figure *exactly* square having been constructed on the extended base. It is true, of course, that one *could* attempt to learn about the

geometrical relationship between the side-lengths and areas of squares in this manner, through empirical measurement. But that is not how we seem to reach the judgment in this case. For when we reflect on the possibility that, strictly speaking, the empirical drawn figure has none of the metric properties relevant to the truth of our geometrical judgment, this does not seem in any way to undermine our warrant for making that judgment. After all, our judgment concerns not the figure, but the geometrical situation it represents. Just as in Plato's example of recollection, being visually confronted with approximate equals serves to bring equality itself before our mind, in this case, the approximate rendering of the geometrical situation seems in some way to present us with the situation itself. When we consider this situation, in which the large square's base is (exactly) twice the length of the small one, we seem to grasp clearly that the large square will *have* to be (exactly) four times the area of the small one. Our phenomenology suggests that the warrant for this judgment is based in some way on our visual experience of geometrical situation presented by the diagram, but it is not based on our experience of the precise metric features of the diagram itself, taken as a physical, empirical object.

How are we to assess the difference between the two conflicting judgments that have now been given by the student? Both can be regarded as 'intuitive' judgments, in the specific sense that the warrants supporting these judgments seem to be—at least in part—inherently *private*. That is, in neither case is the student in a position to fully articulate his reasons for judging in the way he does, such that he could provide an argument capable of convincing someone else to arrive at the same

judgment. In the first case, this is because (as I've suggested) the heuristic-based processes that result in his judgment are entirely unconscious, so he lacks any awareness of the warrant for his judgment. We might say that in this first case his judgment issues from *blind intuition*.

In the second case, the student is similarly unable to fully articulate why he arrives at his judgment, though for a different reason. Here he seems simply to *grasp*, on the basis of what he perceives in the diagram, that the stipulated relation between side-lengths enforces on the pair of squares a certain relation between their areas. Of course, having grasped the relationship for himself, he might be able to go on to lead someone else to arrive at the very same recognition. For instance, he could employ Socrates' own method, of drawing a suitable diagram, and drawing attention to its relevant features. But this would never amount to an argument that could compel belief in those who fail to see the relationship for themselves.⁶ The situation is reminiscent of Sibley's comments about art criticism. A critic can never *prove* to an audience that the artwork has the aesthetic features the critic claims it does; all that can be done is to employ various indirect means to prompt or encourage the audience to see what the critic has seen aesthetically in the work (Sibley 1959, 439-45). The recognition of geometrical truth in the diagram, like the recognition of beauty in a painting, remains within the sphere of private experience.


⁶ Cf. Zwicky's comment on the experience of teaching the full demonstration to her undergraduates: "Yes, there are some who grasp the *Meno* proof with a gasp, but there are others who don't see it the first, or even the second, time. If they don't get it, there's little I can do but say the same thing—walk through the demonstration...—again." (2006, 8)

While the second judgment is therefore also *intuitive* in the specific sense of having an inherently private warrant, it does not issue from *blind* intuition in the same way that the first judgment does. Again, reflecting on our own experience, of arriving at the geometrical judgment on the basis of our visual experience of Figure 1.1, it would be clearly false to say that we lack any awareness of the judgment's warrant. For not only do we *judge* that the large square is four times the area of the original one, we can quite clearly *see why* it has to be so, by our lights—it's just that we are unable to *articulate* the reason. In this case, the judgment issues not from *blind* intuition, but from what seems (to us) to be a kind of *presentational* awareness of the very geometrical situation that constitutes the subject matter of the judgment. That is, we seem to have an awareness *of* the geometrical situation, which provides an immediate ground for making propositional judgments about it, in something like the way that our visual awareness of our surroundings provides an immediate ground for making propositional judgments about what is going on around us. There is a phenomenology of the geometrical situation being *there*, presented to us, such that we can form judgments about it.

To return to the demonstration, the student has now realized that in order to construct a square twice the area of a given one, one cannot simply construct a square on a base twice as long. The initial problem remains unsolved. The discussion now progresses through a second iteration of the procedure we have just witnessed: Meno's slave proposes another answer, also plausibly heuristic-based, and is again shown a diagrammatic construction that refutes his answer (83e). He is now

thoroughly puzzled about what the solution could be (84a). Socrates now takes a diagram just like the one in Figure 1.1, and draws in diagonals of the smaller squares to connect the midpoints of the adjacent sides of the large square, yielding the diagram shown in Figure 1.2 (85a).

Figure 1.2: Step two of Plato's proof



By posing a series of questions, directed at this diagram, Socrates leads the student to recognize that the four diagonals can be seen as the sides of a new square, which is obliquely oriented. Since the triangles created by drawing in the diagonals are all equal, the original square is made up of two triangles, and the new square is made up of four triangles, the new square has to be twice the area of the original square. The student now recognizes for the first time that the square twice the area of a given one is the one constructed on the diagonal of the given square (85a-b). The demonstration is complete.

What has it shown? Socrates, at least on the face of it, takes the demonstration to show that Meno's slave has learned this truth of geometry by a process of

‘recollection’. That is, the true opinions were already *in* him prenatally, in latent form, and needed only to be “stirred up” by presenting him with suitable diagrams and posing suitable questions about them. After all, as Meno confirms, the slave has never been educated in geometry, and has answered only with his own opinions, at no point having been given the answer by Socrates himself (85b-e). Of course, given the leading nature of many of Socrates’ questions, one might wonder whether the slave has truly been giving his own opinions throughout, as opposed to simply telling Socrates the answers he thinks he wants to hear. As we discussed in Section 2.1, however, this worry misses the point. Moreover, it does so in essentially the same way in which we would miss the point in worrying that the ‘squares’ in the empirical, drawn diagram do not in fact possess the precise metric properties of the perfect, geometrical squares they are intended to stand for. It is not important that Meno’s slave has *in fact* recognized the truth of the proposition for himself. What is important is that Socrates has shown how he *could* have done so. We, the readers of the dialogue, grasp this possibility directly, in virtue of following through the demonstration for ourselves. It is our own phenomenology that shows us that we do, or at least can, *recognize* the truth for ourselves, as opposed to merely acquiring *information*, by means of measurement or any other form of testimony. Even if we resist Socrates’ story about recollection of prenatal knowledge—and he himself expresses some doubt on the matter (86b)—the demonstration should convince us that we have strong phenomenological grounds for believing in the possibility of recognizing geometrical truth ‘from within’.

I now want to suggest that Plato's picture of geometrical insight, which we have just considered, allows us to see the way towards a resolution of the puzzle we encountered in Section 2.5: Namely, how are we to make sense of geometrical insight from a psychological perspective, given that it is held to be at once presentational as well as *a priori*? After all, the clearest example we have of presentationally grounded knowledge is that of ordinary sensory experience, which is clearly not *a priori*. And the clearest example we have of *a priori* knowledge is that of ordinary discursive reasoning, which seems to lack any 'presentational' character of the sort at issue. So how can we make sense of a mental act of insight that combines these apparently opposing characteristics? I think considering Plato's view of geometrical insight as recollection suggests an appealing answer, which I now want to explain. The first task is to unpack Plato's view of *a priori* justification.

4.3 *Plato's notion of a priori justification*

To ask about Plato's view on the *a priori* is somewhat anachronistic, since Plato's writings predate the distinction between the *a priori* and the *a posteriori* as we now understand it. Nonetheless, I think that Plato does provide a clear characterization of a form of justification that is readily recognizable as corresponding to what we would now generally call '*a priori*'. Moreover, Plato's own characterization of *a priori* justification turns out to be rather more perspicuous than the common contemporary gloss as justification 'independent of experience'. In brief, I think Plato's notion of '*a*

priori justification’ should be understood as *justification independent of testimony*. I will now try to spell out more clearly what I take Plato’s view here to be.

Plato’s main concern in *Meno*, as I read the dialogue, is to establish that there can be knowledge that does not rely on testimony. But what exactly is ‘testimony’? In the first place, Plato is concerned with testimony in the most familiar, and literal sense: that of second-hand reports from others supposedly ‘in the know’. This is why, as we noted in Section 4.1, Plato portrays the character of Meno, who serves as a foil for Socrates, as someone who apparently believes that *all* knowledge (at least of eternal matters) arises from testimony. Meno, not seeing the possibility of discovering the truth for oneself, simply wants to be told the answer. The point of the geometrical demonstration is to show that some eternal truths can be grasped in a way that does not rely on the testimony of others already presumed to have knowledge.

The dialogue is replete with textual examples that reflect the dim view taken by Plato of the epistemic practice of believing testimony in this familiar, literal sense. For instance, in an interlude that occurs after the geometrical demonstration, Anytus enters, expressing a highly critical opinion of the teaching methods of the Sophists. When Socrates inquires about what justification Anytus has for taking this view, Anytus admits that he has never met any of the Sophists—his opinions have apparently been acquired secondhand. Socrates responds with a pointed rhetorical question: “How then... can you know whether there is any good in their instruction or not, if you are altogether without experience of it?” (92b-c). Similarly, later in the dialogue, Socrates draws a pointed contrast between the epistemic positions of a

“man who knew the way to Larissa” and one who merely “had a correct opinion as to which was the way but had not gone there nor indeed had knowledge of it” (97a-b). The clear worry about testimony raised by these examples is that one has to *trust* the information channel—in this case, a person assumed to be reliable—to deliver an accurate correct report. The process that produces that report as output remains opaque—a ‘black box’, from the epistemic vantage point of the recipient of secondhand information.

In cases such as these, when Plato is considering testimony in the most literal sense—that is, testimony in the form of secondhand verbal reports—the contrast that is drawn is to firsthand *perceptual* knowledge. Nonetheless, it would be a mistake to interpret Plato as claiming that ordinary perceptual judgments are free from the sort of blind trust characteristic of testimony in general. For Plato is of course highly distrustful of the deliverances of the senses; in other texts, he places strong emphasis on the tendency of the senses to mislead, especially in relation to visual illusions.⁷

Moreover, it is clear that recollection, Plato’s ultimate contrast to justification by testimony, is not intended by Plato to depend upon the reliability of visual perception in presenting the subject with an empirically accurate representation of the physical diagram itself. This is perhaps most evident in *Phaedo*, where Plato’s examples of recollection make it clear that while the physical object may *prompt* or *trigger* the occurrence recollection, this does not involve taking one’s perception of the physical object itself as evidence. Rather, the physical object serves a ‘reminding’

⁷ E.g., 602c-d in *Republic* (Plato 1992).

role, bringing to mind content that is already internally available. For instance, Socrates says the following in *Phaedo* by way of characterizing recollection:

Well, you know what happens to lovers: whenever they see a lyre, a garment, or anything else that their beloved is accustomed to use, they know the lyre, and the image of the boy to whom it belongs comes into their mind. This is recollection, just as someone, on seeing Simmias, often recollects Cebes, and there are thousands of other such occurrences. (73d)

Closer to our concerns regarding geometry, Plato also mentions the example of seeing a pair of approximately equal sticks, and being reminded of equality itself (74a-75c). The upshot is that while ordinary sensory perception plays a role in recollection, on Plato's view, it is in no way an *evidential* role. While empirical perception might serve to remind one of things with which one is already familiar, it does not carry any justificational weight. So recollection, in these cases, does not rely on what we often aptly describe as the 'testimony of the senses'. The same lesson, presumably, applies to the diagram in the *Meno* demonstration: Its role consists in merely bringing to mind an *internal* presentational awareness of the relevant geometrical situation, which itself grounds justification for the geometrical judgment.

In addition to testimony by secondhand verbal report and the testimony of empirical perception through the senses, it is clear that Plato is also concerned with *internal* forms of testimony, which he similarly intends to contrast with the sort of

genuinely testimony-free justification provided by recollection, in his view. As we saw when we considered the *Meno* demonstration above, the student's initial, incorrect answer seemed to express an intuitive judgment that issued from an unconscious, and plausibly heuristic-driven, process. Part of Plato's point in taking us through the demonstration in the way he does is to emphasize the contrast between such 'blind intuition' and the sort of presentational awareness we enjoy through recollection. The former is clearly a case of reliance on testimony, for the intuitive judgment one reaches in such a case is the output produced by a process whose operations remain an opaque 'black box' from the vantage point of the subject. In another example of blind intuition, Socrates later in the dialogue ironically suggests that it is "right to call divine" those soothsayers and prophets who receive divine dispensations of information "without any understanding" (99c-d).

The distinction Plato is emphasizing here is sufficient to counter a certain skeptical complaint often raised against the very idea of insight: that what is claimed to be 'insight' is merely a case of *intuition*, and intuition is known to be highly untrustworthy. Plato might well respond that *blind* intuition, which merely delivers a judgment as output without allowing the subject to 'see into' the process that produces it, is indeed untrustworthy, because it is a variety of testimony. On the other hand, geometrical insight, construed as recollection, does not depend on testimony.

The final form of testimony that Plato considers is testimony provided by *memory*. At one point in the discussion between Socrates and Meno about virtue, Socrates pauses to raise a doubt regarding a conclusion they had previously reached

in their discussion (that virtue is a kind of knowledge). Perhaps, Socrates suggests, they were wrong in arriving at this conclusion. Meno is resistant to reassessing the previous judgment, pointing out that “it seemed to be right at the time” when they originally considered the matter. Socrates responds by saying: “We should not only think it right at the time, but also now and in the future if it is to be at all sound” (89c). I read Plato here as suggesting that a mere *memory* to the effect that one has previously arrived at a confident judgment about some matter is itself a form of testimony, to be contrasted with the kind of justification provided by recollection. In order to enjoy this latter sort of justification, it is not enough to merely recall *that* one has previously judged so-and-so to be the case, for what one presumes to have been good reasons. Rather, one must presently have those reasons before one’s mind.

The general picture Plato presents of justification by testimony, then, is a multifaceted one. The category of testimony includes not only the literal sort of testimony we encounter in secondhand verbal reports, but also the ‘testimony of the senses’, the testimony of blind intuition, and the testimony of memories of what one has oneself previously judged to be the case. What unites all these cases is that the subject is presented only with the *output* of an information-delivering process; this process itself remains a ‘black box’, opaque to the subject. As such, in relying on testimony, of any form, the subject has to *trust* in the reliability of the process that delivers the relevant information. The promise of recollection is to provide justification that is wholly independent of reliance on testimony, in any of these ways.

I submit that this condition, of being independent of reliance on testimony, captures the sense in which Plato regards recollection as providing *a priori* justification.

4.4 *Recollection as recognition*

The previous subsection characterized recollection negatively, in terms of its independence of reliance on testimony. Now, I want to offer a positive proposal, one that I claim makes the best overall sense of the ideas Plato advances about recollection. On the face of it, Plato has presented us with a puzzle. On the one hand, *memory* about what one has previously judged or seen to be the case is a form of testimony, and hence is not *a priori* in the sense that matters to Plato. On the other hand, what is ‘recollection’ supposed to be, if not a form of memory? In addition, we are still faced with our puzzle from Section 2.5: How are we to psychologically make sense of a mental act of attaining knowledge, which is at once presentational and *a priori*? I will try to show how both of these puzzles can be resolved, by interpreting Plato’s notion of ‘recollection’ as a kind of *recognition* or *understanding*.

Based on reading *Meno* alone, one could easily read Plato as holding that geometrical insight, construed as recollection, operates on the level of *propositions*. On this interpretation, the propositional knowledge *that* the relevant geometrical proposition is true is already tacitly present within one’s soul, and requires only the right triggering occasion—the right diagram, perhaps in combination with the right questions about it—to be broad to surface of one’s conscious awareness. Indeed, Socrates explicitly says precisely this: that the true *opinions* that Meno’s slave

eventually comes to assert must have been, in some form, already inside him (85c). If we read Socrates' comment at face value, however, as the claim that recollection of geometrical truth consists in recovering latent propositional knowledge, then we are confronted with a view clearly in tension with Plato's commitment to recollection being independent of reliance on testimony. For if recollection consists in remembering propositional opinions one already possesses, it will certainly rely on the testimony of memory.

In the examples in *Phaedo*, Plato presents a different picture, one of *objectual* recollection. As we've already seen, in *Phaedo* Plato's examples of recollection involve not recollecting propositional knowledge, but merely being reminded of objects with which one is already in some way familiar: When one sees a lyre, one may be reminded of one's beloved, and upon seeing Simmias, one may be reminded of Cebes (73d). Immediately after presenting these examples, Socrates proposes a *pictorial* one: A person "seeing a picture of Simmias" may thereby "recollect Simmias himself" (73e). Then, following up on this example, Socrates then suggests that one can see two approximately equal objects (he mentions both sticks and stones as examples) and thereby recollect "Equality itself" (74a-75c).

These latter two examples suggest an appealing view of the role of the diagram in presenting subjects with the geometrical situation their judgments ultimately concerns. Just as in the pictorial case, where one recognizes Simmias *in* a depiction of him, one can also recognize *perfect* equality in a mere approximation to the latter. Presumably, this sort of objectual recognition will be similarly operative

when one is confronted with a geometrical diagram, as in the *Meno* demonstration: In that case, one would recognize a geometrical square *in* the drawn diagram. Of course, on Plato's official view, this ability itself has a rather mystical explanation, being possible only in virtue of the soul already having seen The Square Itself, while in its pre-mortal, disembodied state. But it is unlikely that we need to accept such an extravagant story in order to explain how one can recognize a geometrical square in an imperfect diagram. An alternative proposal, one that will be considered at length in Chapter 3, is that we possess a *recognitional concept* of a geometrical square, which allows us, in effect, to *see* a perfect square *in* an imperfect diagram. Since this will be a central topic of later discussion, for present purposes, I will simply assume that there is some naturalistic explanation for our ability to recognize perfect geometrical objects in imperfect, approximate diagrams, such that the geometrical objects are, in a certain sense, brought 'before one's mind'.

On this proposal, then, the sort of 'recollection' that operates when one is confronted by a geometrical diagram is, in the first place, an *objectual* form of recognition, whereby we recognize geometrical objects *in* the diagram. This proposal places us in a better position to make sense of Plato's ultimate target: 'recollection' of *propositional* truths of geometry. For we can view this propositional 'recollection' as itself a form of recognition, one that is *grounded in* the objectual form of recognition. The idea is that by 'bringing before one's mind' the relevant geometrical objects, we are epistemically in a position to *recognize that* certain propositions are true of those objects, in a way that is *a priori*. In common parlance, we often use the phrase

‘recognize that’ to denote a propositional judgment that is arrived at, not by receipt of any new information, but rather by considering what one already knows in a new light, and grasping for the first time that it fits together in a way that supports the judgment. This sense of *seeing how things fit together* seems to be an apt characterization of the sort of phenomenology we enjoy when we ‘get’ the *Meno* proof. What we experience is an appreciation for how the geometrical situation as a whole is structurally integrated. There is, moreover, a feeling of things falling into their natural place, which is somewhat suggestive of objectual recognition—so it is not altogether surprising that Plato draws a close association between the two.

4.5 *Integrative understanding as a priori*

What I now want to suggest is that this sort of integrative appreciation, which seems to serve as the justificational basis for propositional recognition, is nothing mysterious, but in fact a rather familiar mental act, which we have independent good reason to believe in, and which is (or at least can be) both presentational and *a priori*. This integrative appreciation—the *seeing how things fit together*—is generally described by epistemologists as ‘understanding’. To briefly summarize the current thinking on the topic, understanding (in the relevant integrative sense) is usually held to take as its object a coherent body of information. One cannot capture the epistemic value of understanding reductively, however, by adding up knowledge of all the atomic propositions that collectively comprise the relevant body of information, since this would in some important sense leave out the understanding itself (Elgin 2007,

Gardiner 2012, Grimm 2012). For understanding requires the integration of all these disparate pieces into a unified whole. If one has failed to grasp how all the atomic propositions hang together as a collective whole, one has failed to understand.

The following passage from Gardiner is representative of the current thinking on the nature of understanding:

When we understand an object we also grasp relations among the various parts of that object; we see its structure. When we understand an explanation, we see how different elements of the causal or explanatory web hang together and how various facts interact.... It seems that in general when we think about the nature of understanding, as distinct from true belief, knowledge or other epistemic standings, what springs to mind is coherence among beliefs and a grasping of the relations between parts. (2012, 164-5)

This picture of understanding suggests two important points. First, understanding seems to have a *presentational* character. As Gardiner says, we *see* the structure of the relevant subject matter; we *see* how its various elements hang together. In relation to some relatively abstract subject matter—say, understanding how a winner-takes-all electoral system systemically yields a *de facto* two-party political system as an emergent product—we might perhaps be inclined to dismiss this talk of ‘seeing’ as merely metaphorical. But in relation to our present concern—understanding how the relevant parts of a geometrical situation fit together, based on our visual experience of

a relevant diagram—there is little temptation to deny that our understanding is indeed presentational. Indeed, it seems to be literally *visual* in character.

The second important point is that, if we identify understanding with this integrative grasp of coherence relations, there is a simple argument to the conclusion that understanding is *a priori*. For understanding in this integrative sense is precisely what is *left to achieve*, epistemically, once one is already in possession of all the relevant pieces of *information*. Consider that the following is a clear commitment of the consensus view on the epistemology of understanding just sketched above: A subject can possess *all the empirical information* relevant to understanding how or why *P*, and yet can still fail to understand how or why *P*. Therefore, what is epistemically gained through understanding, in this sort of situation, cannot itself consist in grasping any further empirical information—by hypothesis, one already *has* all the relevant empirical information. It follows that the epistemic contribution of the understanding *itself* must be *a priori*.

Some might balk at this suggestion, given how thoroughly engaged understanding typically is with empirical content. One might think that understanding a scientific explanation of an empirical phenomenon, for instance, must surely be a thoroughly *a posteriori* matter. Here I think it is helpful to distinguish between two senses of ‘understanding’, a broad sense and a narrow one. In the broad sense of understanding, one can only be said to understand a scientific explanation (for instance) if all of the claims integrated within that explanation are themselves known to be true by the subject. Hence a scientific explanation that is internally coherent, but

which is based on mistaken hypotheses, cannot be said to be understood in the broad sense, even if the subject succeeds in grasping its internal coherence. Understanding in the narrow sense, in contrast, refers only to one's grasp of the coherence relations themselves. So in the narrow sense of understanding, one could truly be said to *understand* an explanation, even if that explanation merely integrates spurious misinformation. It is the narrow sense of understanding that I claim to be *a priori*. I claim this for the reason that the grasping of coherence relations themselves—considered independently of the truth or falsity of the pieces of information so integrated—is clearly something that makes a substantive epistemic contribution, in many cases. In particular, it is clearly capable of taking the subject from an epistemic state in which one is not justified in believing that *P*, to one in which one *is* justified in believing that *P*. (Such may well be the case when one grasps an explanation of *P*, for instance.) And just as clearly, this epistemic contribution does not depend on the receipt of any novel information.

Notice that understanding even in the narrow sense is, on its own, sufficient to justify a certain sort of propositional belief, even if the subject remains agnostic about the truth values of the pieces of information being integrated. This belief will merely have to take on a subjunctive form, being conditional on the assumption that those pieces of information are indeed true. That is, the understanding *itself* (taken in the narrow sense) will be able to justify belief in a proposition concerning what *would* be the case, were all those pieces of information to come out true. It is, in short, *subjunctive* propositions whose truth can be grasped *a priori* by means of integrative

understanding. Put differently, the epistemic contribution made by *a priori* understanding is one that is captured by a subjunctive proposition. This result, of course, fits very naturally with the proposal advanced in Section 2.4, to the effect that geometrical insight delivers justification that has a *subjunctive* form.

My suggestion, then, is that this form of integrative understanding is what underlies the recognition of the geometrical proposition at issue in the *Meno* proof. What one recognizes is roughly that *if* there were a square of the kind envisioned, and occupying space of the Euclidean sort tacitly assumed by the subject,⁸ and *if* there were a second square constructed on the diagonal of the former square, the second square would have exactly twice the area of the first one. The subject's recognition that this relationship obtains is grounded in an integrative visual understanding of the relevant spatial relationships, seen to obtain in the diagram. By this means, the subject is able to enjoy justification for a geometrical belief of a sort that is both presentational as well as *a priori*. This, I submit, provides an appealing picture of geometrical insight that is substantially true to the spirit of Plato's own proposal.

5 Conclusion and prospect

I conclude that the general proposal of geometrical insight does indeed stand up to scrutiny, and is capable of surviving the objections that have been raised against it. In Chapter 2, I will briefly confront an additional objection: that there are apparent counterexamples to the reliability of mathematical judgments reached on the basis of

⁸ The subject's tacit imposition of assumptions reflecting the structure of Euclidean space will be a central topic of Chapter 4.

visualization or visual experiences of diagrams. In Chapters 3-5, I turn to the question of the psychological underpinnings of geometrical insight. In brief, I propose what I call the *dynamic imagery account*. I argue that this account enables us to make sense of the psychological bases of basic geometrical knowledge, understood to include the foundational assumptions of Euclidean geometry, as well as the theorems proved in Book I of Euclid's *Elements*. Chapter 3 will set the stage by carefully examining the account of basic geometrical knowledge put forward by Giaquinto. Chapter 4 will provide a detailed examination of the cognitive basis of our grasp of the fundamental assumptions of Euclid's geometry. Finally, Chapter 5 will aim to shed light on our understanding of Euclidean theorems, by addressing the *generality problem*: the problem of how we are able to gain insight into general theorems of geometry on the basis of our visual understanding of diagrams depicting only special cases.

Chapter 2: Reliability

1 Introduction

In the last chapter, I presented an initial defense of the reality of geometrical insight, in which I argued that subjects can attain geometrical knowledge through the integrative understanding of visual diagrams or through visualization. In this chapter, I will briefly address an outstanding objection to the geometrical insight proposal: That appeals to visualization and visual diagrams in mathematics are known to lead to errors in mathematical judgment. If so, this seems to cast a general doubt on the reliability of visual-based insight in mathematics, and thereby undermines the claim that geometrical insight is sufficiently reliable for the judgments that it yields to qualify as knowledge.

A number of well-known cases purport to show that visualizations and visual diagrams are inherently unreliable as guides to mathematical truth. Even the staunchest defenders of diagrams grant that there are many cases in which diagrams are positively misleading. I argue that these ‘problem cases’ have been misdiagnosed. In all such cases, the erroneous judgments are not the result of any problem inherent to visual or diagrammatic methods, but are rather due to an uncritical reliance on a specific set of cognitive heuristics that operate at an unconscious level. In fact, in many of the cases thought to be most damaging to diagrams’ claim to epistemic reliability, the heuristic-based errors prove to be *correctable* by means of the appropriate use of diagram-based visual understanding. There is, I conclude, no

evidence at all to suggest that diagrams are inherently unreliable as guides to mathematical truth.

There is, of course, a long tradition of using diagrams to justify belief in propositions of mathematics. Manders has speculated that in the pre-Euclidean practice of the early Greeks, “‘seeing in the diagram’ must have been the primary form of geometrical thought and reasoning” (2008, 81). For millennia following, the more systematic (though still diagram-based) Euclidean-style practice served as the very paradigm of rigorous mathematical reasoning. Around the turn of the 20th century, amidst an emerging crisis in the foundations of mathematics, the use of diagrams came under attack from various prominent figures, including Pasch, Hilbert, and Russell (see Mancosu 2005 for an overview). One of the major criticisms leveled against diagrams concerned their putative unreliability as a guide to mathematical truth. This criticism was motivated by a number of well-known cases in which mathematical claims that had been previously accepted, apparently on the basis of ‘visual intuition’, were discovered to be false.

In the past two decades, there has been a resurgence of interest in the view that mathematical diagrams may have genuine justificatory value after all. Several philosophers have argued that *in some cases*, even standalone diagrams are sufficient to justify mathematical belief (Brown 1999, Giaquinto 2007, Azzouni 2013). Others have focused on the justificatory role of diagrams embedded *within* mathematical systems (such as that of Euclid) and have argued that claims of the unreliability of diagrams *in these contexts* have been significantly overstated (Manders 1995, Shabel

2003, Macbeth 2010). Nevertheless, even the staunchest defenders of diagrams have granted that *in many cases*, visual diagrams are indeed positively misleading, giving rise to compelling but illusory intuitions, which are at odds with the mathematical facts (Manders 1995, Brown 1999, Giaquinto 2007).

This concession might well be thought to give much ground to the critics of diagrams. For to the extent that diagrams (and the ‘visual intuitions’ they yield) play a similar cognitive and phenomenological role in generating the beliefs, in the ‘misleading’ cases as in the ‘safe’ ones, it becomes difficult to maintain that diagrams are *ever* truly reliable as guides to mathematical truth. As Burgess has recently pressed the challenge:

For all one ever has to go on, if one appeals to intuition, is one’s *apparent* intuitions at the time. If “apparent” intuitions are not all “real” intuitions, then one is going to need something *other* than intuition... to sort out cases and distinguish which apparent intuitions are real ones.” (2015, 30)

The challenge is clear: If the diagram-based means that lead us to true beliefs in some cases cannot be reliably distinguished from the ones that lead us to false belief in others, then we have every reason to regard the diagrammatic route *in general* as inherently unreliable as a guide to mathematical truth.

In this chapter, I attempt to meet this challenge, by arguing that a clear distinction can in fact be drawn between the ‘reliable’ and ‘unreliable’ cases. Moreover (and crucially), I claim that this difference is one that is

phenomenologically manifest. Both kinds of cases involve ‘intuitive’ judgments, in the sense that one may not be in a position to provide a complete *articulation* of one’s reasons for so judging. But in the ‘reliable’ cases, one is nevertheless intuitively *acquainted* with those reasons, which are evident in one’s visual understanding of the mathematical situation the diagram depicts. The ‘unreliable’ cases, in contrast, are characterized by ‘blind intuition’, in which judgments arise in consciousness as a result of cognitive *heuristics* that operate at an unconscious level, according to principles opaque to the subject. If that is the correct analysis, then in order for one to use visual diagrams as a *reliable* guide to mathematical truth, it is sufficient to judiciously restrict oneself (as Descartes would have put it) to what can be seen ‘clearly and distinctly’—while discounting any judgments that have the mere *feeling* of plausibility, whose justificatory basis is not only *inarticulate* but also *opaque*.

A positive account, showing how visual diagrams can justify mathematical beliefs, lies well beyond the scope of this chapter, and I make no attempt to provide one here, aside from the most cursory indication. My focus is instead on diagnosing the ‘problem cases’: those that have been taken to count against the justificatory use of diagrams in mathematics. The burden of my argument is to show that these cases have been widely *misdiagnosed*—that the fault lies not with any inherent frailties of visual diagrams or ‘visual intuition’, but rather with unconscious heuristics of the sort widely studied in cognitive psychology in connection with cognitive biases. In particular, I identify three specific (plausible) heuristics, considering them in relation to a range of the problem cases, and showing how they would predictably generate

just the erroneous judgments observed. I also point out some points of contrast to the ‘good’ cases—in which diagrams are used reliably—and I show that *in all such cases*, the heuristic-based errors seem in fact to be *correctable* by recourse to visual understanding of just the sort diagrams provide.

Two clarifications are in order. First, while I claim that diagrams, properly used, can serve as a *reliable* guide to mathematical truth, I do not claim that this use of diagrams is mathematically *rigorous*. That is because rigor, in the sense relevant to mathematical practice, requires a degree of formality, and explicitness about assumptions, that is fundamentally opposed to the intuitive character of diagram-based justification. In my view, the respective virtues of *rigorous* justification and *intuitive* justification are both important in mathematical practice, with neither displacing the other. Second, I do not deny that there are real *limits* to the scope of diagrammatic methods—indeed, much of mathematics appears simply inaccessible to visual understanding. What I do deny is that a sufficiently judicious use of visual understanding is apt to get us into trouble, by straying *beyond* its proper limits.

2 An idea from Plato

Recall from our discussion of Plato’s *Meno* in Chapter 1 that there seemed to be a special phenomenology that characterized geometrical insight experienced in relation to the relevant geometrical diagram. In particular, we seemed to grasp how the depicted geometrical arrangement necessarily ‘snaps together’, structurally, in a way

that guarantees the result. It is this kind of phenomenology that I take to be generally characteristic of the *reliable* use of visual diagrams as a guide to mathematical truth.

In contrast, recall the initial, incorrect answer Meno's slave gives to Socrates' question, prior to being shown the complete diagram. If we consider a square twice the area of a given square, how long will *its* side be, in relation to the length of the side of our original square? The response was as follows: "Obviously, Socrates, it will be twice the length." There is indeed superficial intuitive plausibility to this answer: *twice the length, twice the area*, we might suppose. Both the swiftness of this intuitive judgment and the apparent lack of any consciously entertained *reason* for so judging led us to speculate that this judgment issued from a *heuristic* of the sort considered within the 'heuristics and biases' research program in cognitive psychology.⁹ These heuristics operate unconsciously, according to 'rough and ready' principles, to deliver rapid judgments that are accurate enough, often enough, to promote human survival. They are, at the same time, notorious sources of persistent errors in human judgment. In this case, the relevant heuristic would seem to be that of 'attribute substitution', in which a readily accessible attribute (length) is substituted for a less accessible one (area).

I think Plato is sensitive to our susceptibility to such 'blind' intuitions, and as we saw in the previous chapter, one of the central themes of *Meno* is the importance of *recognizing* (or 'recollecting') the truth for oneself, rather than relying on the 'testimony' of opaque processes that yield judgments in the absence of

⁹ Again, Tversky and Kahneman (1974) is the classic source; see Gilovich et al. (2002) for a contemporary review.


understanding.¹⁰ For Plato, it is precisely this recognizing-for-oneself that the geometrical demonstration with the diagram is meant to illustrate. It is by facilitating appropriate visual understanding of the diagram that Socrates is able to correct the slave's initial heuristic-driven error. In what follows, I attempt to show that this point generalizes to the range of cases widely thought to undermine the epistemic reliability of mathematical diagrams: the errors in these cases are in fact traceable to *heuristics*, and are *correctable* by diagram-based visual understanding of the Platonic kind.

3 Attribute substitution

We have just considered a case in which an erroneous geometrical judgment is plausibly the result of the *attribute substitution* heuristic. Might this heuristic explain some of the errors in judgment that are often attributed to the putative 'unreliability' of diagrams? Figure 2.1 shows an example from the mathematics educator Fischbein.

¹⁰ In this connection, consider again Socrates' ironic remark in the dialogue, that "it is right to call divine" those "soothsayers and prophets" who "without any understanding" relay supposed divine dispensations of knowledge (99c-d). The phenomenological similarity to heuristic-based judgments is rather striking.

Figure 2.1: Fischbein's example¹¹



It is stipulated that the two horizontal 'strips' have the same width, constant throughout. What is the relation between the areas of the strips? While the tempting intuitive response is that the top strip is 'larger', in fact the areas are identical. The erroneous judgment, which might be taken as evidence of the inherent unreliability of diagrams, is handily explained as a result of attribute substitution: Since the relation of areas is not immediately apparent, but the top strip is clearly *longer*, the heuristic would predictably deliver the judgment that the top strip is also *larger*. But appropriate visual *understanding* of the way the figure 'fits together' can show why this heuristic response is wrong: Notice first that the regions ABFE and DCGH must be equal in area, because the strips have the same constant width. Now 'subtract' the common region DCFE from both, and one grasps immediately that the strips

¹¹ Redrawn from Fischbein (1987, 116).

themselves must be equal. The error here is plausibly attributed, not to the use of a *diagram*, but to reliance on a *heuristic*, and is in fact readily *corrected* by appropriate visual understanding of the diagram.

4 Approximation to perfection


If we visually perceive that elements of a diagram *appear* (at least approximately) to exhibit some ‘perfect’ geometrical property—such as perfect straightness, parallelism, equality, and so forth—we are apt to jump to the conclusion that this property *does* in fact hold perfectly of the geometrical situation depicted. This is a simplifying assumption of just the sort characteristic of heuristics in general. Call this, then, the ‘approximation heuristic’. This heuristic is plausibly responsible for the error in what is perhaps the most famous fallacious geometrical proof: the isosceles triangle fallacy, which ‘proves’ that all triangles are isosceles. As is well known, the spurious plausibility of the ‘proof’ rests on the fact that the exact metric stipulations of the construction (that a given line be drawn to *bisect* an angle, that additional lines be drawn *perpendicular* to the triangles’ sides) are incompatible with the topological properties of the figure as it is (incorrectly) drawn. Nonetheless, the figure *appears* (at least approximately) to satisfy these stipulations, and this appearance is accepted by at face value by the naïve subject.

As Manders observes in his (2008) study of diagrams in Euclid, such ‘exact’ properties (those disrupted by even minimal distortions to the figure) are *never* attributed based on face-value appearance, in proper Euclidean practice. Rather, they

must be shown to follow from prior assumptions, or from the more stable topological or structural (in his terms, ‘co-exact’) properties, which *can* be reliably discerned based on appearance. While Manders emphasizes the importance of implicit norms in restricting attributions of ‘exact’ properties, it is anyway obvious that such appeals to the approximation heuristic would never satisfy the Cartesian test of ‘clearness and distinctness’. For it is phenomenologically apparent that we are unable to visually distinguish *perfect* straightness from *near*-straightness, etc. Indeed, there are well-known visual illusions corresponding to just about every ‘exact’ geometrical property: position (Poggendorff), orientation (Zöllner), length (Müller-Lyer), straightness (Herring), size (Ebbinghaus), parallelism (the café wall illusion), and so forth. In contrast, there appear to be *no* visual illusions corresponding to Manders’ ‘co-exact’ properties, such as containment, adjacency, etc.


The approximation heuristic plausibly accounts for such putative cases of ‘misleading diagrams’ as the numerous variations on the ‘missing piece’ puzzle (Figure 2.2). This fallacious ‘proof’-by-rearrangement ‘works’ because the ‘hypotenuses’ of the whole ‘triangles’ *appear* (approximately) straight, even while sufficient visual probing would reveal that they are not.

Figure 2.2: Missing piece puzzle



Contrast this case with the famous visual proof (also by rearrangement) of the Pythagorean Theorem (Figure 2.3).

Figure 2.3: Pythagorean proof by rearrangement



If one merely accepts, based on appearance, that the interior oblique quadrilateral on the right is a true *square* (with perfect right angles), one is not using the diagram in a reliable way. But sufficient reflection on the structural symmetries at play does, I think, establish that this *must* be the case, with the same degree of certainty that we observed in the similar *Meno* demonstration.¹²

5 Failures of imagination

Our final heuristic is a version of what Dennett has dubbed “Philosophers’ Syndrome: mistaking a failure of imagination for an insight into necessity” (1991, 401). The relevant version for us involves a judgment that something must be *impossible* based on one’s failure to imagine the possibility. This heuristic is a plausible source of those errors in judgment widely thought to be most threatening to diagrams’ claim to epistemic reliability.

In his famous essay “The Crisis in Intuition”, Hahn (1933) draws attention to a number of famous examples of ‘pathological’ mathematical objects, drawn mostly from analysis, and identifies these as cases in which ‘visual intuition’ is shown to be inherently unreliable. Objects of this kind include everywhere-continuous-nowhere-differentiable functions, space-filling curves, ‘sponges’ with zero volume and infinite surface area, and so forth. Prior to their discovery, such objects would very likely have been dismissed as impossible, with the fault apparently lying with the


¹² I leave it to the reader to consider how this proof works—it takes considerable imaginative probing before one can fully ‘get’ the proof, and it is therefore often dismissed much too quickly as being merely suggestive, rather than (as I take it in fact to be) genuinely demonstrative.

unjustified appeal to the visual. In virtue of their similarity, it will suffice to consider just one of these cases: functions that are everywhere continuous, but nowhere differentiable (i.e., admitting of tangents at no point).

The reason why this case does not threaten the claim to reliability of the visual is this: One simply does not visually *grasp* the impossibility of such a function in a clear way. In contrast, Euclid *does* use a diagram to clearly establish, e.g., the impossibility of a triangle having two equal angles without being isosceles. He does so by performing ancillary constructions on the triangular figure, which serve to demonstrate (by appeal to *topological* properties) that the assumption that the relevant sides are *unequal* is contradictory. Nothing like that is true in the case at hand. Rather, here one *fails* to imagine the possibility, and jumps to the conclusion that the function is *impossible*, via a heuristic.

Moreover, one *can* appeal to visual diagrams to grasp the possibility of constructing such a ‘pathological’ function, through a succession of iterations. Indeed, that was precisely the motivation underlying von Koch’s discovery of his famous ‘snowflake’, a curve which is everywhere continuous but has no tangents at any point. The iterative process for constructing this curve is shown in Figure 2.4.

Figure 2.4: The von Koch snowflake



In von Koch's own words:

I have asked myself... whether one could find a curve without tangent for which the geometrical appearance is in agreement with the fact in question.

The curve which I found and which is the subject of this paper is defined by a geometrical construction, sufficiently simple, I believe, that anyone should be able to see through "naïve intuition" the impossibility of a determinate tangent. (1906, 146)¹³

This point generalizes, I believe, to *all* of the cases of 'pathological' functions mentioned by Hahn: Even though the completed object cannot *itself* be visually imagined, appropriate visual understanding of an appropriate iterative sequence of diagrams is sufficient to allow to subject to 'see' that such an object must be possible.

¹³ As quoted in Mancosu (2005, 17).

6 Conclusion

Through examining a range of cases, of the sort often thought to count against the claim of visual diagrams as reliable guides to mathematical truth, I have argued that the erroneous judgments in these cases are the result, not of any inherent unreliability on the part of diagrams themselves, but rather of an unjustified reliance on various cognitive heuristics. By restricting intuitive judgments to those that are properly grounded in visual understanding, these errors can be avoided. Moreover, they seem in fact to be *correctable* by recourse to the appropriate employment of diagram-based visual understanding.

Chapter 3: Towards a Theory

1 Introduction

In Chapter 1, we considered the way in which Socrates, in *Meno*, leads an uneducated slave to the spontaneous recognition of a simple truth of plane geometry, by visually presenting him with a drawn figure depicting the relevant geometric shapes. The discussion there seems to support the claim that genuinely *novel* mathematical knowledge is able to emerge out of the interplay between, on the one hand, constraints that are brought to bear on the perceptual understanding of diagrammatic representations, and on the other, the form and structure belonging to the figure itself. Some of the key components of this view appear to be borne out by recent research in cognitive science, which provides support for the idea of a universal geometrical competence among humans, as well as for a central role played by pictorial representations in the development of basic knowledge of geometry. For instance, there is evidence from cross-cultural studies to suggest that, even in the absence of formal training, humans universally develop concepts for the basic objects of Euclidean geometry (such as points, lines, and angles), and spontaneously converge on an intuitive acceptance of some of the basic principles embodied in the postulates of Euclid's *Elements*, such as the existence of a unique parallel line that can be drawn through a given point placed off of a given straight line (Dehaene et al. 2006, Izard et al. 2011). There is also evidence to suggest that the development of such 'natural geometry' in childhood depends upon the integration of innate contributions from

multiple ‘core knowledge’ systems, and is thereby something of a constructive achievement, rather than merely an innate endowment (Shusterman et al. 2008, Izard and Spelke, 2009, Spelke and Lee 2012, Dillon et al. 2013). Moreover, among the most obvious vehicles capable of serving effectively as *sites* of the “productive combination” of the requisite forms of core knowledge are a set of “widespread but culturally variable cognitive devices, such as pictures, scale models, and maps” (Spelke et al., 2010).

The task that will be taken up in this and the next chapter is to propose an account of the visual understanding of geometrical diagrams in terms of what I will call ‘dynamic visual imagery’. The core contention of this *dynamic imagery account* is that mathematical knowledge can arise by perceiving pictorial representations ‘dynamically’, that is, by employing in a concerted fashion two forms of dynamic visual imagery: dynamic aspectual imagery, which enables the subject to perceptually grasp the mappings among alternative perceptual integrations of a common pictorial surface form, and dynamic transformational imagery, which enables the subject to apprehend the perceived figure against a backdrop of imaginatively rehearsed spatial transformations of various kinds, including reflections, rotations, and translations. I will argue that these basic cognitive tools can serve to generate intuitive knowledge of the most fundamental objects and principles of Euclidean plane geometry, when they are applied to the pictorial understanding of simple line drawings of appropriate character. This will require considering basic features of visual perception alongside the basic definitions and postulates of the system set forth by Euclid in *Elements*.

The task will be to explain how concepts and assumptions of the sort identified by Euclid as fundamental to the system of plane geometry can arise from even more fundamental assumptions that are rooted in the structure of visual perception itself. Only by taking this approach will it be possible to provide a substantive explanation of our intuitive conviction that Euclid's postulates are *true*—for these basic propositions evidently *seem* to be true in a way that, while rather strikingly ineffable, nonetheless has a phenomenology of being in some sense 'presented' to us by our visual experience. That is, we are not at all *neutral* in our intuitive attitudes to these basic propositions, regarding them merely as 'axiomatic' starting points that we opt to take on in order to explore their implications; we rather seem just to *see* their truth, in a way that is so visually obvious as to defy articulate justification independent of ostensive reference to the images themselves. The goal, then, is to explain how this intuitive knowledge arises; the main thesis to be defended is that it does so in a way that essentially involves the use of dynamic imagery in understanding pictorial representations. On the view that will emerge, global spatial structure is dynamically 'unfolded' out of the local structures at play in the attended portion of the picture plane; it is in a real sense the pictorial forms themselves, as perceived, which serve to enforce a recognizably Euclidean character on the space they 'set up' around themselves.


2 Giaquinto's project

In this chapter, we will be concerned with Giaquinto's groundbreaking account of basic geometrical knowledge, set out in several articles (1998, 2005) and Chapters 2 and 3 of his (2007) book. Giaquinto's account serves as a useful starting point for several reasons. First, his account is distinctive among philosophical treatments of human knowledge of geometry in being directly informed by results from the cognitive science of vision. Second, his account aims to capture the epistemology of precisely the sort of 'obvious' geometrical propositions with which we are here concerned. In particular, Giaquinto takes as his case study a proposition that appeared as part of the line of reasoning in the demonstration in *Meno*, and was accepted as true by Meno's slave immediately and without question upon perception of the drawn figure: that the diagonal of a square cuts the square into two equal (congruent) triangles. Third, the core features of Giaquinto's account are broadly compatible with the key ideas of the dynamic imagery account, and their consideration thereby enables us to begin to see how dynamic imagery can play an important role in the acquisition of quite fundamental geometrical knowledge. In particular, Giaquinto's emphasis on the centrality of orientation to the perception of geometric shape, and his associated postulation of a *reference system*, a pair of orthogonal axes of orientation that are standardly applied by the visual system to shapes on the picture plane, will also end up (with refinements) being critical components of the account to be developed throughout this chapter. Finally, some of limitations of the account

Giaquinto provides will help to set the agenda for the more comprehensive account of basic geometrical knowledge to be developed in the next chapter.

By “basic geometrical knowledge” Giaquinto means to refer to knowledge of geometry that is acquired by the subject in a way that is neither by testimony nor by means of inference from prior knowledge the subject possesses (2007, 35). The target piece of basic geometrical knowledge around which Giaquinto develops his account is indicated in Figure 3.1: It is that a diagonal line drawn between opposite corners of a square divides the square into two equal parts that are perfectly congruent (that is, they are the same in both shape and size).

Figure 3.1: Square with diagonal



On the face of it, this does appear to be something that we can know immediately and non-inferentially, simply by perceiving the structure of the drawn figure. Giaquinto’s core contention is that knowledge of this proposition can indeed be acquired from the figure itself (either by perceiving a real picture or by entertaining a visual image), by bringing appropriate concepts to bear on its perceptual understanding. He is


concerned to demonstrate that the knowledge thereby acquired has at least the following three features. First, it should be suitably general, applying in scope not merely to a particular square, but to the entire class of squares. Second, it should be non-empirical, that is, it should not depend on the use of sensory experience as *evidence* for the belief acquired. Third, it should nonetheless depend essentially on the use of sensory experience, though of course this will need to be in some non-evidentiary capacity. Satisfying this third requirement is in many ways the lynchpin of Giaquinto's account, since the epistemic role visual experience is generally regarded as inherently evidentiary in character, which raises the question of how any non-evidentiary contribution of sensory experience to knowledge is even possible, let alone knowledge of a *general* class of objects. In brief, Giaquinto's answer, which depends on Peacocke's (1992) account of concept-possession in terms of belief-forming dispositions, is that if we assume the subject's possession of appropriate geometrical concepts, sensory experience may serve to trigger general belief-forming dispositions, yielding general beliefs which may qualify as knowledge, provided they meet certain further conditions. Since the resulting account thereby identifies basic geometrical knowledge as both non-empirical and essentially sensory-involving, Giaquinto notes that it may be interpreted as a defense of a version of the Kantian idea that geometrical knowledge is "synthetic *a priori*" (2007, 47).

3 Spatial structure and symmetry

Giaquinto opens the exposition of his account with a consideration of some general features of human visual perception of shapes, which will prove important for the discussion to follow in this chapter. To begin with, he takes the perception of *lines* more or less for granted (2005, 31; 2007, 13-14), understanding them as consequences of the parsing of scenes by the visual system, by means of constructing an initial representation of the surface layout of visual scene, consisting of bordered segments (see Nakayama et al. 1995). While visual perception tends to use segmentation of surfaces as a preliminary to constructing visual representations of the three-dimensional spatial structure of objects in space, Giaquinto points out that, particularly in the context of the perception of pictures of Euclidean plane figures, the surface borders can also be interpreted merely two-dimensionally, that is, as lines laying in a flat plane. We will return to consider the visual basis of geometric lines in the next chapter.

From there Giaquinto turns to a consideration of orientation and its effects on the visual perception of shapes, which he will go on to make use of in providing a visual ‘category specification’ for squares, and (on that basis) in specifying a geometrical concept for squares that is grounded in visual perception. Mach (1897) was the first to draw attention to the radical influence orientation is capable of exerting on the perceived forms of objects or figures, as illustrated by the example of the square-diamond (Figure 3.2).

Figure 3.2: Square and diamond



The two figures are congruent under a rotation of 45 degrees, but perceptually, they are apprehended as having quite different ‘shapes’, merely as a result of the difference in orientation. While Giaquinto himself does not make the point explicit, it is worth pausing to consider that it is perfectly possible for a subject to be visually familiar with the shapes of both the square and the regular diamond, and to nonetheless fail to recognize that these shapes are identical (in the sense of being congruent). In fact, the recognition that they are identical might reasonably be regarded as a very elementary example of synthetic *a priori* knowledge of geometry. This is because, at least on the face of it, the insight that this identity holds depends essentially on sensory experience, though in a capacity that would seem to be non-evidentiary, inasmuch as the epistemic force of the insight is not at all undermined by raising the prospect that one’s visual perception may not be veridical; it seems beside the point, that is, that the *appearance* of the square figure may be, for instance, a hallucination, because the recognition that the *apparent* shape is identical to a regular diamond remains sound. The case is also revealing in that it provides a very clear illustration of the way that the two postulated forms of dynamic imagery can interact

in order to deliver knowledge of geometry. Plausibly, the identity of square and regular diamond might be grasped by means of perceptually understanding Figure 3.2 in the following way: First, one engages in a 45-degree “mental rotation” (Shepard and Metzler 1971) of one or the other of the figures, employing what we have characterized as transformational imagery; second, the resulting shift in orientation yields an accompanying shift in *aspect* or interpretive perceptual integration, such that the alternative ‘takes’ of *square* and *diamond* are apprehended as mapping to the same underlying surface form, in an application of what we have called dynamic aspectual imagery.


It is not yet clear, however, what it is to perceptually grasp a perceived figure as square or diamond, or why orientation should exert such a forceful influence on the perception of shape in situations of this kind. Here Giaquinto appeals to research from the psychologists Irvin Rock and Stephen Palmer on the intimate relationship between visual orientation and the perception of reflection or ‘mirror’ symmetry (Rock 1973, 1997; Palmer 1983, 1985; Rock and Leaman 1997). To briefly summarize the key connection between orientation and symmetry this research reveals, the visual system appears to construct representations of shapes by assigning, to perceived forms, *axes of orientation*, typically vertical and/or horizontal; these orientation axes then appear to serve as axes for perceptual judgments of reflective symmetry. (In fact, the interplay between orientation and symmetry is rather more complicated, since an initial ‘first-pass’ evaluation of reflection symmetry seems to play a role in the initial assignment of orientation axes, which thereafter serve to

support subsequent, more precise judgments of symmetry; see Palmer and Hemenway 1978.) The visual system seems to have a general preference for vertical and horizontal orientations, in the sense that subjects tend to be significantly more accurate in making perceptual discriminations across a broad range of tasks—a widely studied phenomenon in psychophysics known as the “oblique effect” (Appelle 1972). For our purposes, the important judgments to consider are those made about reflection symmetry and perpendicularity, both of which have been experimentally shown to be less susceptible to error when stimuli are presented to subjects aligned with vertical or horizontal axes of orientation (e.g., Goldmeier 1972, Ferrante et al. 1997).

Giaquinto, borrowing terminology from Rock, characterizes a *reference system* as “a pair of orthogonal axes, one of which has an assigned ‘up’ direction” (2007, 15). Given his subsequent elaboration of geometric concepts in relation to reference systems, I take this to be a key posit of Giaquinto’s account, so it is worth getting clear on exactly what is meant by this term. One of the key findings of the research by Rock and Palmer is that a reference system is assigned to the visual representation of an object or figure in a way that can be informed by various factors. Retinal orientation plays a role, but tends to be outweighed by the influence of environmental orientation: Subjects viewing stimuli with their heads tilted, for instance, tend to be as accurate in their judgments concerning environmentally upright, but retinally oblique figures as subjects viewing the same stimuli with their heads oriented normally (Rock 1997). Environmental orientation, however, is itself


liable to be overridden by various framing effects, as illustrated by phenomena like that shown in Figure 3.3 (Palmer 1985).

Figure 3.3: Effect of framing on orientation



Here, an embedded figure that would appear diamond-shaped based on either a retinal or an environmental (here, page-based) orientation, is instead perceived as square-shaped, due to the ‘frame’ in which it is embedded. This example illustrates that the assignment of orientation, *via* the alignment of the perceived figure with a reference system, is in certain cases determined by the structural layout of the figure itself, a phenomenon called “intrinsic orientation” by Rock (1997); he provides the cases shown in Figure 3.4 as examples in which figures inherently possess a ‘natural’ vertical axis.

Figure 3.4: Intrinsic orientation¹⁴




Intrinsic orientation is important because it demonstrates the following key fact: While the assignment of a reference system appears to be one of the fundamental structuring assumptions the visual system brings to bear on the perceptual understanding of drawn figures, the *manner* in which the reference system is aligned with the perceived figure is marked by a receptivity to the figure's inherent structure, which in some cases allows the figure itself to effectively 'set up' its own orientation.

In light of this fact, it is slightly odd that Giaquinto seems to characterize a reference system, *as such*, as a pair of *orthogonal* axes. The reason this is odd is because orthogonal structure is not an inherent feature of all perceived forms, and in some cases, it seems to run very much against the structural grain of the figure. For instance, the six-fold symmetry of the snowflake, which is perceptually quite vivid when one is presented with the form, would be rather obscured by the rigid imposition of orthogonal axes, if these were taken to be structurally primary. There remains, however, a clear phenomenological salience to the orientation of spatial forms simultaneously along orthogonal axes, as becomes quite apparent when

¹⁴ Redrawn from Rock (1997, 139).

considering the aesthetics of simple forms. Rudolf Arnheim offers a particularly vivid illustration in his *Art and Visual Perception* (1974, 93-4), in a diagram apparently inspired by comments from Wertheimer, shown in Figure 3.5.


Figure 3.5: Sequence of triangles¹⁵



The sequence of five triangles is to be understood as the succession of salient ‘stages’ of the continuous motion of the right-side vertex in a downward direction, as the two left-side vertices remain fixed in place. The telling point is that the structural *consequences* of the continuous motion of the point are not themselves continuous, but rather are parsed into a series of discrete stages determined by the proximity of the form to one of these five stages, each with natural orientation along orthogonal axes simultaneously. As Arnheim himself points out, however, alignment along orthogonal axes is not perceptually inevitable, even in cases for which it is structurally intrinsic or ‘natural’ to the figure—as illustrated in Figure 3.6, triangles *c* and *d* can, with some degree of effort, be seen ‘against their natural grain’ as, respectively, oblique (for *c*) and noticeably deviating from a right-angled ideal (for *d*).

¹⁵ Redrawn from Arnheim (1974, 93-4).

Figure 3.6: Deviant orientation¹⁶



It therefore seems that Giaquinto's formulation overstates the centrality of orthogonal-axis structure, by apparently building it into the inherent nature of the visual reference system that gets assigned to perceived figures in general. At the same time, it seems clear that there *is* something privileged about orthogonal-axis structure, a fact that is likely bound up with the intimate relationship between orientation and reflection symmetry. For one of the findings already mentioned is that sensitivity to reflection symmetry is greatest when the symmetry is along the primary axis of orientation. If we take the primary axis of orientation (and hence, of reflection symmetry) to be the vertical axis, then, it is striking that the only straight line capable of intersecting this axis, that would also possess reflection symmetry *across* it, would be a straight line along the orthogonal (that is, horizontal) axis. So the mere fact that the primary axis of orientation is associated with reflection symmetry seems inherently to endow the orthogonal axis with salience, which may be reflected in visual processing in various ways. To an extent this speculation is borne out by recent experiments, which have used clever techniques to distinguish twofold reflection


¹⁶ Redrawn from Arnheim (1974, 94).

symmetry *simpliciter* from twofold reflection symmetry *along orthogonal axes*, and which demonstrate a significantly greater visual salience of twofold symmetries with orthogonal axes as compared with twofold symmetries with non-orthogonal axes (Treder and van der Helm 2007). As we proceed in this chapter and the next, it will turn out to be important that the human visual system is not constrained to fix orientation with respect to an orthogonal pair of axes. It will also turn out to be important, however, that the visual system has a natural preference for orthogonal-axis orientation.

We now return to the question of what it is to perceive a figure as a square *vs.* as a diamond, and why this difference in perceptual aspect would be tied to orientation. Giaquinto proposes that to recognize a figure as a particular *kind* of shape is just to perceive it as satisfying the relevant *feature description set*, an entity which he notes should be understood as playing roughly the same role as the *category patterns* posited by Kosslyn (1994) in his integrated account of visual perception and imagery—in essence, these are stored visual templates to which incoming visual information is matched in the process of object recognition; in a secondary use, they can play the role of driving the formation of visual imagery, especially in the absence of (or absence of attention to) visual inputs. Since sets of feature descriptions function to specify the category of the object being perceived, Giaquinto also describes them using the term “category specification”. The question then, thusly framed, is how the category specifications of squares and diamonds, respectively, differ in their feature descriptions. Following Palmer (1983), Giaquinto proposes that in both category

specifications, a prominent feature description will be the reflection symmetries that exist along the primary pair of orientation axes (2007, 21-2). In both cases, the primary axes are vertical and horizontal, but due to the different orientations of the two otherwise-identical figures this same pair of axes imposes different primary symmetries, as shown in Figure 3.7.

Figure 3.7: Square and diamond with main symmetry axes¹⁷



The figure on the left is perceived as having reflection symmetries along orthogonal axes that cut the opposite sides at their midpoints; the figure on the right is perceived as having reflection symmetries along orthogonal axes that bisect the opposite angles. These different sets of perceived reflection symmetries have immediate geometrical consequences that serve to distinguish the two figures, in virtue of the way they enforce mappings between the structural components internal to the figure. The symmetries perceived in the figure on the left map *opposite sides* to

¹⁷ Redrawn from Giaquinto (2007, 22).

one another, and map *adjacent angles* to one another. This effectively enforces rectangularity on the perceived figure, though these primary symmetries do not actually guarantee that the perceived figure is a *square*, as becomes evident when one imagines either the vertical or horizontal pair of sides changing in length; it is apparent that this spatial transformation would not violate symmetry along these two primary axes. Similarly, the symmetries perceived in the figure on the right map *opposite angles* to one another while also mapping *adjacent sides*. This guarantees equality of all the sides, but not equality of all the angles; one can imagine deforming the figure into a rhombus with varying angle measures without violating the primary reflection symmetries.

The empirically motivated proposal that orientation axes serve also as axes for the perception of reflection symmetry, then, turns out to provide a satisfying account of the phenomenological difference between perceiving the common square-diamond feature under the two different orientations. Intriguingly, while the figure in both cases is surely perceived as having all of its sides *and* all of its angles at least *approximately* equal, the difference in symmetry-based feature descriptions in each case ‘locks onto’ only one of these equalities, leaving the other ‘insecure’, or potentially variable. Once this difference is pointed out, it seems intuitively to get our phenomenology right. That is, the left-side figure forcefully asserts itself as rectangular, leaving no doubt as to the perpendicularity of all its angles, but is not so assertively square, that is, equal in respect of the lengths of all its sides. Conversely, the right-side figure patently has all equal sides, but the angles, which do seem at

least *roughly* perpendicular, nonetheless seem as though they could easily slip into a slight incongruence without our noticing. This is interesting insofar as the common figure, which lies at the intersection of the larger classes of rectangles and rhombuses, is perceived differently, effectively in virtue of its alignment with one or the other of these classes. From the perspective of the dynamic imagery view, however, there is a deeper significance, for the phenomenological effect of these different symmetry assignments can be seen to consist essentially in the dynamics those symmetries impose on the figure. Both figures are, we might say, *allowed* to stretch or compress along their vertical and horizontal axes, in the divergent ways that entails under their respective orientations. This exemplifies a general feature of axes of reflection symmetry that will play an important role in the account to be developed: They not only serve to structurally integrate static forms by enforcing mappings amongst their parts; they also serve to structure the *dynamics* of the space, rendering translation *along* the axis salient, in a sense negatively, precisely by enforcing restrictions on translations *across* the axis.

As is evident from the foregoing discussion, reflection symmetry about vertical and horizontal axes is not sufficient to distinguish squares from the broader class of rectangles. This implies that additional feature descriptions will need to be added to the category specification for squares. Giaquinto proposes that in order for a shape to be perceived as square rather than merely rectangular, it is necessary that it be perceived as having additional, secondary reflection symmetries about the two oblique (45-degree) axes, which will be fixed in relation to the initial assignment of

the primary (vertical and horizontal) axes. Giaquinto himself explicitly presents this postulation of a *secondary* sensitivity to symmetry about oblique axes merely as a convenient assumption (2007, 23), but in fact, the psychophysical evidence indicates that the sensitivity to reflection symmetry of the human visual system is distributed across different orientations of axes in just the manner Giaquinto proposes: Vertical symmetry is most salient, followed by horizontal, and then symmetry about the diagonal (45-degree) axes of orientation, with the latter having been found to be significantly stronger than for non-diagonal oblique orientations (Wenderoth 1994). The proposal of symmetry perception along secondary axes of lesser salience is, moreover, able to account for how the square can be perceived as possessing the same symmetries that serve as the primary feature descriptions for the diamond-shape, while still being perceived as primarily square-shaped rather than diamond-shaped.

4 A perceptual concept for squares

Giaquinto is now in a position to provide the full category specification for the visual recognition of squares, which contains the following feature descriptions:

Plane surface region, enclosed by straight edges;

edges parallel to H, one above and one below;

edges parallel to V, one each side.

Symmetrical about V.

Symmetrical about H.

Symmetrical about each axis bisecting angles of V and H. (2007, 23)

“V” and “H” here, of course, refer respectively to the vertical and horizontal axes of orientation. Now, one might expect this category specification itself to be identified as a perceptual *concept* of squares, but on the framework Giaquinto borrows from Peacocke (1992), it is the disposition to form beliefs in certain contexts that is regarded as criterial of concept possession. As such, the perceptual concept for squares will be framed in terms of the disposition to form beliefs, and can be stated concisely in terms of the category specification just provided, as follows:

The concept {square} is the concept C that one possesses if and only if the following holds: When an item x is represented in one’s perceptual experience as [approximately satisfying the category specification] and one trusts the experience, one believes without reasons that that item x has C. Conversely, when one trusts one’s perceptual experience of an item x , one believes that x has C only if x is represented in the experience as [approximately satisfying the category specification]. (2007, 27)¹⁸

¹⁸ In fact, Giaquinto’s full statement of the perceptual concept {square} includes an additional clause designed to capture the fact that perceptual concepts can be applied to support inferences concerning things that are *not* being perceived. Given that our interest here is in things that *are* being perceived, we can ignore this complication for the sake of simplicity.

Note that there are actually two distinct senses ascribed to the term “C” here: The references to an item ‘having C’ presumably indicate that item’s possessing the *property* C that is picked out by the *concept* C. As Giaquinto notes (2007, 28-9), the perceptual concept for squares just characterized is a *vague* concept, since it captures the relationship between a disposition to form the belief that an item is square and the perceptual representation of that item as *approximately* (as he puts it, “nearly or completely”) satisfying the various description features that make up the category specification for squares. This is sensible, insofar as lines can be represented in perception as more or less straight or parallel, figures can be represented as more or less symmetrical about some axis, and correlatively, figures can be represented as more or less square. But of course, a *geometrical* concept for squares should not have this indeterminacy. Therefore, Giaquinto proposes that we can define such a geometrical concept of a (perfect) square based on the initial perceptual concept, by specifying that the item must be represented in one’s perceptual experience as *perfectly* satisfying the description features of the category specification. One might doubt whether human perception could ever represent an item as *perfectly* satisfying the description features of straightness, parallelism, symmetry, and hence as being perfectly square. Giaquinto’s response is to point out that, since there is “a finite limit to the acuity of perceptual experience, there are lower limits on perceptible asymmetry and perceptible deviation from (complete) straightness” (2007, 28). Accordingly, it is possible for our perceptual experience to represent certain items as

perfectly instantiating the relevant description features, and hence, as falling under the geometrical concept {perfect square}.¹⁹

Even granting this point (to which we will return in the next chapter), one might still worry that the concept {perfect square} imposes an unrealistically exacting constraint on the drawn figures that are capable of engaging the concept. In particular, these figures will have to be perceptually indistinguishable from *perfectly* square figures in order to be able to activate the belief-forming dispositions that will issue in general beliefs about geometric squares, on Giaquinto's account. If his account is subject to this stricture, it is doubtful that it will be able to provide a realistic explanation of basic geometrical knowledge as it figures, for instance, in *Meno*. For it is simply implausible that Socrates' demonstration depends on the figures being drawn *perfectly*. Giaquinto anticipates this problem, and responds that a figure that is a visibly imperfect instance of a square might still activate the belief-forming dispositions associated with the concept {perfect square} *indirectly*, via visual imagination (2007, 39). In this case, the subject would perceive the figure as (imperfectly) square (presumably using the *perceptual*, not geometrical, concept for squares), and is thereby caused to imagine a perfect square. While Giaquinto's exposition of his account assumes that the figure itself can be perceived as perfectly square, then, this can be understood simply as a convenient assumption, given that the

¹⁹ As Giaquinto explicitly notes, there are surely other concepts expressed by the word 'square' that are appropriately characterized as geometrical concepts; his claim is just that the geometrical concept for squares that is useful for explaining basic geometrical knowledge about squares is plausibly a modification of the perceptual concept, of the kind just indicated (2007, 29).

concept {perfect square} could be indirectly engaged by an imperfectly square figure in this manner.

Giaquinto's proposal that perception of a visibly imperfect square might cause visual imagery of a square that is perceptually indistinguishable from a perfect one seems plausible, insofar as it arises naturally out of the idea that the perceptual concept for squares is grounded in a "category specification" (in his own terminology) or what Kosslyn (1994) calls a "category pattern". These category patterns can be thought of as stored templates that are matched to incoming visual information in object recognition. Visual imagery, on Kosslyn's view, can arise when a category pattern is activated in the absence of visual input. As such, it is plausible that a figure that only *approximately* satisfied the activation conditions of the category pattern might issue in visual imagery of a more 'perfect' or idealized version of the figure itself, resulting from the activation of the category pattern itself.

Giaquinto's proposal still leaves some important questions unanswered, however, concerning the *content* of the perceptual experience the subject has, when perceiving a visibly imperfect figure in conjunction with visual imagery of a more idealized form. The view defended here will be that the contribution of imagery to perceptual content should be understood in relation to the *depictive content* of the figure. We will return to this point in the next chapter when we discuss the question of what it is to perceive a geometrical line.

5 Basic geometrical knowledge

Having now offered a statement of the geometrical concept {perfect square}, Giaquinto is ready to provide his account of basic geometrical knowledge. Recall that the specific instance of knowledge Giaquinto wants to explain is one of the steps that is taken to be obvious in the geometrical demonstration in *Meno*: that a straight line drawn between opposite corners of a perfect square (that is, along the diagonal) will divide the square into two parts that are perfectly congruent (see Figure 3.1). On his account, acquiring this bit of knowledge, for the *special case* of the particular drawn figure (that is, not yet for the class of squares in general), will depend on two conditions. First, the figure must be perceptually represented as satisfying one of the description features of the category specification for squares: namely, reflection symmetry about (one of) the oblique axes assigned as part of the reference system. Second, the perceiver must possess the concept of *congruence*. Here is what Giaquinto says:

Suppose one has a concept for geometrical congruence. If a figure *a* appears to one symmetrical about a line *l* and one trusts the perceptual experience, one will believe that the parts of *a* either side of *l* are congruent. (2007, 36)

Given that Giaquinto's statement of the geometrical concept {perfect square} links the perceptual judgment that an item is (perfectly) square directly to (perfect) reflection symmetry about the axis on which line *l* is drawn, this does provide a

plausible starting point for understanding how we can acquire knowledge of the relevant truth in the special case of the particular figure. This is, however, all that Giaquinto says about knowledge of the special case. One would like to hear more, in particular, about the nature of the concept for geometrical congruence, and how its belief-forming dispositions are engaged by the perception of reflection symmetry. Presumably, the statement of the concept for congruence (the relation between two figures with both the same shape and size) would not explicitly mention reflection symmetry itself, since the general property of congruence applies in many cases that lack reflection symmetry. And while one might be tempted, pre-theoretically, to think that the mere perception of reflection symmetry in some sense inherently ‘involves the idea’ of congruence, that claim itself stands in need of a more precise statement, and seems difficult to reconcile with Giaquinto’s claim that symmetry, in the sense in which it occurs as part of the category specification for squares, is intended to be *non-conceptual*, merely a feature “that the visual system can detect and represent” (2007, 33-4, note 39). In the next chapter, we will return to the problem of the perceptual basis of a geometrical concept for congruence, and its relation to symmetry.

The other feature that is missed in Giaquinto’s rather brief explanation of how we acquire knowledge of the special case has to do with a point that he makes vividly elsewhere in his book: “The creative heart of the discovery process lies in viewing a form in two ways at once” (2007, 158). The description ‘two ways at once’ is especially apt because indeed, it really is crucial for the perceiver to be able, not merely to view a figure under different aspects at different times, but further, to

represent the single figure as *viewable* in both ways. Indeed, this unified representation of the coincidence of multiple aspectual integrations is precisely the role played by what I have termed dynamic aspectual imagery. What is not fully clear is whether or not Giaquinto regards the same point as applying in the case of *basic* geometrical knowledge—whether seeing Figure 3.1 in ‘two ways at once’ is taken to be important in acquiring knowledge of congruence (in the special case), on his account. For Figure 3.1 *does* seem to have the kind of pictorial ambiguity that Giaquinto regards as playing a key epistemological role. One way of seeing this figure grasps it as a single square, with an interior line that can be thought of as placed onto the square. The other way of seeing it grasps it as a pair of adjacent triangles that share a border. Intuitively, these different ways of seeing the figure seem to parse it into different sets of ‘objects’, and correspondingly, they implicitly endow the figure with different dynamic profiles: On the first interpretation, the diagonal line (a positive ‘object’ instead of a negative border) could be placed differently onto the square; on the second interpretation, the triangles (as distinct ‘objects’) could be arranged into different configurations that are not necessarily square-shaped. On the account of the concept of congruence to be given later, being able to perceive the figure in these ‘two ways at once’ will turn out to be crucial, because the judgment of *congruence of parts* will require imaginatively rehearsing a spatial transformation that necessarily destroys the original square.

In any case, having indicated how knowledge of congruence in the special case is to be acquired, Giaquinto proceeds to consider how this knowledge could be

extended to cover the general case, that is, how we can acquire knowledge that the congruent-parts property holds for *any* square. Here is the full passage in which Giaquinto first characterizes knowledge of the general case, beginning with the two sentences quoted above, which explain knowledge of the special case:

Suppose one has a concept for geometrical congruence. If a figure *a* appears to one symmetrical about a line *l* and one trusts the perceptual experience, one will believe that the parts of *a* either side of *l* are congruent. We can further say that if *a* appears to one symmetrical about *l*, regardless of whether one trusts the experience one will believe that if *a* were as it appears (in shape), the parts of *a* either side of *l* would be congruent. With this antecedent condition, it is only the apparent shape of *a* that is relevant: having that shape, the shape that *a* appears to have, suffices for the attributed property. So one has a more general belief, about *any* figure having the apparent shape of *a*, that it has the attributed property. This is the level of generality that we require for geometrical truths. (2007, 36)

Here Giaquinto assumes that the subject who perceives figure *a* (i.e., Figure 3.1) can, as he puts it, “think of its apparent shape demonstratively, as *that* shape” (2007, 35). Of course, if an account along these lines is going to be able to explain how the subject can come to form the belief that *all* perfect squares have the congruent-parts property, the content of the demonstrative “*that* shape” will have to abstract away

from differences in size, orientation, and position. That Giaquinto understands the scope of the demonstrative “*that shape*” in this way is clear from his characterization of the following belief-forming disposition (“PS” for “perfect square”), which he imputes to the perceiver:

(PS) If you were to perceive a figure as perfectly square, you would believe of its apparent shape S that whatever has S is perfectly square, and that whatever is perfectly square has S. (2007, 36)

It is the possession of this disposition (PS), on his account, which licenses the perceiver to treat the demonstratively grounded property “having S” and the geometrical property “being a perfect square” as “cognitively equivalent” in the sense that they will automatically be taken as substitutions for each other in inferences (2007, 38). Giaquinto appears to take the disposition (PS) to follow automatically from the dispositions inherent in three different concepts he assumes the perceiver to possess: the demonstrative concept {*that shape*} (or “S”),²⁰ the geometrical concept {perfect square}, and finally, an additional concept for restricted universal quantification, or {r.u.q.}. Giaquinto doesn’t provide a complete statement of {r.u.q.}, but offers the following key fact about it:

²⁰ Giaquinto himself doesn’t explicitly state that the demonstrative expression “*that shape*” expresses a demonstrative *concept*, but this seems like a reasonable interpretation. Note that, unlike possession of the other concepts to which he appeals, possession of this one will not only be specified relative to some (potential) perceptual experience, but will itself depend upon the actual occurrence of a perceptual experience.

If one has the concept of restricted universal quantification, one will believe the proposition “Every F has G” when and only when one would find cogent any given inference of the form “ x has F, so x has G”. (2007, 35)

With these tools in place, Giaquinto is almost ready to provide his explanation of how the subject perceiving figure *a* is able to arrive at the general belief that all perfect squares have the congruent-parts property. Recall that he has already explained the forming of a belief about the special case (i.e., that the property holds for the particular square in the drawn figure) in terms of the following disposition: “If a figure *a* appears to one symmetrical about a line *l* and one trusts the perceptual experience, one will believe that the parts of *a* either side of *l* are congruent” (2007, 36). In order to state the more general belief-forming disposition, about *any* figure with the same shape, he will need to say something along these lines: One will believe that for *any* figure with the *apparent* shape of figure *a*, the parts either side of *that* figure’s ‘line *l*’ would be congruent. Since “*l*” itself refers to particular line in figure *a*, though, that formulation will not do—Giaquinto needs a way for the subject to think about the line that *corresponds* to *l* in the generic case. He proposes, therefore, that the subject will need an additional concept for perfect correspondence. He provides a truth-conditional definition of correspondence of lines:

A line k through b corresponds to line l through a if and only if some similarity mapping of a onto b maps l onto k . (2007, 37)

Giaquinto is now ready to characterize the belief-forming disposition (called “C” for “congruence”) that will enable the subject to grasp the general truth about all figures with the same shape that figure a is perceived as having:

(C) If one were to perceive a plane figure a as perfectly symmetrical about a line l , then (letting “S” name the apparent shape of a) one would believe without reasons that for any figure x having S and for any line k through x which would perfectly correspond to l through a if a were as it appears, the parts of x either side of k are perfectly congruent. (2007, 37)

Since (C) includes a conditional disposition to form a belief about “any figure x having S”, it must therefore depend on {r.u.q.} operating in conjunction with the demonstrative concept {*that* shape}. Beyond that, the possession of (C) will of course depend on the possession of concepts for perfect congruence and for perfect correspondence. As stated, (C) itself is not yet a disposition to form a belief about geometrical squares *as such*, but given that the antecedent of the conditional (that a is perceived as symmetrical about some line l) is satisfied by perceptions of figures as square (taking l as the diagonal), and given also that (PS) licenses the subject to regard “having S” and “being perfectly square” as cognitively equivalent, we can take

the subject to possess as well, the following disposition—effectively, a paraphrase of a special case of (C):

(C*) If one were to perceive a plane figure *a* as perfectly square, one would believe without reasons that for any perfect square *x* and for any diagonal *k* of *x*, the parts of *x* either side of *k* are perfectly congruent. (2007, 38)²¹

Having reached this point, Giaquinto pauses to comment on the significance of what he takes this to demonstrate:

The point here, the truly remarkable point, is that if the mind is equipped with the appropriate concepts, a visual experience of a particular figure can give rise to a general geometric belief. In short, having appropriate concepts enables one to “see the general in [the] particular”. One cannot have those concepts without having a disposition to form a general belief as a result of a certain kind of visual experience. In the example at hand the general belief is the target belief that the parts of a square each side of a diagonal are congruent. (2007, 38-9)

²¹ Strictly speaking—and as Giaquinto himself notes—possessing (C*) will further require one to have a concept for *diagonals*, so that one will be able to “think of a line through perfect square *x* which would correspond to a diagonal of *a* if *a* were perfectly square as, simply, a diagonal of *x*” (2007, 38).

Two comments are in order here. First, Giaquinto takes his account to have provided a solution to a restricted version of what was introduced at the conclusion of Chapter 1 as the *generality problem*—the problem of how to grasp truths about a general class of objects by perceiving a depiction of only a particular instance from that class. The reason Giaquinto’s solution applies only to a restricted version is because squares are all geometrically similar (in shape) to one another—a feature that enables his account to derive a general claim from {r.u.q.} in connection with {*that shape*}, and also to appeal to a restricted concept of correspondence that is defined in terms of a “similarity mapping”, or “shape-preserving transformation”. The problem is more difficult in the case of claims quantifying over all triangles, for instance, since triangles are *not* all geometrically similar to one another; the account as it stands will not be equipped to handle these more difficult cases.²² Nonetheless, even the restricted version of the problem that arises for squares is far from trivial.

The second comment concerns Giaquinto’s claim that the solution to the generality problem (in its restricted version) is entirely a matter of possessing the right concepts. While it is true that a perceptual experience must occur in order to “trigger” the belief-forming dispositions entailed by possession of those concepts, it is indeed remarkable that on his account, the concepts themselves are taken as sufficient to dispose the subject to form a fully general geometrical belief immediately upon perceiving the particular diagram. Of course, according to DIA, dynamic imagery,

²² Giaquinto does offer an account of how he thinks these difficult cases can be handled. That account will be discussed when we consider the generality problem directly in Chapter 5.

and not just concepts, plays a critical role in deriving geometrical beliefs from perceived figures. Therefore it is worth a careful look to see whether concepts are really able to accomplish what Giaquinto supposes they are. We have already identified a lacuna in Giaquinto's account of the special case of congruence-of-parts, which seemed to point to a role for imagery of the sort posited by dynamic imagery account; we will now see that the same point applies to his account of the general case as well.

6 Criticism of Giaquinto's view on generality

Consider that on the account just provided, generality (in the sense of the subject being disposed to form a belief about all members of a general class of geometrical objects) arises primarily due to the interaction of {r.u.q.} with the demonstrative concept {*that shape*}. These are the ingredients in common between the bases for the two general belief-forming dispositions (PS) and (C), which are themselves ultimately used in conjunction to determine (C*). In the case of (PS), the concept {perfect square} is also involved; in the case of (C), the concepts of *correspondence* and *congruence* are. It's not difficult to see why {r.u.q.} and {*that shape*} would be the crucial ingredients in getting to a general geometrical belief (or a disposition to form such general beliefs): Effectively, the former is just a disposition to draw generalizations from an arbitrary case, and the latter is just a way of pointing to the feature that the perceived case has in common with the intended class of objects. In particular, {r.u.q.} disposes the subject to believe that "Every F has G" whenever the

subject finds cogent some inference of the form “ x has F , so x has G ”. In the cases Giaquinto presents, “ F ” would be the property of having “*that shape*” (demonstratively fixed in relation to a perceptual representation of figure ***a***) and “ G ” would be either the property of being perfectly square (in order to determine the first disposition in (PS), that anything with *that shape* is a perfect square), or the congruent-parts property (in order to determine (C), the disposition to believe that anything with *that shape* has the congruent-parts property).

What is unclear, however, is that the perception of figure ***a*** as having either of these properties “ G ” should be counted as a case in which the subject finds the inference “figure ***a*** has *that shape*, so figure ***a*** has G ” cogent. The “so” here, of course, is crucial: The universal generalization to “Every F has G ” is only logically valid provided that an inference has been made from “ x has F ” to “ x has G ” *without any other* assumptions having been made about x . Has the perceiver, in the cases Giaquinto describes, really made such an inference to “figure ***a*** has *the property of being perfectly square*” or “figure ***a*** has *the congruent-parts property*” *merely* on the basis of the judgment that “figure ***a*** has *that shape*”? It depends on how we understand the content of the demonstrative “*that shape*”. If this is understood simply to mean “*that perceptual appearance*” (i.e., the *total* manner in which figure ***a*** is represented in perception), then it does seem plausible that the perceiving subject’s subsequent judgments, about further properties figure ***a*** is perceptually represented as having, can be regarded as inferences made *only* on the basis of the demonstrative judgment “figure ***a*** has *that appearance*”. Given that figure ***a*** does depict what seems

commonsensically to be just a certain shape, it is tempting to think that when it comes to the appearance of figure *a*, we can treat “*that shape*” and “*that appearance*” as equivalent.

It is clear, however, that Giaquinto intends “*that shape*” to pick out a more specific property of figure *a* (as represented in perception), which it has in common with other figures that do not necessarily have the same appearance. This is evident in the latter of the two belief-forming dispositions that make up (PS): that upon perceiving a figure as perfectly square, and also thinking of its apparent shape demonstratively, as “*that shape*”, one will believe that anything that is perfectly square will have *that shape*. If this is to be a perspicuous belief-forming disposition, it has to be the case that all perfect squares have the same *shape*. But we have already seen that all perfect squares do not have the same perceptual *appearance*: If a perfect square is oriented obliquely, it may appear not as square but rather as rhombic, or diamond-shaped. This shows that “*that shape*” will have to be understood as a picking out a property that abstracts away from differences in orientation, as well as differences in spatial location and size—indeed, if this were not the case, it could hardly serve to underwrite geometrical beliefs concerning properties that hold for all geometric squares. But this means that the content of “*that shape*” will have to abstract away from many of the features that are part of the perceptual appearance itself.

As such, it seems that the belief-forming disposition {r.u.q.} does not apply to these cases in quite the way Giaquinto supposes. In particular, it is not directly

engaged by the general demonstrative concept *{that shape}*, but only by the more specific demonstrative concept *{that appearance}*. This is for the straightforward reason that the universal generalization to “Every F has G” is only valid (and is only licensed by Giaquinto’s {r.u.q.}) provided that “x has G” has been inferred by assuming *nothing* about *x* aside from its having F. This condition is easily met in the context of formal logic, where we can begin by introducing an ‘arbitrary *x*’ simply as a free variable, and can then reason about *x* in a fully general manner, by ensuring that any assumptions made about *x* (on which our subsequent judgments of the form “*x* is G” will be based) are discharged in the generalization step. In the present context, however, it seems that “*x*” is just figure *a*, *as represented in the perceptual experience of the subject*, and the judgments corresponding to “*x* is G” are *perceptual judgments* to the effect that figure *a*, as perceived, has the properties *being square* and *congruent-parts*. The absence in perception of anything corresponding to a truly arbitrary figure *x* (about which nothing at all is assumed) might lead one to think the whole attempt to apply {r.u.q.} to perceptual judgments is misguided from the outset. After all, we don’t even seem to be capable of perceptually representing a figure adequate to the description “arbitrary figure *x* is *that shape*”—for we will always have “assumed” in our perceptual experience certain features, such as orientation, that are not mandated by the figure possessing the property of being *that shape*.²³ Upon reflection, however, it is straightforward to recognize that {r.u.q.} *can*

²³ One possible response on behalf of Giaquinto might be to say that when he speaks of the subject thinking of the figure’s apparent shape demonstratively, what he has in mind is that the subject thinks of the figure’s perceptual *appearance*, as a whole,

be made to apply to perceptual judgments, provided that we specify our assumption “ x is F ” as the assumption that the (“otherwise arbitrary”) figure-as-it-appears x has precisely the appearance represented in the subject’s own perceptual experience.

For it *does* seem true that subjects are able to infer that figure a should have the properties *being square* and *congruent-parts* on the sole basis of the demonstrative judgment “figure a has *that* appearance”; after all, *this* demonstrative judgment encompasses all the features of the perceptual representation of figure a on the basis of which those further properties are perceived to hold. The subsequent universal generalization will then be valid: The subject possessing {r.u.q.} will then be apt to believe that *any* figure that was exactly as figure a appeared to be (not just in ‘shape’, but in general) would have the geometrical properties that figure a appears to have. This is not enough to get us to a belief about squares in general, of course, but the belief formed by the “triggering” of this disposition by perception of an appropriate figure *is* nonetheless a belief with some nontrivial scope of generality, and it is of essentially the sort Giaquinto is aiming at; insofar as it depends crucially on sensory experience, but in a thoroughly non-evidentiary role, it can reasonably be

under the description “*arbitrary* figure of *that* shape”. In that case, {r.u.q.} *could* be used to validly generalize over perceptual judgments the subject goes on to make in relation to the figure, while thinking of its appearance in this way. Moreover, the idea is not implausible, given that we often *do* seem to think of a particular drawn figure as (depicting) “an arbitrary square”, for instance. But this response only defers on the real question at issue, which now takes the form of how it is possible for the subject perspicuously to think of the figure’s appearance under the description ‘*arbitrary* figure of *that* *shape*’. How is the subject able to recognize that the figure as it appears *does* serve suitably as an arbitrary instance of that shape? This is just the question for which Giaquinto’s account lacks a clear answer; on the alternative account sketched below, it is suggested that dynamic imagery can underwrite the perceptual recognition that the particular square figure serves as an arbitrary instance of its shape.

described as “synthetic *a priori*”. That sensory experience is playing a non-evidentiary role here is implied by the fact that the justification for the belief—that any figure that *was* exactly as this one *appears to be* would have the property or properties in question—does not depend on the assumption that this figure *is* as it appears to be (that perception is veridical, that is).²⁴ This is the important grain of truth in Giaquinto’s claim that *concepts* are sufficient to take us from the visual perception of a particular figure to knowledge of a general truth of geometry.

The problem is that the belief in question, while it *does* ascribe geometrical properties to objects of a certain general class, is only able to pick out that class of objects by means of demonstrative reference to one’s own private perceptual experience. Moreover, the demonstrative concept {*that appearance*} seems too fine-grained to be able to engage the belief-forming dispositions of our general geometrical concepts, such as {perfect square}, which abstract away from various features of particular appearances. The concept {*that shape*} *is* pitched at the right level to comport with our concepts for geometrical objects, but as we’ve seen,

²⁴ One might wonder whether this fact, that the belief’s justification does not depend on veridical representation of anything in the *external* world, is really sufficient to establish that the belief is truly “*a priori*” in the sense that it does not depend on the use of sensory experience as *evidence*. An alternative view might try to construe the belief’s justification as ultimately empirical, in the sense that it depends on one, in effect, *observing* oneself having a perceptual experience as of the figure, and perceptually judging that the figure (as it appears) has this or that property; on this sort of view, one might argue that it is a condition on justification that one’s higher-order (perhaps ‘introspective’) awareness *represents* the (first-order) perceptual appearances in a way that is veridical. This challenge raises general questions concerning the epistemology of perceptual experience that are beyond the scope of this chapter; they will need to be addressed carefully, at a later point. Giaquinto himself considers and responds to a version of this challenge (2007, 56-9).

ascriptions of {*that shape*} are unable to support universal generalizations from perceptual experience in accordance with {r.u.q.}, precisely because they abstract away from some of the features necessarily present in the perceptual representation of the figure, which is itself what grounds the ascriptions of geometrical properties that we would seek to generalize. Of course, it might well seem *obvious* to the subject perceiving figure *a* that the geometrical properties observed to hold of it depend only on the figure's *shape*, and not on its size, position, or orientation—but this cannot be taken for granted in the present context, because the task at issue is precisely to shed light on the psychological underpinnings of our intuitive knowledge of just such 'obvious' truths. It appears that Giaquinto's purely concept-based approach is unable to fully account for the intuitive obviousness of the claim that *any* square has the *congruent-parts* property.

Nonetheless, the mere fact that it seems so *obvious* that the relevant geometrical properties of figure *a* depend only on its *shape* suggests that perhaps Giaquinto was correct in supposing that the generality of this belief hinges on something like the ability to think demonstratively about "*that shape*", in abstraction from the more fine-grained property picked out by "*that appearance*". It is unclear, though, what it is for a subject to think demonstratively about a given shape in itself. Given the central role played by {*that shape*} in Giaquinto's account, it is somewhat surprising that he doesn't offer an explicit account of what it is to possess a general concept of *shape*, such that one could come to form the belief that two figures, with different appearances, are the same in shape (the square and the diamond, for

instance). Nonetheless, he does offer a suggestive hint, in an endnote that clarifies the notion of a “similarity mapping”, in terms of which the concept of *correspondence* was defined (as we saw earlier). Here is the endnote, which is worth quoting in full:

A similarity mapping is a shape-preserving transformation, such as uniform expansion or contraction, rotation, translation, or any composition of these. (We include the null transformation among similarity mappings.) Let *a* and *b* be similar, i.e. figures with the same shape. Imagine *a* contracting or expanding uniformly until it forms a figure *a'* the same size as *b*; then imagine *a'* moving so as to coincide with *b*. Any such similarity mapping maps each line through *a* onto a line through *b*. (2007, 47, note 4)

While Giaquinto’s aim in this endnote is merely to specify the notion of “similarity mapping” that figures in his definition of *correspondence* (of lines), the passage provides a clear indication of how a general concept of *shape* might be specified, in terms of the ‘shape-preserving’ transformations that constitute a similarity mapping. Suppose, for instance, that one possesses the concept *shape* if and only if one is disposed to believe a figure to have the *same shape* as a given figure just in case one is disposed to believe that the former can be transformed into the latter (or into coincidence with it) by means of some composition of uniform expansion or contraction, rotation, and translation. Giaquinto even offers a description of the kind of perceptual experience that could give rise to a judgment that *x* and *y* have the same

shape, namely, an imaginative rehearsal of the sequence of spatial transformations that are required to make the figures coincide.

This proposal helps to flesh out the idea that one can think of the shape of figure *a* demonstratively, as “*that shape*”. What one will be thinking of here is the property shared by a general class of figures, such that they can be mapped onto coincidence with figure *a* (or *vice versa*) by a sequence of expansion/contraction, rotation, and translation. Suppose we then postulate, along the lines of the dynamic imagery account outlined earlier, the existence of *dynamic transformational imagery* that enables the subject to imaginatively rehearse spatial transformations of these kinds, in such a way that the representation of the figure-as-transformed is integrated with the representation of the figure-as-untransformed in perceptual content. We will then have at least a sketch of a plausible explanation for why the subject perceiving figure *a* should find it *obvious* that the geometrical properties the particular figure is perceived to possess are ones that depend only on the figure’s *shape*. For the subject will be able to perceive the figure against a backdrop of imagined transformations of the shape-preserving kinds, and will thereby be able immediately to see that transforming the figure in any these ways will still leave a figure that can be perceived as (for instance) a perfect square.

This proposal provides at least a plausible route to explaining how the perception of a *particular* figure as having a certain geometric property (such as being perfectly square) can lead to a judgment that *any* figure with *that shape* would have that geometric property. It is importantly different from Giaquinto’s purely

concept-based proposal, which supposes that {r.u.q.} can license the generalization “any figure with *that* shape is perfectly square” on the grounds that the subject has drawn the inference “*x* has *that* shape, so *x* is perfectly square”. The problem, as we have seen, is that {r.u.q.} licenses no such generalization, because the subject has drawn no such inference: The subject never perceptually represents an *otherwise arbitrary* figure *x* with *that* shape—rather, the subject represents the total *appearance* of figure *a* and judges of the particular figure, as it appears, that it is perfectly square. Since the appearance includes the ‘assumption’ of features aside from *shape*, {r.u.q.} will not automatically attribute the subsequent property ascription to the shape itself. Moreover, Giaquinto’s account as it stands is unable to explain how the subject might be able to apprehend the distinctive dependence of the geometric property on the figure’s *shape*, as opposed to its size, orientation, or position. That is just what the present proposal is able to do—by appealing to dynamic imagery. It proposes that the general concept of *shape* can be specified in terms of invariance under a certain set of spatial transformations. By imaginatively rehearsing those transformations, applied to the particular figure, the subject is able to see that the properties seen to hold of the particular figure are themselves invariant under the transformations. Accordingly, the subject is in a position to judge that the properties (such as *congruent-parts*) seen to hold in the special case (that is, of the particular figure, as it appears to be) will hold for all figures of the same shape. This is just an initial sketch of how dynamic imagery can be applied toward a solution of the restricted version of the generality problem that arises in cases like the one Giaquinto considers. The purpose in

sketching the account here is to illustrate that the appeal to dynamic imagery is equipped to do real explanatory work at precisely the point where the explanations provided by Giaquinto's concept-based approach appear to encounter problems.

Chapter 4: Basic Knowledge of Geometry

1 Introduction

Having in the previous chapter considered and critically evaluated Giaquinto's account, we proceed in this chapter to consider how the insights thereby gained can help to set the agenda for developing a more comprehensive account of basic geometrical knowledge, one that is capable of explaining our intuitive belief in the truth of Euclid's postulates.

2 An assessment of Giaquinto's account

There is a good deal of extant literature concerning the visual perception of geometrical forms, on the one hand, as well as concerning the epistemological foundations of Euclidean plane geometry, on the other. The account we have just examined in the previous chapter, however, is distinctive in that it considers basic geometry from both psychological and epistemological vantage points, and aims to connect these two domains in order to explain how our knowledge of the fundamental propositions of Euclidean plane geometry could be grounded in features of human visual perception. As we have just seen, however, Giaquinto's account as it stands fails to provide a fully satisfying answer to some of the key questions it aims to address in relation to the specific geometrical proposition it takes as its case study. Even to the extent that it is successful in handling this case, it is unclear how its approach might be generalized to apply to other, even more fundamental,

propositions of geometry. This chapter takes up the project Giaquinto has set out on, and attempts to provide a more comprehensive account of basic geometrical knowledge, which aims to explain our intuitive knowledge of the *most* fundamental propositions of Euclidean geometry, corresponding to Euclid's own postulates. It will be argued that the postulation of dynamic imagery plays an important role in the acquisition of this knowledge. The account to follow is based in part on features of Giaquinto's own account, and is at the same time partly motivated by the limitations of this account. Accordingly, we begin by reviewing some of his account's key virtues, along with some of its limitations.

One important virtue of Giaquinto's account is the central role he assigns to the perception of orientation and symmetry in the understanding of geometric figures. To begin with, Giaquinto draws attention to the manner in which the perception of form depends on the assignment of orientation, by means of aligning the perceived figure with a *reference system* that consists of certain axes of orientation, whose assignment in relation to the figure is sensitive to the structural features inherent to the figure itself. This latter point will prove important, because it will help to explain how Euclidean spatial properties can arise 'locally', in the space the figure is perceived to set up around itself. The fact that a reference system can also be assigned by the conscious direction of attention is important as well, since it provides a mechanism by which a perceiver can actively determine a shift in the perceived aspectual integration of a figure. Giaquinto also draws attention to the crucial fact that the axes of orientation serve additionally as axes of *symmetry*. As illustrated by the

square-diamond example, reflection symmetry can be shown to influence the perception of certain geometrical properties, such as the equality of the lines or angles ‘mapped’ to one another by reflection about an axis. As we also noted, reflection symmetry appears to structure the perceived *dynamics* of the figure, in the sense that certain spatial equalities are ‘locked in’ by the reflection symmetry, while others are left ‘variable’, depending on the assignment of the axis of reflection. In particular, the perception of reflection symmetry seems to determine the range of spatial *translations* that can easily be imagined: Translations *across* the axis, which would violate the reflection symmetry, are precluded, while translations *along* the axis, which respect the reflection symmetry, are perceived as possible or even salient. Symmetries in general will play a central role in the application of dynamic imagery to basic geometry, precisely because of the deep relationship between the perception of symmetry and the imagination of spatial transformations.

The other key virtue of Giaquinto’s account that bears special mention is the way that he combines ideas from Kosslyn and Peacocke in order to explain how the possession of *perceptual* concepts for geometrical objects or properties (such as *square*) can dispose the perceiver to form beliefs about the geometrical properties of a figure automatically, on the basis of the visual detection of features such as symmetry. By building (Kosslyn-style) visual *category specifications* directly into the content of geometrical *concepts* (understood, Peacocke-style, as belief-forming dispositions), Giaquinto provides a framework in which geometrical concepts can be identified in terms of the visual experiences that would be apt to generate beliefs

employing those concepts.²⁵ This general approach will prove especially fruitful in combination with dynamic imagery, since imaginative rehearsals of spatial transformations will themselves be included among the visual experiences capable of engaging our geometrical concepts. We will then be able to characterize our concepts for geometrical relations like *parallelism* and *congruence*, for instance, in terms of our dispositions to make *judgments* of parallelism and congruence on the basis of imaginatively rehearsing spatial transformations of certain specific kinds.

Giaquinto's account also has some significant limitations, however. For it is unable to provide a complete explanation of our intuitive knowledge of the proposition Giaquinto takes as his case study: that every square divides into two congruent parts along its diagonal. The judgment that this proposition holds can usefully be thought of as involving two stages: In the first stage, the subject judges that the *congruent-parts* property holds in the special case, that is, for the particular

²⁵ As we saw earlier, Giaquinto stresses the distinction between “perceptual concepts” (which are vague and therefore apply to approximate instances) and “geometrical concepts” (which are still based in the same visual category specifications, but which properly apply only to perfect instances). While the distinction he draws here is quite real, it is unclear that it is best captured as a distinction between different *concepts* (or different *kinds* of concepts). After all, the vague “perceptual” concept already implicitly characterizes the sort of object that would satisfy the more exacting “geometrical” concept. Indeed, the view that *one and the same* concept engages with both approximate *and* perfect instances seems to provide the most natural means for explaining how a visual experience of a perceptibly *imperfect* drawn figure can activate visual imagery of a *perfect* instance, which is imaginatively ‘projected onto’ the visual experience of the drawn figure, along the lines Giaquinto himself suggests (2007, 39). Instead of drawing a distinction between two sorts of concepts, then, we will speak more inclusively about ‘perceptual concepts for geometrical properties’, and will understand these concepts to apply properly just to perfect instances, but to be responsive to imperfect (or approximate) ones as well, and thereby to support the judgment that a line is *approximately* straight, for instance.

figure, as it appears; in the second stage, the subject judges that the first judgment holds in general—that the *congruent-parts* property is possessed by *all* squares. As we've seen, Giaquinto's account faces unresolved problems at both stages. In explaining the first stage, he assumes that the perception of reflection symmetry issues directly in a belief about congruence, but this leaves a mystery concerning how the belief-forming dispositions of the *general* concept for congruence come to be engaged by the perception of symmetry. This is precisely where dynamic imagery can be useful: If we understand the concept *congruence* in terms of the rigid spatial transformations that would suffice to demonstrate congruence, we have a natural link to the perception of reflection symmetry, given that *reflection* (which can be imagined as folding across the axis) is among the relevant spatial transformations. Giaquinto's explanation of the second, generalization stage is similarly incomplete, because his account turns on an appeal to the demonstrative content {*that shape*}, which proves unable to engage the belief-forming dispositions of his {r.u.q.} in the way he supposes. As we've seen, however, putting pressure on his appeal to {*that shape*} leads quite naturally to a view in which the concept *shape* can be engaged directly by the imaginative rehearsal of dynamic spatial transformations, in a way that can provide the subject with an intuitive basis for generalizing over the class of figures with the same shape. The general impression, then, is that if we take seriously the question of how the more *fundamental* concepts he invokes, such as *shape* and *congruence*, can be constituted so that they are engaged by our visual experiences in

the appropriate ways, we find ourselves pushed in the direction of granting an important role to dynamic imagery.

This consideration already highlights an important limitation of Giaquinto's account: It is not pitched at as fundamental a level as one might desire for an account of basic geometrical knowledge. Giaquinto takes as his case study a proposition about squares, and accordingly, considers {perfect square} as his key example of a geometrical concept. It is evident just from examining his category specification for squares, however, that this concept presupposes more fundamental ones, for it specifies that the edges are "straight" and the opposite ones "parallel". It seems reasonable to suggest that, in the absence of a prior account of what is to perceive a line as *straight*, or two straight lines as *parallel*, one has not really been given a complete account of what it is to perceive a figure as *square*. Moreover, we have just seen that, in attempting to explain our knowledge of this proposition about squares, Giaquinto's account encounters problems precisely around the appeal to the more fundamental concepts of *shape* and *congruence*, which are simply taken for granted. This suggests that what is needed is an epistemological account that approaches the subject matter of geometry beginning with its most fundamental concepts and principles—only on that basis will we be able to provide a truly comprehensive account of geometrical knowledge. Of course, the classical statement of the fundamental concepts and principles of Euclidean plane geometry is given by Euclid himself, in the list of definitions and postulates that are stated at the beginning of

Elements. In order to set the agenda for the investigation to follow, we will look to Euclid as our guide.

3 Applying dynamic imagery to Euclid's starting points

In the rest of this chapter, then, we will put dynamic imagery to work explaining our most fundamental intuitive knowledge of Euclidean plane geometry, and we will take knowledge of Euclid's postulates as our target. The five postulates are stated as follows:

1. To draw a straight line from any point to any point.
2. To produce a finite straight line continuously in a straight line.
3. To describe a circle with any center and distance.
4. That all right angles are equal to one another.
5. That, if a straight line falling on two straight lines make the interior angles on the same side less than two right angles, the two straight lines, if produced indefinitely, meet on that side on which are the angles less than two right angles.²⁶

Of course, Euclid's postulates are only meaningful provided that one has a grasp of the concepts they employ: *straight line*, *right angle*, and so forth. In considering each postulate, we will begin by specifying a perceptual concept for each of the

²⁶ Here and throughout, I follow Heath's edition of *Elements* (Euclid, 1956).

geometrical terms appearing in the statement of the postulate. These perceptual concepts will be specified along the lines indicated by Giaquinto, in terms of the conditions that must be met by the perceptual experience of a figure, in order that it should be apt to issue in a *judgment* ascribing the relevant geometrical property to the figure. Having in this manner specified the concepts involved in the postulate, we will proceed to explain how the subject can acquire intuitive knowledge of the postulate, by means of perceiving an appropriate figure, ‘animating’ the perceptual experience of the figure by employing dynamic visual imagery, and bringing the relevant concepts to bear on the experience, thereby arriving at the judgment that the postulate is true. Throughout, our specifications of the concepts will be informed by a careful reading of Euclid’s own definitions; at the same time, we will be guided by the constraint that our geometrical concepts should be constituted in such a way that they are apt to be engaged by possible perceptual experiences of appropriate figures. Moreover, rather than assuming that the subject antecedently possesses concepts for geometrical objects and properties, we will, when possible, aim to show how these concepts themselves can arise synthetically out of the interplay between the visual system’s inherent sensitivity to certain spatial features, on the one hand, and the perceived forms themselves, on the other.

In explaining intuitive knowledge of Euclid’s postulates (or their equivalents), we will rely only on a core set of minimal assumptions about the visual perception of drawn figures. We will not rely on *any* assumptions about the geometry of the subjective ‘visual space’ in which the figures are perceptually ‘given’ to the subject:

In particular, we will not assume that visual space inherently possesses the properties of *zero curvature* (the ‘flatness’ of Euclidean space guaranteed by Euclid’s fifth postulate) or even *constant curvature* (the uniformity-in-all-directions guaranteed by his fourth postulate, which characterizes not only Euclidean, but also elliptic and hyperbolic spaces). This is just as well, given the compelling reasons—already suggested by informal observations regarding the many well-known geometrical-optical illusions—for doubting that visual space can be adequately described by Euclidean geometry (Suppes 1977), or indeed, by *any* single, consistent geometry (Wagner 2006). On our approach, Euclidean spatial properties are not held to be *a priori* in the sense that the perception of a drawn figure already involves or includes its being perceived as located ‘within’ some preexisting space that has Euclidean properties already ‘built in’. Instead, it is the *figures* that are held to come first; Euclidean space will always arise locally, *in* the figures, and will be ‘unfolded’ out of them, *as* the space they effectively *set up* around themselves. In fact, in the context of the present account, it is in a certain sense misleading to talk of “Euclidean space” at all, since we will only ever properly be concerned with Euclidean *figures*, and with the manner in which they enforce Euclidean structure onto the (extended) figures that can be ‘unfolded’ or constructed out of them. We will never have a need for a Euclidean “space” (or an intuition thereof) to serve as a ‘venue’ that provides for the possible presence of Euclidean figures ‘within’ it.

The key assumptions we *will* make concerning the visual perception of drawn figures can be roughly indicated as follows. We inherit our first two assumptions

more or less directly from Giaquinto's account, with some modifications. The first is that the visual system assigns to any perceived figure a *reference system*, consisting in the alignment of the figure with one or more *orientation axes*, which themselves can serve secondary roles as axes of perceived symmetries. Based on observations made earlier, we will assume that a privileged "up" direction is inherent to any reference system, and that the visual system strongly prefers to assign the *primary* axis or axes of orientation vertically and/or horizontally with respect to this "up" direction (and otherwise, along one of the diagonals).²⁷ Secondary axes may also be assigned, at various orientations. In addition, we assume that the assignment of reference system can be influenced by both environmental orientation and the figure's 'intrinsic' orientation, as well as by deliberate attention.

We also borrow from Giaquinto the assumption of *perceptual concepts* for geometrical properties. We will construe these rather more narrowly, as consisting in dispositions that visual experiences representing certain kinds of objects as having certain *spatial symmetries* should issue directly in judgments ascribing the correlative *geometrical properties* to those objects. (That is, we will attempt to establish a much tighter connection between geometrical concepts and the perception of *symmetry* in particular.) Similar to Giaquinto's own view, we will allow that these concepts are

²⁷ The proposal that *any* of these four orientations (and not just the vertical) can be assigned as primary with respect to "up" comports well with the psychophysical data reported by Wenderoth (1994), which shows that sensitivity to reflection symmetry peaks at these four specific orientations, with greatest sensitivity at the vertical orientation, followed by the horizontal, and finally the two diagonals. It also explains our apparent ability to perceive straight lines (which have only one salient axis) *as* horizontal or oblique. Presumably other orientations can be accommodated to one of these four, by means of shifting the assignment of the "up" direction.


vague, such that they can support judgments of both *approximate* and *perfect* straightness (for instance), and will also suppose that figures that *approximately* satisfy the relevant visual category specifications are apt to give rise to visual imagery of an idealized or ‘perfect’ version of the (imperfect) figure, which is overlaid onto one’s visual experience of the figure itself. In this manner, the perceiver is able to *see* the perfect (ideal) object *in* the imperfect (drawn) figure—in something akin to the way one can look at an ordinary picture and *see* its depictive content *in* the surface markings.²⁸ In general, we will understand judgments of *perfect* straightness, etc., to be ascriptions of these properties, not to the drawn figures themselves, but rather to the ideal geometrical objects that are their depictive contents.

Our two final assumptions are just the two forms of dynamic visual imagery postulated by the dynamic imagery account. As is by now familiar, dynamic *aspectual* imagery captures the perceiver’s ability to see a figure (as Giaquinto put it) ‘in two ways at once’—that is, to see it as *seeable* in multiple ways. Of course, this presupposes the ability to see a figure in *some or other* ‘way’, which in the present context comes down to three things: parsing a figure into constituent parts, assigning it a reference system, and perceiving symmetries. It also presupposes the ability to attend to the continuity of certain parts *across* aspectual shifts. For instance, in understanding the proof of Euclid’s Proposition I.1, the perceiver will need not only to represent each of the three lines, alternately, as *radii* of the relevant circles, and

²⁸ See Wollheim (1980) for the original statement of the ‘seeing-in’ idea in relation to pictorial perception.

also as *sides* of the triangle, but moreover, to grasp in each case that it is the very *same* line that belongs to both structures (see Figure 4.1).

Figure 4.1: Diagram for Euclid’s Proposition I.1



In order to accommodate this perceptual awareness of continuity *across* alternative aspectual integrations, we will assume that visual perception includes something like the “object files” that have been postulated to explain subjects’ visual tracking of persisting individual objects, independently of the properties that subjects ascribe to those objects in perceptual representation (Kahneman and Treisman 1984, Kahneman et al. 1992).

Dynamic *transformational* imagery captures the ability to visually imagine the spatial transformation of a figure or one of its parts—in effect, to see it ‘in two *positions* at once’ (the position it *does* occupy, and the one it *would*, under the relevant transformation). While we can visualize spatial transformations of various

kinds, for the purposes of this chapter, the relevant transformations are just the *rigid motions* (or ‘isometries’) of reflection, translation, and rotation, which we take as basic. We will assume that *reflections* of planar objects are imagined as rotations in three-dimensional space about the axis of reflection—effectively, the object gets ‘flipped over’ (or ‘folded over’) so that it comes to rest again in the original plane. Translations and rotations, on the other hand, are imagined as motions *within* the plane, either along an axis (in the case of translation) or about a point (for rotation). In addition to the ability to imagine spatial transformations applied to the *objects* represented by perception, we also include, as part of dynamic transformational imagery, the ability to imagine the visual consequences of adjustments in *subjective* perspectival vantage point: for example, visually *scanning* across an image, or visually ‘zooming in’ on, or ‘zooming out’ from an image.²⁹ These latter transformations do affect the subject’s perceptual representation of the relevant object (e.g., a line may *appear* to grow smaller when one imaginatively ‘zooms out’ from it), but they are thought of as transformations of one’s perspective on a stable object, rather than transformations of the object itself. Finally, we will assume that different transformations can be imagined concurrently, and that transformational imagery in general can operate concurrently with aspectual imagery, in such a way that the two kinds of dynamic imagery interact in perception. For example, given that dynamic *aspectual* imagery may involve parsing a figure into different sets of component

²⁹ See Chapter 10 of Kosslyn (1994) for discussion of the inspection and transformation of visual images, including the operations of imaginative ‘scanning’ and ‘zooming in’.

parts, the choice of parsing will have implications for the range of dynamic *transformations* that can be applied to these parts. Conversely, since the perception of *symmetries* will on this account be understood in terms of the rehearsal of *isometries* (that is, rigid motions), transformational imagery will have implications for aspectual integrations, which are themselves partly structured by perceived symmetries.

These are all the assumptions, then, that we will need in order to explain our intuitive knowledge of Euclid's postulates: the ability to assign a *reference system*, to apply *perceptual concepts*, and to employ *dynamic aspectual imagery* and *dynamic transformational imagery* in an integrated fashion. If this small set of core assumptions is indeed sufficient to this task, it is a remarkable fact, since these assumptions exclusively concern processes of *visual perception*, rather than any distinctively 'cognitive' processes. As we will see in the next chapter, this same set of core assumptions is in fact sufficient to explain, further, how our intuitive knowledge of plane geometry can be extended *beyond* the postulates, to the propositions proved in Book I of *Elements*. For the time being, however, the focus is just on fundamentals.

4 The common notions

Of course, the 'first principles' of Euclid's system are only *partly* comprised by the five postulates; these stand alongside his five "common notions", so called because they are meant to capture assumptions that are not specific to geometry, but which

geometry shares in common with other sciences.³⁰ The common notions are stated as follows:

1. Things which are equal to the same thing are also equal to one another.
2. If equals be added to equals, the wholes are equal.
3. If equals be subtracted from equals, the remainders are equal.
4. Things which coincide with one another are equal to one another.
5. The whole is greater than the part.

The common notions are noticeably more abstract than the postulates: Instead of specific objects, they concern the general notions of equality, addition/subtraction, greater/lesser, and whole/part, which will end up being applied by Euclid to quite different *kinds* of objects—*lines*, *angles*, *triangles*, and so forth—in relation to which these notions will take on quite different spatial meanings. (For instance, the *greater-than* relation as it applies to lines is quite different from the corresponding relation that applies to angles.) This abstract character might seem to make the common notions inherently bad candidates for being grasped by means of pictorial perception, and hence to place them outside the scope of our account. One might therefore naturally be led to conclude, against what has just been claimed, that we will need to take on further assumptions that are not of a distinctively perceptual character, but

³⁰ This goes back to Aristotle's division of first principles into the *postulates* that are peculiar to a given science, and the *axioms* that are common to all sciences. See Heath's editorial commentary in Euclid (1956), pp. 117-122.

instead concern the subject's possession of prior knowledge of a more abstract, logical character. Upon further reflection, however, the common notions can be seen to accord much more naturally with the perception-based approach of our account than initial impressions might suggest.³¹

It is useful to begin by considering Common Notion 4, which is quite plausibly read not as an axiom at all, but rather as an attempt at *defining* the abstract notion of *equality* (as it figures in Euclidean plane geometry, that is) in terms of the perceptible spatial property of *coincidence*. Matters are somewhat complicated here by the fact that Euclid employs the notion of *equality* in plane geometry in two distinct senses. First, there is equality as in the sense of *congruence*, which applies univocally to lines, angles, and plane figures. For instance, Proposition I.4 states that if two triangles have two sides equal, and the angle contained by those sides also equal, the triangles will themselves be “equal” to one another (Euclid 1956, 247). Here the claim that the triangles are equal is plausibly read as the claim that the triangles are *congruent*—that is, identical in both shape and size. On the other hand, Proposition I.38 reads as follows: “Triangles which are on equal bases and in the same parallels are equal to one another” (Euclid 1956, 333). In this case, the equality asserted between the triangles cannot be understood as congruence, because triangles

³¹ Of course, if the task were to explain the subject's knowledge of *unrestricted* versions of the common notions, which would apply outside of spatial contexts, the challenge of providing a thoroughly perceptual basis for this knowledge would be much greater. Here our task is merely to explain knowledge of the common notions as they apply in the context of plane geometry; as such, we will only be concerned to explain knowledge of the equality and relative magnitude of the relevant objects, namely: lines, angles, and plane figures.

on equal bases and in the same parallels are not, in general, congruent; here “equality” must rather be understood to refer to equality of *planar area*.

Let’s consider, first, equality in the sense of congruence. If we read Common Notion 4 in this sense, it states that if two things coincide, then they are congruent. Here we may take *coincidence* to be the relation two objects have when they align with each other perfectly, such that one object is directly ‘on top of’ the other in the plane. It is standard, however, to take Common Notion 4 as licensing the *method of superposition*, in which one plane figure is ‘applied to’ another, so that they coincide.³² This is the method Euclid uses to prove Propositions I.4 and I.8; in both cases, he reasons that *if* one triangle is applied to another, such that certain of their parts coincide, the triangles’ other parts will coincide as well, and will therefore be equal. If we interpret Common Notion 4 together with the method of superposition, we can see that it follows that two things are congruent, not merely if they *do* coincide, but if they *can be made* to coincide by means of the ‘application’ of one onto the other. Moreover, Euclid’s own reasoning in Propositions I.4 and I.8 appeals not only to Common Notion 4, but also to its *converse*: In particular, Euclid reasons that if one straight line is applied to another that is equal to it, then the two straight lines will coincide. This indicates that (at least restricting ‘equality’ to the sense of *congruence*, which is the sense that applies to straight lines) Euclid accepts a *biconditional* form of Common Notion 4. Putting these points together, we have the basis for a definition of congruence: In particular, two objects are *congruent* if and

³² See Heath’s commentary in Euclid (1956), p. 225.

only if they can be made to *coincide* with one another by means of the ‘application’ of one object onto the other. The important thing to notice here is that the property in terms of which congruence is defined—namely, potential coincidence under some possible ‘application’—is itself a perceptible *symmetry* of the kind that is able to engage a *perceptual* concept for geometrical congruence, on our approach. In particular, a concept for congruence can be specified in terms of the disposition to believe that two planar objects are congruent if and only if one believes that some sequence of rigid motions would map the objects into spatial coincidence with one another. The former belief is a geometrical one; the latter is perceptual, grounded in the capacity to imaginatively rehearse visual imagery of the spatial transformation of one of the perceived objects, by reflection, translation, and rotation. So far, then, Common Notion 4 seems to comport quite naturally with the core assumptions of the dynamic imagery account.

As already indicated, there is an additional complication, resulting from Euclid’s use of the term “equal” sometimes to refer not to congruence but to *equality of area*—a use that first occurs in the statement of his Proposition I.35. This proposition is the first in a series, in the course of which Euclid develops what is effectively an account of *area-preserving transformations* (which, unlike the ‘basic’ transformations considered in this chapter, involve deformations of planar objects). The culmination of this series is Proposition I.45, in which Euclid provides a general procedure for constructing a parallelogram (of any angle-size) that is equal in area to *any* given rectilinear figure. While a complete discussion of area-preserving

transformations will need to be deferred to the next chapter, those transformations, which can be specified in a way that is independent of the notion of *area*, will enable us to specify a *complete* perceptual concept for *equality of area*, along the same lines as our perceptual concept for *congruence*. In particular, possession of this concept will consist in a disposition to judge two plane figures to be equal in area just in case one judges that the two plane figures can be made to coincide with one another other under the appropriate spatial transformations. For the time being, however, we can understand Common Notion 4, on the reading under which it refers to equality of *area*, simply as it is stated: that is, as asserting that coincidence is a *sufficient* condition for equality (of area). Taking account of the associated method of superposition, we can say that Common Notion 4 corresponds to a *part* of our perceptual concept for equality of area, which consists in a disposition to judge two plane figures to be equal in area whenever one believes that they can be made to coincide by applying some sequence of rigid spatial transformations.

In the context of this understanding of Common Notion 4, it is now possible to interpret Common Notion 5 along similar lines, as playing a definitional role, rather than an axiomatic one. In particular, it can be read as an attempt to define the abstract *greater-than* relation (and, by implication, the *lesser-than* relation) in terms of perceptible *part-whole* relations. Common Notion 5 states: “The whole is greater than the part.” Now, in Euclidean plane geometry, there are three different kinds of magnitude to which the *greater-than* relation applies, and correspondingly, three different kinds of part-whole relationships: A whole line is greater (in length) than

one of its parts (e.g., one of the lines into which it divides, when intersected by a different line), a whole angle is greater (in angular size) than an angle contained within it (e.g., one of the acute angles into which a right angle divides, when a line is drawn bisecting the right angle), and a whole plane figure is greater (in area) than one of the plane figures it contains (e.g., one of the two triangles contained by a square that has been divided along its diagonal). Of course, so far we have only said that standing in the *part-whole* relation in the right manner is a *sufficient* condition on instantiating the *greater-than* relation; we will need the converse in order to have a genuine definition of *greater-than* in terms of the *part-whole* relation. Here it is useful to borrow, from Common Notion 4, the idea of spatial coincidence and the associated method of superposition. We can then say that one object is greater than another if and only if the latter can be applied to the former, so that it coincides with one of its parts. As was the case with Common Notion 4, this will give us, initially, only a restricted concept of *greater-than*: Because ‘application’ refers to rigid motions only, we will need to restrict *greater-than* to apply only to pairs of objects that are commensurable (such that their parts can be made to coincide) by means of rigid motions. We could accomplish this by restricting our definition to apply only to *straight* lines, *rectilinear* angles (that is, angles contained by straight lines), and *similar* plane figures (or more broadly, pairs of plane figures such that one is congruent to some part of the other). The second of these three restrictions is innocent, since the only angles Euclid ever refers to in *Elements* are rectilinear ones, and it is unclear that there is any other useful notion of *angle* (given that the ‘angles’

formed by curves are sensibly regarded as identical to the angles formed by the *straight lines* tangent to those curves at the point of intersection). The other two restrictions, however, are not so innocent. Eventually, we would like our concept of *greater-than* to be able to figure in judgments about a straight line being longer than a curved line, say, or a square being larger than a triangle, even when these two plane figures are of such shape and size that no set of rigid motions will make one to coincide with a part of another.

The problem is that rigid motions do not provide a general means for comparing all such pairs of objects. Like the concept of *equality* we associated with Common Notion 4, the concept of *greater-than* suggested by Common Notion 5 thereby stands in need of further elaboration, specifically concerning the spatial transformations by which, respectively, two lines of arbitrary curvature, and two plane figures of arbitrary shape, can be made *commensurable*—in the specific sense that one can be made to spatially coincide with a (not necessarily proper) part of the other. We do, it seems, have at least a rough intuitive grasp of what it would look like, visually, for a curved line to be ‘pressed flat’ so that it can be compared to a given straight one, and similarly, what it would look like for the ‘matter’ of an arbitrary plane figure to be spatially reconstituted, by means of deformation, into (say) a rectangular shape such that it can be compared to a rectangle of given dimensions.³³ But what is lacking in these cases is a grasp of the specific constraints

³³ These intuitions have clear connections to the ‘conservation tasks’ used by Piaget to demonstrate children’s achievement of the ‘concrete operational stage’ on his theory of cognitive development, which elicit, for instance, judgments about whether

that guarantee that these transformations will conserve the magnitudes of linear length and planar area. In contrast, we are supposing that our intuitively *basic* transformations (the rigid motions) are instrumental in fixing the *concepts* of these magnitudes in the first place, *via* the concept of equality as applied to these magnitudes; since equality of magnitude is *defined* in terms of coincidence under rigid motions, there is no room for a question to arise about whether the rigid motions guarantee conservation of magnitude.

To summarize where we have gotten to, Common Notion 4 (together with the rigid motions implicated by the associated method of superposition) can be understood as *defining* equality as (potential) spatial coincidence under some set of rigid motions. Common Notion 5 then associates the *greater-than* relation with the spatial *part-whole* relation; together with Common Notion 4, it effectively defines the *greater-than* relation as (potential) spatial coincidence, of one object with a *part* of another, under some set of rigid motions. Neither definition is *complete*, however, because we would like our notions of (equal and relative) magnitude to apply, more broadly, to pairs of objects not commensurable by means of rigid motions, and we have only a rough intuitive sense of the requisite spatial transformations, which will involve deformations of planar objects. It is plausible, I believe, that non-rigid transformations can be specified in terms of the basic, rigid ones (drawing additionally on dynamic aspectual imagery). In particular, certain non-rigid spatial

liquid volume is conserved when the contents of a tall, narrow container are poured into a shorter, wider one. The connection between the intuitive understanding of conservation of magnitudes in concrete and geometrical contexts is a topic I hope to take up in future research.

transformations can be shown to preserve equal area, by assuming only the standard of equality (of area) that is provided by potential coincidence under the *rigid* motions. For the time being, however, we do have in hand acceptable definitions of *equal* and *greater-than*, so long as we restrict the application of these notions to pairs of geometrical objects that are commensurable by means of rigid motions: that is, pairs of straight lines, pairs of rectilinear angles, and pairs of plane figures such that one is congruent to some part of the other. That will be more than sufficient for the purposes of the present chapter—indeed, it will be enough to get us through the majority of the propositions that make up Euclid’s Book I.

This leaves us, then, with the first three common notions, all of which employ the concept of *equality* implicitly defined by Common Notion 4, and two of which employ, respectively, the new notions of *addition* and *subtraction*. Again, these read as follows:

1. Things which are equal to the same thing are also equal to one another.
2. If equals be added to equals, the wholes are equal.
3. If equals be subtracted from equals, the remainders are equal.

As was true of Common Notions 4 and 5, these three will need to be treated differently, depending on whether or not their scope is thought of as restricted to objects that are commensurable by rigid motions. Ultimately, they will need to be understood as having scope not bound by this restriction: In particular, the

unrestricted versions of Common Notions 2 and 3 will need to be employed (in Euclid's proof of I.35) in order to establish the validity of the area-preserving transformations that will serve to underwrite the possession of a complete, unrestricted notion of *equality* (of planar area). That more general application of these common notions will be taken up in due course; for present purposes, it is sufficient to provide an interpretation of the first three common notions in their restricted versions, that is, where 'equal' is taken in the sense of *congruent*, a sense that has been fully specified by the way we have interpreted Common Notion 4.

Under this restriction, then, Common Notion 1 says that if two objects are both congruent to some third object, then they are congruent to each other. When we consider this conditional alongside the perceptual concept for *congruence* we have articulated above, we can see that no need will arise for assuming Common Notion 1 as a *general* principle; rather, the subject will already be able to perceive that it holds in any of the particular cases for which the principle might be invoked. For on the account just provided, if the subject perceiving a figure is disposed to believe that the antecedent of the conditional holds of some particular case—for instance, that each of two straight lines is congruent to a third straight line—then the subject will also be disposed to believe that (different) rigid motions can be applied to the third straight line so as to map it onto perfect spatial coincidence with (respectively) the first and second straight lines. That much follows directly from the subject's possession of the perceptual concept of *congruence*, of the sort that we argued was implicitly specified by Common Notion 4. All that is now required is for the subject to rehearse

transformational imagery of the third line being applied to each of the first two;³⁴ the subject will already be disposed to perceptually judge these imagined applications to be rigid motions that bring the lines into perfect coincidence. If the subject now perceives the two separate motions (mapping the third line onto the first and second, respectively) as conjoined, that is, as parts of a single sequence of rigid motion that maps the first line onto the coincidence with the second (*via* coincidence with the third), the subject will thereby undergo a perceptual experience of the total figure that includes all these lines, which will be apt to issue in a perceptual judgment that the first line is congruent to the second. To put the main point more concisely, the visual experience by which the subject sees the third line *as* congruent to both the first and second³⁵ needs only to be reparsed in order to constitute an experience in which the subject sees the first line *as* congruent to the second. The use of dynamic imagery allows the subject to perceive that any *particular* application of Common Notion 1 will hold true; there is no need, then, to invoke the *general* common notion as an additional assumption.

A similar point can be seen to hold for Common Notions 2 and 3, once the notions of *addition* and *subtraction* they respectively invoke can be provided with a basis in visual perception. This is relatively straightforward: As is suggested by the

³⁴ This imagery is likely to be cued by the activation of the *belief* that rigid motions will map the third line onto coincidence with each of the first two, which is itself plausibly prompted by the judgment that the third line is congruent to each of the first two.

³⁵ By “visual experience” here, I mean to refer inclusively to both the visual imagery of the relevant motions, as well as their appraisal as *rigid* motions that map the objects onto *perfect* coincidence.

statements of these two common notions, these notions can be understood in terms of part-whole relationships. In particular, the *addition* of one object (line, angle, or plane figure) to another can be understood in terms of the inclusion of the former object within the larger object that results from considering the two objects as parts of the same whole; similarly, the *subtraction* of one object from another can be understood in terms of the consideration of the latter as constituted by two parts, one of which is the former object.³⁶ No transformational imagery is required in order to perceptually grasp addition and subtraction; the only imagistic requirement is the ability to entertain alternative part-whole integrations, which of course falls within the scope of dynamic aspectual imagery. The perceptual understanding of addition and subtraction of geometrical objects does, however, require that the subject be able to arrive at the following perceptual judgment: *that the parts are collectively equal to the whole*. This does not itself need to be assumed as a general principle, however, for it can be seen to hold in any particular case in which such a principle would apply. Shabel considers the case of an angle that is divided into two smaller angles by an interior line; she observes that in this case, “the diagram shows us that the parts coincide with the whole in the sense that they additively and exhaustively determine the whole” (2003, 20). Indeed, this spatial coincidence can be seen to hold by employing dynamic aspectual imagery to the perception of the diagram, seeing it alternately as depicting a

³⁶ Given the restrictions presently in place on our interpretation of these common notions, we will stipulate that the relevant parts and wholes should all be of the same category. For instance, we will take it as a constraint on addition of two straight lines that they be collinear, and should thereby additively constitute a whole that is itself a straight line.

larger angle with a line placed inside it, and as depicting two smaller angles that are adjacent and hence share a border. Given that the subject thereby has a perceptual experience of perfect spatial *coincidence*, Common Notion 4 applies straightforwardly (that is, without the need to invoke the rigid motions of the method of superposition), leading the subject to be apt to judge that the parts that are seen to (collectively) *coincide* with the whole are (collectively) *equal* to it (in the sense of being congruent).

While no motion needs to be rehearsed in order to grasp the way parts additively determine wholes (or the way the subtraction of parts from wholes determine other parts, as remainders), perceptual judgments to the effect that Common Notion 2 or 3 holds in some particular case *will* require the use of transformational imagery, since these common notions involve the mapping of part-whole structures onto other part-whole structures. If we consider the application of Common Notion 3 to a pair of arbitrarily positioned straight lines, for instance, it says in effect that if the two straight lines are equal, and each can be divided into two parts such that the first parts are equal, then the second parts will be equal as well. Similar to what we observed for Common Notion 1, in this case too a subject who judges that the antecedent of the conditional holds will be apt to have a perceptual experience that is itself apt to trigger a judgment that the consequent holds.³⁷ In particular, a

³⁷ Note that it is not being claimed that the subject who makes the initial judgment will *necessarily* have this perceptual experience, nor that the subject who *does* undergo this perceptual experience will *necessarily* arrive at the latter judgment. This is because there is no guarantee that the subject making the initial judgment will activate the relevant visual imagery, nor that the subject will succeed in the aspectual

subject who sees the whole lines *as* equal will do so in virtue of rehearsing the rigid motion of one whole onto perfect coincidence with the other. Throughout the course of this rigid motion, however, the whole line remains available to be seen alternately as a conjunction of two adjacent parts, themselves straight lines. Given that by hypothesis, the subject believes the *first* parts of both wholes are equal, the subject will be disposed to believe that the rigid motion of the first part of the *first* whole maps it onto perfect coincidence with the first part of the *second* whole. By employing dynamic imagery of both transformational and aspectual varieties in an integrated fashion, the subject will thereby be apt to perceptually judge that the *second* part of the first whole maps onto perfect coincidence with the second part of the second whole. This is because the subject will grasp the same rigid motion, alternately, as the motion of the *whole* line and as the motion of the conjoined parts. Since the wholes and the first parts are both assumed to coincide perfectly under the motion, the coincidence of the second parts will be perceived as a direct consequence of the fact (which the subject perceptually judges to hold) that the parts are collectively equal to the whole. The subject will then be apt to judge that the second

reconfiguration of that imagery, nor that the subject will attend to the experience of the reconfigured imagery in a manner that triggers the latter judgment. This contingency, however, does not undermine the claim that the subject *can* attain genuine knowledge that the conditional holds, nor does it reduce the account provided here to a mere psychological generalization. This is because the manner in which the initial judgment is “apt to” generate an experience, and the manner in which the (subsequently reconfigured) experience is “apt to” trigger a second judgment, is a disposition grounded in the possession of perceptual concepts for geometrical properties.

parts are themselves equal. A parallel explanation applies to the grasp of particular applications of Common Notion 2.

The upshot, then, is that in spite of initial appearances, explaining the basic knowledge captured by Euclid's common notions does not require us to posit anything beyond the resources already provided by the dynamic imagery account.³⁸ For Common Notions 4 and 5 appear to be best understood as implicit *definitions* of equality and relative magnitude, for which perceptual concepts can be specified in accordance with our approach. And Common Notions 1-3 need not be invoked as independent, general principles, for our account is already capable of accounting for the perceptual knowledge of the equalities they assert, in any given particular case. In the remainder of this chapter, then, we proceed to apply this account to the task of explaining our intuitive knowledge of Euclid's postulates. In the next section, we consider the basic objects to which all of the five postulates (including, tacitly, the third) appeal: namely, *straight lines*. Our investigation of the perceptual underpinnings of Euclidean geometry will begin with the fundamental question of what it is to have a visual experience of a geometrical straight line.

5 Straight lines

Straight lines are the 'basic' objects of Euclid's geometry in the sense that throughout the *Elements*, straight lines tend either to serve as the objects assumed as given in the

³⁸ Strictly speaking, this has only been established so far for restricted-scope interpretations of the common notions, which suffice for present purposes, but not for the treatment of planar area as magnitude, which occurs beginning with Euclid's I.35. This will be addressed in the following chapter.

*setting-out*³⁹ phase of the proof (thus Proposition I.1 begins: “Let AB be the given finite straight line...”), or else the objects assumed to be given in the *setting-out* are ones that can be constructed out of straight lines, in ways established by previous propositions. An exception to this rule is of course the circle, which might reasonably be regarded as comparably basic, since Postulate 3 directly warrants its assumption in the *setting-out*, and it is not constructible out of straight lines in the sense in which rectilinear angles and plane figures are so constructible (that is, the placement of finitely many straight lines in a certain configuration does not yield a figure that can be seen as a circle). Nonetheless (as we will see in more detail later), Euclid’s definition of the circle is given in terms of the equality of the straight lines that serve as its *radii*, and the construction of the circle assumed in Postulate 3 (which is specified relative to a given radial ‘distance’) is naturally understood in terms of the *rotation* of a straight line about one of its (fixed) endpoints—effectively the same procedure used to mechanically draw a circle using a compass. Moreover, the *exclusive* role of circles in Book I⁴⁰ (which is our focus in this chapter and the next) is to establish properties of the straight lines that serve as the circles’ radii: for instance, that two straight lines are equal, given that both can be seen as radii of the same

³⁹ Classical Greek proofs had a characteristic six-step form that Euclid respects: First is the *enunciation* of the general proposition to be proved, second the *setting-out* which provides the particular figure; the third step, the *specification*, restates the proposition in relation to the particular figure. Thereafter follows the *construction*, which extends the figure (e.g., by drawing additional lines), the *proof* proper, and finally the *conclusion*, which echoes the original enunciation. See Heath’s commentary in Euclid (1956, pp. 129-31).

⁴⁰ Euclid considers circles in their own right in Book III of *Elements*.

circle. In such cases, the circular parts of drawn figures can be thought of, intuitively, as explicit depictions of the rotational trajectories that are available to map these straight lines onto coincidence with one another. For these reasons, it seems appropriate to treat straight lines as *more* basic than circles.

One might also be inclined to regard *points* as the truly ‘basic’ Euclidean objects. Indeed, points might seem to be the obvious choice from a contemporary standpoint, which tends to identify continuous lines with *sets* of (infinitely many) points. This viewpoint, though, is obviously anachronistic as applied to Euclid, and seems unlikely to shed light on our basic *intuitive* understanding of points and lines. There are, however, other reasons we might be inclined to take points as ‘most basic’, some of which stem from consideration of Euclid’s text itself. The point is, for instance, the only object that Euclid defines in entirely negative terms, as “that which has no part”, and it is the only object whose placement may be directly assumed in the *setting-out* steps of his demonstrations, *without* being warranted by any postulate or previous proposition—both indications that Euclid regards the point, in contrast to the line, as too fundamental to be provided a meaningful analysis.

One might be tempted to draw a similar conclusion about points and straight lines as the one drawn above regarding straight lines and circles: Namely, just as the construction of the circle in Postulate 3 can be thought of in terms of the *rotation* of a *straight line*, so the construction in Postulate 1, “to draw a straight line from any point to any point”, can be thought of in terms of the *translation* of a (third) *point*, which takes it from spatial coincidence with the first fixed point onto coincidence with the

second. Now, it does seem intuitively very plausible that curves in general (for our purposes: straight lines and circles) are often thought of in terms of motions applied to points; indeed, this is effectively how they are drawn in ruler-and-compass constructions, which would hardly be possible without moving the (physical) *point* of the pencil across the page. But of course, the trajectory of this movement is always dictated by either compass or straightedge—in either case, by what is a physical approximation to a straight line. In the first case, we hold the ‘straight line’ fixed at one of its endpoints, and thereby enforce its rotation in such a way that the other endpoint is constrained to draw a circle; the circle arises from the dynamics inherent to a straight line that is fixed at one end. In contrast, a straight line does not arise out of the ‘internal’ dynamics of a point—rather, the point-in-translation rather needs to be *externally* constrained by the straightedge, which is needed to ensure that the axis of translation is itself straight. So we cannot get straight lines from points in the same way we can get circles from straight lines.

Now, Euclid’s definition of *straight line* does indeed mention points; his definition of the general notion of *line*, however, does not. This suggests at least that Euclid does not suppose that lines are to be reduced to, or exclusively understood in terms of, points. Indeed, Euclid’s opening three definitions might well give the impression of points and lines as having equal standing as regards fundamentality:

1. A **point** is that which has no part.
2. A **line** is breadthless length.

3. The extremities of a line are points. (Euclid 1956, p. 153; **bold** in original)

The first two are ‘proper’ definitions, with the *definienda* (indicated in bold text) characterized independently of each other; the third definition is ‘improper’, in that it merely describes one central way in which points and lines are related to one another, and thereby serves to further characterize both notions simultaneously, in terms of each other. Definition 3 can thus be read in two different ‘directions’: as characterizing *lines* in reference to their endpoints, or as characterizing *points* as the extremities of lines. One plausible motivation Euclid might have had in offering Definition 3 is to explain how objects of *both* kinds can arise ‘synthetically’ in the construction phase of a proof, in a way that depends on the prior presence of an object of the *other* kind. On the one hand, given the presence of two points, Postulate 1 licenses the construction of a novel straight line, drawn from one point to the other. On the other hand, given the presence of a straight line, Postulate 2 licenses its extension, which yields a new terminal endpoint. In addition, if Postulate 1 is used to draw a new line across an existing one, this again yields a novel point by Definition 3, since the intersection itself is the extremity of each of the four branches that result, when we consider each of the original lines as divided into parts by the line that intersects it. In two different ways, then, points can be ‘constructed’ in the course of a Euclidean demonstration, and indeed, constructed *out of* lines.

While the foregoing discussion does not definitively resolve the question of the relative fundamentality of points and lines, it ought to be sufficient to dispel the

impression that our decision to regard the straight line, and not the point, as the basic object of Euclidean geometry is clearly misguided. In treating (straight) lines as basic, we will not be denying that points can be grasped independently of the lines that terminate in them. In particular, we can allow that points can be visually grasped simply as locations in planar space,⁴¹ relative to a planar object at the current focus of attention (to which a reference system is assigned). For instance, points in the space surrounding an attended square figure can be indicated by small dots drawn in the surrounding region of the page. On our approach, however, we will only ever consider points in relation to the straight lines which are or which could be drawn to them, in the manner captured by Euclid's first postulate.

Before addressing Postulate 1's construction of a straight line, we need to consider how we might specify a perceptual concept for geometrical straight lines, since it is only by engaging such a concept that a figure drawn in accordance with Postulate 1 will be perceived by the subject *as* (depicting) a straight line, in such a manner as to permit visually-based knowledge concerning the geometrical properties of straight lines. In a recent paper, Giaquinto offers a brief comment on perceptual and geometrical concepts for *straight line*, which serves as a useful starting point for this investigation:

[We] start with a perceptual concept of straight line which applies to
apparently straight uninterrupted surface edges or surface marks that have

⁴¹ Recall, however, that we are *not* making the assumption that the planar space depicted by the surface of the page comes with its Euclidean properties 'built-in'.

length but negligible or imperceptible breadth. A line falling under that concept looks perfect if it has no perceptible breadth and its deviations fall below visual acuity. Perception or visual imagination thus supplies us with the material for a concept of a perfectly straight line: it is a line that has the spatial properties that any perfect-looking straight line must appear to have in order to look perfect and perfectly straight. This, I suggest, is our initial geometrical concept of straight line. (2011, 291)

This passage applies to *straight line* the idea that Giaquinto previously defended in relation to *square*: In short, there is a bridge between the vague ‘perceptual’ concept and the more exacting ‘geometrical’ concept, in such a way that the latter can be derived from the former, by restricting its scope of application to instances of *perfect* accordance with the relevant feature descriptions. It is clear from the passage that there are two independent components of the feature description set, corresponding respectively to the properties of *straightness* and *linearity*, and that Giaquinto intends that both properties should be *perfectly* instantiated by the objects to which the (strict) ‘geometrical’ concept *straight line* properly applies. Notably, the first property, *straightness*, is simply assumed in his characterization of the perceptual concept for *straight line*, so there is a complete absence of the kind of informative reduction of a geometrical property to more basic features that was evident in his earlier treatment of the *square*. This may reflect nothing beyond the brief treatment Giaquinto gives to *straight line*—which after all is offered here in the service of making a more general

point about perceptual and geometrical concepts—in contrast to the thoroughgoing analysis of the *square* concept he provides elsewhere. A more intriguing possibility is that Giaquinto regards the perceptual concept for *straightness*⁴² as ‘basic’, in the sense that its associated feature description set (unlike that of *square*) includes just the visually detectable feature of *straightness* itself. It is not fully clear how this idea of basic-level visual features might be cashed out in psychological terms, but one option would be to maintain that there are no more fundamental features, in virtue of which straightness is visually represented, that are themselves accessible to visual awareness at the personal level. Rather, the representation of straightness would be a task exclusively reserved for sub-personal visual processes. This proposal would seem to establish a major difference in epistemological status between, on the one hand, judgments about the properties straight objects possess in virtue of their straightness (e.g., that they can be ‘produced’ in accordance with Euclid’s second postulate), and on the other, corresponding judgments about square objects. For on Giaquinto’s account, the latter judgments can appeal to properties that square objects are *conceptually* guaranteed to possess in virtue of being square (in his example, reflection symmetry about the diagonals). On the proposal that straightness is a basic-level visual feature, there would be no comparable properties to which judgments about straight objects could rationally appeal, since the perceptual concept *straight* would merely point to a nonconceptual simple. The main task in Section 7 will be to

⁴² Here I am taking some interpretive license in assuming that Giaquinto’s view supports the postulation of an independent concept of *straightness*, which might apply to lines as well as non-lines (e.g., bars or strips).

motivate the contrary proposal that straightness need not be regarded as basic in the above sense, but rather can itself be visually represented in terms of certain symmetries, and in particular, the *motions* associated with those symmetries.⁴³

6 Linearity in general

Before proceeding to a consideration of how perceived straightness might be grounded in perceived symmetry, it is worth briefly remarking on the other component feature of the description set for Giaquinto's *straight line* concept: namely, *linearity* itself. What is it to perceive a *drawn* line as a perfect, geometrical line (or as depicting one)? The property of *perfect linearity* poses a distinctive challenge for Giaquinto's attempt to derive precise 'geometrical' concepts from vague 'perceptual' ones. Recall that Giaquinto's general strategy here appeals to lower limits on perceptual discrimination, which allow that drawn figures may in certain respects be perceptually indistinguishable from perfect instantiations of geometrical properties. Even in the absence of such a close approximation to the geometrical ideal, a visibly imperfect figure might still yield a perceptual experience as of its perfect geometrical counterpart, by triggering the visual imagination of an

⁴³ This is not to deny that it might *also* be true that straightness is often visually represented as a 'nonconceptual simple', in virtue of the operation of exclusively sub-personal visual processes. The claim being advanced here is just that it is *possible* to visually represent something as straight in virtue of other properties that are accessible to personal-level visual understanding. A similar point could be made about reflection symmetry itself: That is, there are very plausibly sub-personal visual processes that issue directly in consciousness visual experiences of symmetry. But this in no way precludes the intuitively 'thicker' representation of reflection symmetry *as* a property an object possesses in virtue of being invariant under the *motion* of reflection, which can itself be rehearsed in visual imagination.

apparently perfect instance. In this way, Giaquinto is able to explain how visual experiences of imperfect drawn figures can engage concepts that apply to properly geometrical objects or properties. This strategy applies quite straightforwardly to his examples of *square* and *straight*, given that it is intuitively plausible that subjects can perceptually entertain, in either veridical experience or in visual imagination, figures that are perceptually indistinguishable from perfectly square or perfectly straight figures. The same strategy confronts immediate difficulties, however, when it comes to the property of *geometrical linearity* that Euclid points to in his definition of a *line* as “breadthless length”. While we can readily imagine having a visual experience as of a perfectly *straight* line (or linear strip), it is far from clear what it would be to have a visual experience as of a perfectly *thin* line—one altogether lacking in breadth altogether. In the passage quoted above, Giaquinto suggests that our initial, *imprecise* perceptual concept of *line* applies to “uninterrupted surface edges or surface marks that have length but *negligible or imperceptible* breadth” (2011, 291, emphasis added). It is the second part of the disjunction that is puzzling here. We can grant that a subject might have a visual experience as of a roughly linear object whose perceptible breadth is not significant enough to ‘count’ in the given context, but it is rather mysterious how a subject might have a visual experience as of a line with *imperceptible* breadth, since the very supposition that the line itself is visible seems to entail that it should have at least *some* perceptible breadth. This consideration makes Giaquinto’s assertion that a line “looks perfect if it has no perceptible breadth” difficult to accept, and seems to require his claim that “perception or visual

imagination... supplies us with the material for a concept of a perfectly straight line” to be at least significantly qualified, given that he understands this concept to apply to lines that are not only perfectly straight but also perfectly linear. For there would seem to be no possible visual experience that represents a line as entirely without breadth.

If this is right, and the objects of geometrical knowledge are, in this manner, imperceptible in principle, one might be tempted to conclude that the sort of thoroughly perception-based account of basic geometrical knowledge being defended here cannot possibly succeed. Since providing a complete response to this challenge would take us beyond the scope of the present chapter, only a sketch of such a response will be offered here. In brief, the idea is to pursue the *first* part of Giaquinto’s disjunction, which ties the perceptual concept of *line* to a visual experience of “negligible” rather than “imperceptible” breadth. Provided that we can give a sufficiently precise account of the sense in which the drawn lines comprising Euclidean diagrams are perceived as having *negligible* breadth, that will in fact suffice as an account of how those diagrams can provide a perception-like acquaintance with the Euclidean objects themselves. In this connection, it is useful to take seriously the identification of Euclidean diagrams as *pictorial* representations of a certain kind. It is now a familiar idea in the philosophy of art that the aesthetic properties of an artwork are determined not only by the full set of (non-aesthetic) perceptual features of the work in question, but also by which of these properties are “standard” as opposed to “variable” relative to the conventions of the general

category (that is, the artistic medium or genre) to which the work belongs (Walton 1970).⁴⁴ For example, the flatness and rectangularity of the picture plane, being standard properties for the category of paintings, do not figure thematically in the overall aesthetic effect generated by a given painting: Since paintings are, as a matter of course, *expected* to be flat and rectangular, the presence of such features fails to strike the observer at all, instead serving to implicitly direct aesthetic attention to the *variable* features of the work, which concern *how* paint is arranged on the flat, rectangular canvas. This implicit recognition of certain properties as standard for the relevant category plays an important role, in Walton's view, in determining the depictive contents of pictorial works. This is because it is only by, effectively, *looking past* the standard features of flatness and rectangularity that the viewer is able to discern the *resemblance* between the picture and the object(s) depicted that Walton regards as a necessary condition on the depiction relation itself:

The properties of a work which are standard for us are ordinarily irrelevant to what we take it to look like or resemble in the relevant sense, and hence to what we take it to depict or represent. The properties of a portrait which make it *so* different from, so easily distinguishable from, a person—such as its flatness and its *painted* look—are standard for us. Hence these properties just do not count with regard to what (or whom) it looks like. It is only the

⁴⁴ Walton also mentions a third sort of category-relative property—the “contra-standard”—which comprises those properties that *subvert* the conventions of the artistic medium or genre to which the work in question nonetheless belongs.

properties which are variable for us, the colors and shapes on the work's surface, that make it look to us like what it does. And these are the only ones which are taken as relevant in determining what (if anything) the work represents. (Walton 1970, pp. 344-5)

Walton's view here can help us make sense of the idea that the drawn lines of Euclidean diagrams might issue in a visual experience as of *negligible* breadth, in a way that provides a kind of perceptual acquaintance with genuine Euclidean lines, *via* their role as the depictive contents of the drawn lines—without any need for a problematic appeal to a visual experience of *imperceptible* breadth. In this connection, it is worth emphasizing several features of Walton's account. First, the sense in which the standard properties of a pictorial work “just do not count” in determinations of resemblance and depictive content in no way depends on those properties lying beneath or even near the threshold of perceptual discriminability. On the contrary, their very perceptual obviousness typically plays a key role in achieving their *exclusion* from the experienced depictive content, by facilitating their swift assignment to the class of standard properties. Second, the fact that these properties may be in principle incompatible with certain properties of the depicted objects places no limitations on the *degree* of the apparent resemblance between the objects and the picture that represents them, given the manner in which category-relative understanding of the picture delimits the range of features relevant for resemblance. As such, a picture can be properly said to be a “perfect likeness” of an object,

notwithstanding the obvious discrepancies when it comes to those features that are standard for the relevant category (Walton 1970, 344). Third, the mediation of pictorial experience by category-relative determinations of certain features as standard or variable operates on a *perceptual* level, affecting not merely what a picture is understood to represent, but rather, in a more basic sense, what the picture *looks* like.

Provided that we do take seriously the identification of Euclidean diagrams with a certain category of pictorial representations, we can apply Walton's distinction between standard and variable properties, and can thereby regard the visible breadth of the *drawn* lines as a standard property of depictions belonging to this category. On this view, the visible breadth of the drawn lines should not, in itself, preclude their serving as the basis for a pictorial experience in which the perceiver *sees* in the diagram something truly breadthless. Moreover, the Waltonian perspective on Euclid suggests an appealingly charitable way of interpreting Euclid's much-maligned definition of *line* as "breadthless length" (among others), which in this light can be read *alongside* the diagram itself, as an instruction concerning the conventional depictive character of the latter. By *stipulating* breadthlessness of the drawn lines, which do of course have visible breadth, Euclid can be understood to be laying down the convention that line-breadth is a standard as opposed to variable property of diagrams belonging to the relevant category. In this manner, the definition, in conjunction with demonstrative reference to the perceived diagram, is capable of

conveying the idea of a Euclidean line rather precisely: To a first approximation, it is an object like *that* one, where the breadth is understood not to count.⁴⁵


The proposal that a *roughly* linear strip drawn on the picture surface can in this manner depict something perfectly breadthless might seem like a dubious one, given how different it is from ordinary cases of depiction, in which the depicted object is something that is, in principle, perceptible. If there can be no such thing as the *appearance* of a true Euclidean line, it is difficult to make sense of the idea that category-relative perception results in the line *looking* breadthless, even in the very special sense in which a cartoon caricature, say, can *look just like* the subject it depicts. In that case, even granting that the diagram *represents* a geometrical line, the *mode* of representation might be viewed not so much as pictorial, but rather as resting on a stipulated identification with an object that can be grasped ‘by thought alone’, as one might put the point. Here again, however, we can appeal to dynamic aspectual imagery in order to explain the depiction of breadthlessness. For it is a general feature of the geometrical diagrams at issue that they are, in a certain specific sense, pictorially ambiguous.

Figure 4.2, for instance, could be seen either as a line (or more precisely, a linear strip) ‘drawn inside’ (or perhaps ‘placed onto’) a disk, or else as a pair of half-disks that fit together perfectly to determine the shape of a whole disk. On the former interpretation, the line is apprehended as, roughly, an *object*—that is, something that could be moved or handled directly—and in that capacity, as something that takes up

⁴⁵ Shabel (2003, 13-17) offers a somewhat similar interpretation of Euclid’s definitions.

at least a bit of space on the disk's surface. On the latter interpretation, it is only the half-disks that are apprehended as objects in this sense; the line has assumed a purely negative existence, now appearing merely as the border shared by the half-disks where they meet one another. On this second pictorial interpretation, the line is indeed perceived as breadthless; the breadth of the drawn line does not *count* pictorially, because the only spatial objects being depicted are the half-disks whose mutual border it serves to indicate. The identification of a *line* with this negative role is made explicit in Euclid's Definition 6, which reads: "The extremities of a surface are lines". The problematic idea of there being visual experiences as of perfectly breadthless lines has been replaced by the entirely unproblematic idea of there being visual experiences as of surface objects with perfectly determinate boundaries.

Figure 4.2: Circle with diameter



Of course, given the fundamental role of (straight) lines in Euclidean constructions, it will not suffice to regard lines *exclusively* as, in this manner, negative *non-objects*. For Euclid's constructions are performed precisely by treating lines as objects in their own right—by drawing or placing them in such a way as to allow the

‘synthetic’ emergence of rectilinear angles and plane figures, which appear as soon as one entertains the alternative pictorial interpretation.⁴⁶ Lines thus appear to have a special status in Euclid’s system: Effectively, they straddle two alternative pictorial vantage points or ways of seeing the diagram, inheriting their object-hood from one, and their breadthlessness from the other. In other words, lines as geometrical objects can be regarded as constituted by a fundamental convention governing the visual understanding of Euclidean diagrams, which permits the *integration* of the two alternative interpretations just illustrated in relation to Figure 4.2. This convention can be plausibly regarded as implicitly set out in Euclid’s definitions, by means of the *sui generis* notion of ‘containment’ that relates lines to both angles and plane figures, for example, in Definitions 9, 14, 15, and 19:

9. And when the lines *containing* the angle are straight, the angle is called **rectilinear**.

14. A **figure** is that which is *contained by* any boundary or boundaries.

15. A **circle** is a plane figure *contained by* one line such that...

19. **Rectilinear figures** are those which are *contained by* straight lines...

(Euclid 1956, pp. 153-4; *italics* mine, **bold** in original)

⁴⁶ Cf. Kant’s famous passage in which he describes the geometer performing the construction for Euclid’s proof of Proposition I.32, the triangle angle-sum theorem: “...he *extends one side* of his triangle, and *obtains two adjacent angles* that together are equal to two right ones. Now he divides the external one of these by *drawing a line parallel* to the opposite side of the triangle, and *sees that here there arises an external adjacent angle* which is equal to an internal one, etc.” (Kant 1998, A716/B744, my emphasis)

It is worth remarking on how different this *containment* relation is from the *part-whole* relation considered previously in connection with the common notions. The latter relation applies only between objects of the same type, that is, between two lines, angles, or plane figures. For instance, the half-disks in Figure 4.2 (on the second pictorial interpretation) are both parts of the whole disk; taken together, they *coincide* with and hence are collectively equal to the whole. Containment, on the other hand, obtains exclusively between objects of different types: Specifically, it is lines that contain both angles and plane figures. The striking fact about the containment relation is that it purports to relate coplanar geometrical *objects* of distinct types, but this relation between coplanar objects cannot be observed within a single pictorial interpretation of the diagram, as we just observed in the case of Figure 4.2: On one interpretation, the line-as-object appears to lie either ‘on top of’ the complete disk or ‘within’ the circumscribed space, while on the other interpretation, the line inevitably becomes something negative—a mere boundary—as the half-disks assert themselves as objects. By *defining* angles and plane figures in terms of their containment by lines, Euclid is introducing a relation that spans *across* pictorial interpretations, in what might be termed the ‘foundational equivocation’ of Euclidean plane geometry. From this perspective, there is a fundamental duality in the very nature of Euclidean lines: They serve both as positive *objects* that can be directly manipulated—“set up on” other lines (Definition 10), “drawn through” circles (Definition 17), and so forth—and also as the (negative) delimiting boundaries of

angles and *plane figures*, which then themselves take on the ‘positive’ character of objects: that of being subject to rigid motion, and hence able to be judged *congruent* to other angles or plane figures, for instance.

If this is the right view to take, then it follows that visual experiences of geometrical lines are *inherently* pictorial experiences of a distinctive, ‘complex’ character, involving the integration of alternative pictorial interpretations of a single drawn figure. While this may seem initially like a rather obscure notion, it has precedent, once again, in the aesthetics of pictorial perception. In particular, Brown (2010) argues at length for the central and ubiquitous role, in the aesthetic appreciation of pictures, of what he terms “separation seeing-in”, so called because it involves seeing-in-the-picture a pictorial content that is *separate* from the official subject matter. In a pencil sketch of a human figure, for instance, one such “separation subject” might be a human body that is graphite-colored, partly translucent, and with skin decorated by contours and cross-hatchings that are well defined in some places, while becoming indeterminate in others. Clearly none of these visual features could be properly ascribed to the primary or ‘official’ pictorial subject, which is after all a human figure, presumably with an ordinary appearance. When we see this *official* content in the picture, features of the *sketch* itself, such as its monochromatic coloration and use of cross-hatching, become *discounted* in the manner proposed by Walton, as we saw earlier. Brown’s point is not at all to deny the primacy of the Waltonian seeing-the-official-content in the picture, but he resists the usual assumption that the features that are discounted in such seeing-in can inform

our pictorial experience only *qua* surface features of the sketch itself. Rather, some of them become (as we might say) ‘inherited’ by the alternative subject matters of separation seeing-in, which are entertained concurrently, alongside the official depictive content. As he argues, not only can we “maintain a lively awareness of both in any well-conducted course of pictorial seeing,” but in general, the “full depictive character of a painting is revealed only by the ensemble of seeings-in it offers” (Brown 2010, 213; 217).

Brown’s account of separation seeing-in, then, seems to provide at least a good initial basis for making sense of the dual character of Euclidean lines articulated above. A geometrical line must be visually experienced, not as a single depictive *content*, but rather within the complex depictive *character* that arises from the integrated awareness of distinct depictive contents. To visually grasp a true geometrical line will thereby require, in Brown’s terms, that we “maintain a lively awareness” of *both* pictorial interpretations: one on which the line inherits its visible breadth from the *drawn* line that depicts it, and another on which this breadth is discounted in the manner articulated by Walton. This makes the visual experience of a geometrical line a rather eccentric sort of visual experience, but nonetheless one that arises relatively straightforwardly out of our core postulation of dynamic aspectual imagery. To a rough approximation, then, we can say that a perceptual concept for a geometrical line consists in something like the following bidirectional belief-forming disposition: When one trusts one’s visual experience, one will be apt to perceptually

judge that x is a *line* if and only if one has an integrated visual experience that represents x as both a *linear strip* and as a *determinate surface border*.

7 Straightness

Having thereby indicated how a subject could have a visual experience as of a geometrical line, we are ready to consider the question of how the geometrical property of *straightness* might be represented in visual experience. It will simplify matters somewhat if we restrict our attention just to the straightness of *lines*, given that we will then be able to approach the question of x 's straightness under the assumption that x has constant breadth.⁴⁷ Now, it seems plausible that straightness is represented in early visual processing, perhaps in multiple, convergent ways. This seems likely, given that quite basic visual features, like *direction*, *orientation*, *distance*, etc., effectively presuppose *straightness*. So we should expect that there is a thoroughly nonconceptual representation of straightness, which is encountered in visual experience simply as 'given'. As we noted earlier, however, if we specify our perceptual concept for the *geometrical* property of straightness by appealing to straightness *qua* 'basic-level' visual feature in this sense, this will place epistemological limitations on the justificatory character of any judgments *about* the properties of straight lines that are formed by employing this perceptual concept. For

⁴⁷ This is because a *linear strip* is assumed to have constant breadth. Note that if we can specify a perceptual concept for *straight line* in this manner, it should be a trivial matter to derive from this a broader concept of *straightness*, simply by building in the assumption of constant breadth directly. This doesn't appear to limit the generality of the concept *straight*, which intuitively, applies only to things with either *zero* breadth (like edges and contours) or else *constant positive* breadth.

if we appeal only to straightness as it is directly ‘given’ in perceptual experience, we will have no conceptual knowledge about straightness as a geometrical property.

Therefore, our task is to specify a conceptually ‘thicker’ means of visually understanding straightness, preferably in terms of dynamic imagery.

The common definition of a straight line as ‘the shortest distance between two points’ will not do for our purposes, because if we simply *presuppose* distance, we will confront essentially the same problem as if we took straightness as a basic visual feature. While it is no doubt true that we have a direct visual experience of distance, we lack a conceptually ‘thick’ way of understanding distance—it is seemingly directly ‘given’ in visual experience. Not only would presupposing distance in this way limit the justificatory status of our judgments about distance in virtue of its ‘thin’, nonconceptual character, but it would commit us to the assumption that visual space is Euclidean, and as we noted earlier, there is considerable psychophysical evidence that counts against this claim. A better strategy is to begin by specifying a ‘thick’ perceptual concept for straight lines. Given that straight lines are subject to rigid motions, we will thereby be in a position to apply the perceptual concept for congruence discussed in earlier, and on this basis, we will be able to derive a concept for equal distance in terms of congruent straight lines. This reflects a basic commitment of our approach, which is to explain judgments about Euclidean properties in the first place with reference to constructed geometrical *objects*, and only derivatively to the ‘space’ they occupy.

A more promising avenue proceeds by way of Euclid's own definition: "A **straight line** is a line which lies evenly with the points on itself." This formulation is notoriously hard to interpret, as discussed in Heath's notes on the definition (Euclid 1956, pp. 165-9). Heath provides a careful analysis of the wording and grammar of the Greek text, and argues that the crucial term 'evenly' should be understood as roughly equivalent to 'indifferently' or 'without bias' (among other synonyms); he concludes:

While the language is thus seen to be hopelessly obscure, we can safely say that the sort of idea Euclid wished to express was that of a line which presents the same shape at and relatively to all points on it, without any irregular or unsymmetrical feature distinguishing one part or side of it from another.
(Euclid 1956, p. 167)

This exposition of Euclid's definition gives us just the kind of conceptually 'thick' characterization of a straight line that we are seeking, for it specifies a precise set of visually detectable features whose recognition can serve to provide content to a perceptual judgment to the effect that a perceived line is straight. In particular, both sides of the line need to be such that their visual appearance is invariant as one directs visual attention continuously along its length in a process of visual scanning,⁴⁸ and in

⁴⁸ This scanning-based method is perhaps the most straightforward means of verifying continuous symmetry along the length of the line (as seen from one side); alternatively, one could direct attention to a proper part or segment of the line, and

addition, the line needs to appear the same regardless of which side it is being viewed from.⁴⁹ This characterization, which is in terms of perceived invariance under changes in subjective vantage point (as one visually scans, or regards the line from alternative sides, that is) should be understood to fall within the scope of our account, given that we are understanding transformational imagery in a broad sense, as encompassing not only imagery corresponding to potential motions of attended objects, but also to potential shifts in the vantage point of visual attention (scanning being the most obvious example). As such, we could use the above characterization as a basis for a perceptual concept of straight lines. Alternatively, we could frame our perceptual concept in terms of object motion, which is the approach we will take here.

It is clear upon consideration that if a line possesses reflection symmetry, it will have both of the properties required for straightness, for there will be in this case a mapping between the two opposite sides, which will apply across their entire lengths. This suggests that an appeal to reflection could provide the basis for a perceptual concept of straight lines. One might wonder whether such an approach would introduce a problematic circularity, given that reflection symmetry is observed relative to an axis that must itself be straight. If we understand the motion of

imagine this segment rigidly moving along the length, while the whole line remains intact.

⁴⁹ Since the first requirement is satisfied by any line of constant curvature (including not only straight lines but also circles and circular arcs), we need this second requirement to capture the more specific property of *zero* curvature or straightness.

reflection merely as a ‘folding over’, however, we won’t face any such problem.⁵⁰

Suppose one imagines the line in question being lifted out of the plane, ‘flipped’ in three-dimensional space, and set back into the plane, while ensuring that its endpoints remain in their original locations.⁵¹ The question is then whether the line coincides with its original position⁵² under the imagined transformation; if so, it can be judged *straight*. Importantly, the imagined transformation of ‘folding over’ can be entertained on either of the two pictorial interpretations whose integration was held just above to constitute the experience of geometrical linearity in general. On one interpretation, it is the *linear strip* that serves as the object of motion, and which, supposing it to be straight, is seen to be self-symmetrical, being mapped by the transformation onto a perfect coincidence with itself in its original position. On the alternative interpretation, it is the two planar surfaces that open out from either side of the line that are seen to be mutually symmetrical; assuming straightness, when one is

⁵⁰ Intuitively, consider the way that folding a piece of paper haphazardly (that is, without bothering to align any straight edges) is able to *produce* a straight fold.

⁵¹ This is to ensure that the line is not rotated while being reflected, in such a way that its endpoints become mapped onto their opposites. That composition of reflection with rotation would allow a circular arc, for instance, to be mapped onto coincidence with itself, in spite of not being straight. In case it seems question-begging to assume that *both* endpoints can remain in place under the motion of reflection, it is sufficient to require that only one endpoint remain in place—the question of straightness will then turn on whether any rigid motion that respects *that* constraint (which will be limited to rotations about the fixed point) is capable of bringing the line into coincidence with its original position.

⁵² Of course, given that the *drawn* line remains on the page while one imaginatively rehearses it being moved, one simply needs to judge whether its *potential* position under the imagined motion corresponds to its actual, present position. In this sense, the experience is somewhat like that of manipulating a ‘copy’ or ‘duplicate’ that starts off in a position of coincidence with the ‘original’.

‘folded over onto’ the other, the boundaries of the two will be seen to coincide perfectly.

In this manner, the visual imagination of reflection can serve to ground a perceptual concept for geometrical straight lines, without taking straightness itself as a ‘basic-level’ visual feature. For a subject *S* to possess this perceptual concept is roughly for *S* to have the following pair of belief-forming dispositions: (1) If *S* believes that *x* is a line, and has (as well as trusts) an experience that consists in visualizing *x* being reflected into a state of coincidence with itself, then *S* will believe that *x* is a straight line; (2) If *S* believes that *x* is a straight line, then *S* will believe that *x* is a line and that if *x* were to be reflected so that its endpoints coincide with their original positions, then *x* as a whole will coincide with its original position under the imagined transformation. We have, then, succeeded in showing that the resources of the dynamic imagery account are able to provide the basis for a perceptual concept of geometrical straight lines. While this symmetry-based perceptual concept almost certainly does not exhaust the ways that straight lines are represented as such in visual experience, it does provide a plausible account of the way a perceived line’s *seeming* straightness can be ‘fulfilled’ by visual experience, rather than being *merely* present as part of its content. When one inspects a drawn figure of an apparently straight line with an eye to straightness, the general phenomenological experience is of having a quasi-observational acquaintance with those spatial features *in virtue of which* the line qualifies as straight, which gives the perceptual experience of apparent straightness a

‘transparent’ (as opposed to ‘opaque’) character.⁵³ Even if one is in no position to articulate in any informative terms what it is about the figure that gives rise to an impression of straightness, the experience of *seeming straight* nonetheless includes as part of its phenomenology that one is somehow visually acquainted with such straight-making features. The present account supports this phenomenological impression, by explaining (one case of) the visual experience of straightness as consisting in the active recognition of reflection symmetry by means of the imaginative rehearsal of the motion of reflection.

In the discussion to follow, we consider whether the reflection symmetry that forms the foundation of this experience of a straight line can serve to underwrite the recognition of further properties of straight lines, in a way that yields similarly ‘transparent’ or visually ‘fulfilled’ intuitive judgments. In particular, our focus will be on the question of whether such a perceptual concept might underlie intuitive belief in Euclid’s first three postulates.

⁵³ In that respect, the phenomenological experience of apparent straightness is quite *unlike* the initial (erroneous) judgment made by the slave at the *outset* of the geometrical demonstration in *Meno*. As we noted in Chapter 1, this latter judgment captures the kind of ‘intuition’ that is plausibly ascribed to an encapsulated, heuristic-based process that merely delivers a judgment without making the underlying reasons for the judgment accessible to consciousness or otherwise available for rational consideration. In the experience of apparent straightness, in contrast, we have a judgment that still qualifies as ‘intuitive’ insofar as it arises spontaneously and without mediation by conscious deliberation, but which is accompanied by a conscious visual awareness of the features in virtue of which the appraisal is an apt one.

8 From straight lines to construction postulates

The propositions proved in *Elements* can be categorized into the *problems*, which establish a procedure by which a figure satisfying certain properties can be constructed, and the *theorems*, which establish that some property holds for all figures belonging to a certain class (where existence is ultimately to be justified by providing a general construction procedure).⁵⁴ A parallel distinction can be applied to Euclid's postulates themselves: The first three serve to justify the basic constructions of Euclid's plane geometry—the drawing and extending of straight lines, and the drawing of circles—while the latter two serve justify basic assumptions about the general properties of right angles and (in effect) parallel lines.⁵⁵ In the next several sections we will consider the three construction postulates, whose focus is primarily on the *local* spatial properties of Euclidean figures.⁵⁶ The fourth and fifth postulates,

⁵⁴ The choice to frame a given proposition as either a problem or a theorem appears to be arbitrary to at least a significant degree. Consider, in this regard, the Pythagorean theorem, which appears as Euclid's Proposition I.47: "In right-angled triangles the square on the side subtending the right angle is equal to the squares on the sides containing the right angle" (Euclid 1956, p. 349). While Euclid thereby opts to frame what we might neutrally describe as 'the Pythagorean relation' (that between the squares' areas) as a *theorem*, he could alternatively have taken this relation as the topic for a construction problem, *viz.*: "To construct a square equal in area to two given squares." For further discussion of the theorem/problem distinction, see Heath's notes on the topic (Euclid 1956, pp. 124-9).

⁵⁵ Strictly speaking, this characterization ascribes the wrong content to Euclid's fifth postulate, which in spite of being commonly called 'the parallel postulate', makes no reference to parallel lines, but instead to lines whose internal angles with a transversal are equal to less than two right angles. The postulate tells us that such lines are *not* parallel.

⁵⁶ This is not to say that the construction postulates are without *any* implications concerning the global structure of Euclidean space. On the contrary, the second

which respectively impose the global properties of *homogeneity* and *flatness* on Euclidean planar space, will be considered thereafter. Throughout this discussion, our aim will be to explain how intuitive belief in the postulates can arise out of the employment of dynamic imagery in the visual understanding of straight lines.

We begin with Euclid's first and second postulates, which respectively justify the drawing and extending of straight lines. Again, these read as follows:

1. To draw a straight line from any point to any point.
2. To produce a finite straight line continuously in a straight line.

In practical terms, these postulates characterize what can be drawn by placing an (arbitrarily long) straightedge in alignment with two given points, or with an existing straight line. A few initial comments are in order. First, it is an unstated assumption of Euclid's system that points exist, and that they can be placed in arbitrary locations in planar space; accordingly, the first postulate can be taken to justify the construction

postulate implies that the plane is unbounded. Perhaps less obviously, it is actually the *first* postulate (on the standard interpretation according to which it implies that two points determine a *unique* straight line) that rules out the elliptical model (in essence, the surface of the sphere) as a possible interpretation of Euclid's geometry, and thereby implies that Euclidean space has either zero (Euclidean) or negative (hyperbolic) curvature. While Euclid's fifth postulate does rule out hyperbolic geometry as an interpretation, it is, as stated, *true* in elliptical geometry. Thus, contrary to what is widely claimed, the fifth postulate is *not* always rejected by (even the classical) non-Euclidean geometries, nor is it sufficient to establish the flatness of the Euclidean plane; the first postulate is needed in order to rule out the elliptical model, with its positively curved space.

of an *arbitrary* straight line.⁵⁷ The fact that any straight line whatsoever is constructible by the first postulate does not render the second postulate redundant, however, for even the first postulate together with the unstated assumption that points can be arbitrarily placed does not justify the existence of a point precisely collinear with a given straight line; it is the second postulate that permits the *construction* of such a point, as the new extremity of the ‘produced’ line. Relatedly, the second postulate implies that a straight line divides the plane as a whole, permitting Euclid to speak unambiguously about a point placed arbitrarily within the space on a given *side* of a given straight line.

Second, as stressed by Heath (Euclid 1956, pp. 195-9), both postulates should be read as implying *uniqueness* of the straight lines drawn or produced on their warrant. In the case of the first postulate, we can understand uniqueness as requiring that *any* straight line drawn between the extremities of a given straight line will coincide with the latter. This assumption is needed for the proof of Proposition I.4 (which establishes the side-angle-side congruence criterion for triangles), in which Euclid justifies the coincidence of the triangles’ bases with the parenthetical remark that if they do not coincide, “two straight lines will enclose a space: which is impossible” (1956, p. 248). In the case of the second postulate, uniqueness intuitively requires that “the straight line can only be produced *in one way* at either end” (Heath in Euclid 1956, p 196). One way to make this idea more precise is in negative terms that mirror the statement that *two distinct straight lines cannot enclose a space*: In

⁵⁷ The proof of Proposition I.2 thus begins: “Let A be the given point, and BC the given straight line” (Euclid 1956, p. 244).

this case, the formulation would be that *two distinct straight lines cannot share a part*. Alternatively, we could formulate the requirement positively: *If a straight line is produced twice at one end, the produced parts will coincide until an extremity is reached*.

The final comment concerns the assumption, made explicit in the second postulate but already implied by the first, that lines can be drawn of *indefinite* length. Notwithstanding Euclid's apparent appeal to a "given infinite straight line" in the statement of Proposition I.12 (1956, p. 270), it is not actually necessary that we should read the second postulate as warranting the construction of an infinite straight line that can be assumed as "given"—nor would this be desirable, in light of the remoteness of such an 'actual infinity' from any possible visual experience. Rather, the second postulate need only be taken to underwrite the 'potential infinity' that consists in the possibility of extending *any* line, however long, by some finite amount. To summarize the content jointly conveyed by the first and second postulates, then, they establish: (1) that two points determine a *unique* straight line, (2) that a straight line can be produced from either end *indefinitely* and *in only one way*, and (3) that a particular drawn line is (therefore) adequate to represent an *arbitrary* straight line. This provides us with a sufficiently precise target for the intuitive knowledge our account is intended to capture in relation to the first two postulates.

While (1) and (2) should ultimately be read as applying, respectively, to *any* two points and to *any* straight line, it is convenient to begin by considering these claims merely in reference to a particular given instance, deferring the question of

generality until it can be addressed directly in relation to (3). Consider, then, a diagram depicting a particular pair of points, which (1) claims will determine a unique straight line. Recall from the earlier that a sufficiently thin *linear strip* can serve conventionally to depict a geometrical line, because it can be seen, in accordance with Definition 6, as the extremity of a surface region. Similarly, points can be conventionally depicted by sufficiently small (circular) dots, which can be seen as (potential) “extremities of lines” in accordance with Euclid’s Definition 3. A dot, that is, can readily be seen as something to which, or out of which, lines may be drawn.⁵⁸ Assuming our two points to be positioned along a horizontal axis with respect to the picture plane, then, the relevant diagram will look like Figure 4.3.

⁵⁸ It is only by perceiving drawn points in connection to lines that one can attain any visual appreciation of their being ‘perfectly dimensionless’—the key geometrical property that Euclid captures negatively, in Definition 1, in terms of the inapplicability of the part/whole relation to points. This lack of dimensionality can be thought of as being ‘downward inherited’ from surfaces to lines to points, *via* the containment-by-extremity relation discussed earlier. Note that points, being the extremities of lines, are not *parts* of lines in Euclid’s sense of ‘part’; that is, they do not ‘share space’ in the plane—something Euclid indicates by figuratively describing points as being *on* lines. Similarly, the line drawn through the circle in Figure 2, when interpreted as a linear strip with some degree of breadth, is not seen as sharing planar space with the circle, but rather as being ‘placed onto’ the underlying, complete circle. It is only when being viewed *as* extremities that points (or lines) are seen as properly coplanar with lines (or surfaces).

Figure 4.3: Pair of points



If we understand this image as a depiction of two geometrical points, can we see that the points determine a unique straight line? To begin with, it is relatively easy to explain how we can see that *at least one* straight line can indeed be drawn between these points. For even in the absence of any suggestive framing, the cue provided by the presence of dot-like elements in an otherwise blank space is often sufficient to provoke a spontaneous visual impression of a continuous contour, through a kind of amodal completion; there is, moreover, psychophysical evidence that the visual system implicitly expects contours to continue in a direction that minimally departs from the local tangent direction (with smaller divergence from the tangent direction being treated as more probable), creating a natural tendency for the visual system to represent contours as *straight*, when that possibility is consistent with the stimuli (Feldman 1997, Elder and Goldberg 2002, Yuille et al. 2004). In light of these features of visual contour detection, it is plausible that the subject perceiving an image like Figure 4.3 will be apt to have a visual experience in which an apparently straight contour seems immediately to ‘pop out’. Since contour completion is in this

case amodal, however, the experience won't involve a visual *illusion* as of a contour (the sort of thing that occurs in the Kanisza triangle illusion, for instance). The phenomenology, then, is not of *seeming to see* a drawn line; the contour is experienced rather as something closer to a feature of the structural integration of the two-dot figure, which provides an immediate awareness of where the line *would* be seen if it were to be 'drawn in'. In this way, the contour can be thought of as playing a role somewhat analogous to that of the straightedge in ruler-and-compass constructions. While the amodally-completed contour doesn't itself have the determinate appearance of a drawn line, this appearance can be fulfilled in visual imagination, yielding something like Figure 4.4.

Figure 4.4: Straight line drawn between points



While the straightness of this imagined line is in the first place experienced nonconceptually, this experience can readily issue in an augmented one that engages the perceptual concept for *straight line* described earlier, provided that the subject—at least tacitly—appreciates that the drawn line would self-coincide under the relevant

reflection. We have, then, an explanation for the perceiver's intuition that at least *one* straight line can be drawn to connect the points depicted by Figure 4.3.

In order to see that this straight line is *unique*, it suffices to recognize that any line that connects the given points, but which does not coincide with the line in Figure 4.4, is asymmetrical, and hence cannot be straight, given the conceptual link between straightness and reflection symmetry established perviously. Consider first the particular case illustrated in Figure 4.5, with the candidate alternative straight line placed above our original straight line.

Figure 4.5: Step one of uniqueness demonstration



Of course it is immediately apparent to us that the top line is not symmetrical in the manner necessary to qualify as straight. If we left the original in place, and ‘flipped’ a duplicate while keeping it fixed at the endpoints, the duplicate would clearly fail to coincide with the original. The result would look somewhat like Figure 4.6.

Figure 4.6: Step two of uniqueness demonstration



If this judgment of asymmetry were made merely on the basis of ‘eyeballing’ the figure, however, it would rest on dubious epistemological grounds, for in that case it would depend on an auxiliary *empirical* assumption concerning the veridicality of our visual experience, one that is undermined considerably by our evident susceptibility to visual-geometrical illusions of various kinds. It would, in any case, be question-begging to judge that the top line is not straight, on the grounds that it can be seen to be asymmetrical due to its *shape*. Fortunately, we need not appeal to the shape of the top line to see that it lacks the requisite symmetry for straightness, for it is sufficient to attend in the right way to the pictorial ambiguity associated with Euclid’s fundamental notion of spatial containment. In particular, in addition to being able to see Figure 4.5 as a pair of *linear strips* that share a pair of endpoints, we can alternatively see in it the *plane figure* contained by the top and bottom lines. And since we *already* know—as it were, by hypothesis—that the bottom line is straight, we know that it enforces reflection symmetry on the spatial regions that open out from it on either side. We can thereby appeal to the assumption that the bottom line is straight to warrant the reflection of the enclosed plane figure above it, taking the line

itself as the axis of reflection. This gives us the same Figure 4.6, now seen as a pair of plane figures symmetrical about the middle line that forms their shared border. Since the enclosed figures are symmetrical to each other, so are the boundaries that contain them, including in particular the pair of top and bottom lines. And since these lines *do* coincide at their endpoints, the perceptual concept for straight lines implies that they must themselves coincide if they are indeed straight. We can see in Figure 4.6 that they do not coincide, because they enclose the larger plane figure comprised of the two symmetrical halves. So we can indeed see that the top line is not straight—and we do so by visually rehearsing the reflections associated with our perceptual concept for *straight line*, while at the same time attending to the relationship between the alternative pictorial interpretations of the drawn figure.

This is not, strictly speaking, enough to justify the belief that two points determine a *unique* straight line, because we have only considered one case of a non-coincident line drawn between the given points. The succession of visual inferences that convinced us that the top line is not straight, however, depends only on a feature that holds in the general case as well: that the (by hypothesis) straight line and the additional, non-coincident line enclose some (that is, at least one) plane figure. Indeed, they *must* enclose a plane figure, if indeed they coincide at their points but not completely. While this statement is both true and obvious, it is not an analytic or ‘conceptual’ truth in any ordinary sense—rather, it is true in virtue of the distinctive spatial relationship that obtains between plane figures and the lines that are their boundaries, as characterized by Euclid’s special notion of ‘containment’ and

understood in reference to the fundamental pictorial ambiguity of the drawn diagram. As such, it is reasonable to wonder how, precisely, one is able to discern that the statement indeed holds true in *all* possible cases. This question will not be pursued further at the present point, but will be taken up in earnest in Chapter 5, when we turn our focus to the generality problem.

9 Indefinite extensions of straight lines

We now turn to claim (2)—that a straight line can be produced from either end *indefinitely* and *in one way only*—beginning with the latter component, which captures a distinctive property that straight lines share only with circular ones: A given part of the line, however small, is sufficient to uniquely determine the trajectory of the whole. The recognition that straight lines are, in this manner, *uniquely* extendable, turns on appreciating a relationship that we encountered earlier, between reflection symmetry and translation symmetry. As we noted, we can define reflection symmetry for lines *without* presupposing that a (straight) symmetry axis is already given: Instead, the axis (which coincides with the straight line) can arise out of the *performance* of reflection, in roughly the manner in which a straight crease can be imposed on a piece of paper by the action of folding. This is not the case for translation—without prior reference to a given straight axis of translation, it could not be clearly distinguished against the general backdrop of what we might term ‘intra-planar’ rigid motions: those that remain within the plane throughout (excluding reflections, which need to be ‘lifted out’ of the plane to be ‘flipped over’). Reflection

symmetry can therefore be regarded as delimiting the scope of translation, by setting up a straight axis in the first place. So the relationship between the two symmetries runs quite deep. The aspect of this relationship with relevance for the present target property of straight lines is this: Given any reflection-symmetrical spatial context, and any embedded figural component, itself reflection-symmetrical about the same axis, the set of intra-planar rigid motions of the part that will leave overall reflection symmetry intact is identical to the set of translations *along* the very same axis. While the general principle is not easily *stated* in a way that makes its truth immediately apparent, even the general truth *does* seem obvious when encountered in relation to a concrete case, whereupon respect for the overarching reflection symmetry seems, on a phenomenological level, to assert itself as a felt demand that constrains permissible movements.⁵⁹

In the case at issue—which we can take to be the straight line depicted by Figure 4.4—the important recognition is that the only possible symmetry-respecting intra-planar motions to which any proper part of the line can be subjected are translations along the line itself, which remain invisible so long as they do not breach the endpoints.⁶⁰ Recognizing this fact requires no knowledge of the line’s properties

⁵⁹ The kind of translation under consideration here is, of course, not *symmetry* in the sense typically applied to the context of geometry. Given the general definition of ‘symmetry’ as ‘invariance under transformation’, however, it can be thought of as a distinctive kind of (higher-order) symmetry, insofar as it constitutes a transformation under which reflection symmetry itself is invariant.

⁶⁰ As before, we are considering this motion to be applied to a ‘duplicate’ of the original part, the latter of which remains in place when the duplicate is moved out of a state of spatial coincidence with it.

aside from its identification as a *straight line* in the sense captured by the perceptual concept specified earlier. The experience of reflection symmetry as an overarching constraint, together with the immediate detectability of any violation, combine to make the possibility of translating linear parts *without* violation quite visually salient. It is this tacit visual appreciation of the symmetry-respecting character of the translation of straight-line parts that can serve to justify the intuitive belief that the straight line can be uniquely extended. One needs only to imagine such a part being continuously translated to a position that straddles one of the endpoints, such that it now partly overlaps the original line, and partly extends beyond it.⁶¹ The altered figure can then itself be seen alternatively as a single whole line, which can immediately be seen to possess the following properties: (i) It is *greater* than the original line (by Common Notion 5), because it includes the latter as a proper part; (ii) it is *straight*, because the reflection symmetry of the original line has been preserved by the translation used to construct the new one; (iii) it is the *unique* way of extending the original straight line to its new endpoint, because of the way in which

⁶¹ Note that since a *particular* part of the moving part remains coincident with (some or other) part of the original line throughout, there is no need to assume that any form of symmetry (or symmetry axis) is *already* defined or determined (that is, prior to linear extension) beyond the original line. Rather, the moving part's point of contact with the translation axis can be identified in a way that always falls within the spatial scope of the reflection symmetry bound up with the original line. The overlap, together with the fact that the moving part—itsself a straight line—carries its own *local* symmetry axis along with it as it rigidly moves, allows the line to be extended *beyond* its original boundaries in a way that assumes only symmetries defined *within* its original boundaries.

the symmetry-respecting constraint has restricted the extension-constituting motion to translation along the axis marked out by the original line.⁶²

The visual-imaginative experience just described can serve as the basis for the intuition that our given straight line can be produced or extended in a unique way from one of its endpoints, but of course it only permits the line to be extended by a certain amount, which cannot exceed the length of the original line itself. In order to see that the line can be extended *indefinitely*, one needs to be able to apply the procedure in an iterative manner. Fortunately, the use of aspectual imagery to reappraise the result of translating the part as a new, whole straight line already gives us the starting point for a new iteration of the extension procedure. The difficulty that now arises is that our newly constructed straight line is longer than the original, and the outcome of the subsequent iteration will be longer still. Whether we construct the extended lines on paper, or merely in imagination, the successively longer lines will swiftly reach the upper limits of what can be represented by vision or in visual imagination. Even if this were not a problem, however—even if, say, visual experience had no such upper bound—any succession of iterations, however long, will still only support the intuition that the line can be extended by some specific, finite length, namely the length of the total extension realized across the finitely many iterations we have rehearsed so far. This way of approaching the matter is guaranteed never to get us to an intuition that the line can be *indefinitely* extended.

⁶² Uniqueness follows in any case from the straightness of the new line on the assumption of claim (2), whose intuitive justification has already been explained.

This is a difficult problem, and it is one that we will not be able to fully address until Chapter 5, when we provide a general consideration of the *generality problem*, of which this is a specific instance. It is, however, worth briefly previewing the solution that will be given to this version of the problem, given its importance for understanding Euclid's second postulate. The solution does not require us to look any further than the very first iteration of the extension procedure; it turns on the ability to represent the line that results from the first iteration as both (i) greater than the original line, and also (ii) able to be depicted by the same *drawn* line that served to depict our original line. The integrated recognition of (i) and (ii) is based in the same ability to meta-represent one's own visual experience that is implicit in the use of dynamic aspectual imagery, conceived as enabling the perceiver to 'see the same form as *seeable* in two ways'. Taken alone, (i) can of course be grasped directly on the basis of the extension procedure outlined just above. This leaves (ii), whose recognition depends on a form of dynamic transformational imagery in which the change in a figure's appearance is understood to result not from motions applied to the figure, but rather from a shift in subjective, perspectival vantage point. In particular, one can imagine visually 'zooming out' from the figure, such that it appears 'smaller', and indeed, can be made to coincide in imagination with (the appearance of) the original drawn figure.⁶³ By integrating the contents of (i) and (ii) in visual experience, the subject is given a distinctively visual basis for understanding

⁶³ The property of *depictions* of geometrical lines that was discussed earlier, whereby the visible breadth of the drawn line does not 'count' in recognizing the (breadthless) line depicted, plays a crucial role here, since it effectively allows visual-imaginative 'zooming out' to compress visible *length* without compressing visible *breadth*.

that the application of a single iteration of the extension procedure yields a *longer* line, while nonetheless leaving *entirely* untouched the initial (depiction-mediated) appearance of the line, on the basis of which the possibility of applying the procedure was apprehended in the first place. In bringing this understanding to bear on Figure 4.4, one can thereby see not only that the depicted line can be extended in a unique way, but also that such extension will in no way affect or disturb the (visually apparent) possibility of so extending the line. This allows the perceiver, then, to *see* that the line can be extended *indefinitely*.

The solution just sketched to the ‘indefinitely’ clause of claim (2) already gets us on the way to an explanation of intuitive knowledge of claim (3): that a particular drawn line—again, we can consider Figure 4.4—is adequate to represent an *arbitrary* straight line with full generality. For what this solution shows us is precisely that Figure 4.4 is adequate to represent a straight line of arbitrary *length*. Moreover, the crucial move in the solution—the ‘neutralization’ of the extension-induced shift in figural appearance by a compensatory shift in subjective vantage point—can be generalized to capture arbitrary *orientation* as well. In particular, one can imagine the line depicted in Figure 4.4 as rotating about a fixed endpoint for a full turn, thereby encompassing all possible orientations. One can then see, further, the possibility of neutralizing the shift in appearance that accompanies a given increment of rotation, by allowing the reference system assigned to the line to drift along with the rotation; the line in rotation, then, continues to appear horizontally oriented with respect to the reference system it sets up around itself, which allows the perceiver immediately to

intuit the adequacy of the original Figure 4.4 to represent all possible orientations of the line, were it to be subjected to rotation. This explanation depends on several assumptions concerning the properties of reference systems. First, the line itself must be able to set up its own ‘intrinsic’ reference system. Second, the assignment of this *local* reference system must be independent of a distinct, *global* orientation assignment (corresponding to the orientation of the picture plane itself) which remains constant when the line is (imagined to be) rotated—for if this second condition were not met, then the rotation-induced change in orientation could not be seen as neutralized, but rather would be merely invisible. Fortunately, both of these assumptions are highly plausible in light of the characteristics of reference systems we noted earlier. Indeed, both of them would seem to be operative in the simple (and quite immediate) experience of seeing a picture frame as hanging on a wall with a not-quite-upright orientation. For the perceiver who notices this discrepancy presumably does so by representing both the natural ‘up’ direction intrinsic to the frame, and the global, embedding orientation of the wall itself, in an integrated visual experience that coordinates the two reference-system assignments. The only additional feature in the case at hand is the continuous *change* in the local orientation structure as the line’s rotation is visualized.

We have thereby provided a basis in visual experience for the judgments that Figure 4.4 can serve, respectively, to depict straight lines of arbitrary *length* and also of arbitrary *orientation*. There is still a lacuna in our account of the intuitive knowledge of claim (3), however, for we haven’t yet explained what basis the subject

has for assuming that arbitrariness in both length and orientation is sufficient to encompass the full generality implicit in the term ‘arbitrary straight line’. As such, we need to return to the general construction procedure for straight lines provided by Euclid’s first postulate, which tells us that a straight line can be drawn from *any* point to *any* (other) point. The initial ‘point of departure’—the point *from* which we draw our line—can trivially be taken as arbitrary, given that the spatial properties peculiar to a given (geometrical) point are exhausted by its location in the plane, and it is easy to see that the shift in appearance that accompanies change in location can be neutralized simply by means of visual tracking—by keeping visual attention directed at the point. This arbitrariness of location will thereafter be inherited by any figures constructed out of the point, including our present target: the arbitrary straight line. We can suppose that when we visually attend to a point in space, a reference system is applied, with orientation axes meeting at the point’s location, thereby providing some minimal structure to the planar space that opens out locally around the point, enabling the subject to at least roughly discriminate locations about the point itself.

Suppose we now draw a second point at some particular relative location—say, along the horizontal orientation axis. This gives us a diagram that looks like Figure 4.3. As we’ve seen, this diagram effectively invites the perceiver to imagine drawing in the straight line that would connect the two points, which would give us a diagram like Figure 4.4. Now, whichever of these diagrams the perceiver presently beholds, it takes only a little effort to see in it the spatial content that is most naturally depicted by the *other* diagram. That is, the perceiver can not only see the connecting

line 'in' Figure 4.3, but can also see the separated points of Figure 4.3 'in' Figure 4.4, by suppressing visual awareness of the line. The ready availability of this aspectual shift allows the perceiver (to borrow Brown's apt phrase) to "maintain a lively awareness" of both interpretations *while* imagining the free movement of the second point as the original point remains fixed. While grasping the movement of the point on the 'Figure 4.3' interpretation, the motion can then be reappraised according to the 'Figure 4.4' interpretation, so that the movement is then interpreted as a combination of linear rotation (about the fixed point) and linear contraction or extension, as appropriate. This makes visually apparent the fact that the very same motion-appearance *could* have been generated on the 'Figure 4.4' interpretation, by imagining the appropriate combination of linear rotation and contraction or expansion. Of course, this requires the subject to simultaneously keep track not only of alternating aspectual interpretations, but *also* of movement along both of the two 'dimensions' of *length* and *orientation*. This demand on attention can be somewhat mitigated, however, provided the subject has developed a degree of expertise with the visual imagination of straight-linear extension (and contraction); since straight lines are extended/contracted along a *unique* trajectory, the expert perceiver can visualize a line produced all the way to the somewhat indeterminate outer edges of visual awareness, while understanding that the *true* endpoint of the line could potentially lie anywhere along this trajectory. Then, simply by visualizing the (indeterminate) line rotationally sweeping through the space surrounded the central 'reference point', one can visually grasp that all the points lying within this region have been 'captured':

Draw attention to a sufficiently *local* point, and one immediately understands how to construct a line to it, by taking our original horizontal line, imaginatively producing it along its trajectory of extension to the point of indeterminacy, rotating it until it meets the point, and then contracting it appropriately.⁶⁴ Points that are *insufficiently* local, finally, can be handled using a version of the *neutralizing* strategy used to solve the problem of ‘indefinitely’ extending a line. Putting all these components together, then, one can see that a straight line can be drawn from a given point to *any* point, and also that the straight lines so constructible can alternatively be constructed by applying some appropriate combination of rotation and extension/contraction to a given particular line. Since one can also see the possibility of ‘neutralizing’ any change in appearance that results from transforming the line in these ways, the upshot is that one can arrive at an intuitive belief in the truth of claim (3), whose justification is grounded in the visual understanding of the diagrams shown in Figure 4.3 and 4.4. This completes our discussion of the intuitive foundations of Euclid’s first and second postulates.

10 Circles

The third and final construction postulate warrants the drawing of “a circle with any center and distance” (Euclid 1956, p. 199). As indicated previously, circles play something of a subsidiary role in the constructions of Book I of the *Elements*, serving

⁶⁴ Unlike linear extension, contraction is rather trivial, since it can be performed simply by shifting attention to the *part* of the line that runs from the origin point until the position where the second point cuts it.


primarily as ‘scaffolding’ for what is, in effect, the rotational movement of straight lines. By considering two straight lines that share an endpoint as radii of a common circle, Euclid is able to establish their equality to each other. By composing such circle-mediated ‘rotations’, Euclid is able to move straight lines around in planar space while keeping length constant. On our explicitly dynamic approach to Euclidean geometry, the order of priority is reversed: We *assume* the free mobility of straight lines, and can appeal to it to *justify* the construction of various plane figures, including circles themselves.⁶⁵ Since we won’t need circles for our epistemological account of the content of Book I in the next chapter, we won’t devote the degree of careful attention to the third postulate as we did to the first two; we briefly consider it here for the sake of completeness.

The circle has various distinctive properties that might be taken to serve as a basis for a perceptual concept. Euclid’s definition focuses on the existence of a unique central point, to which equal lines can be drawn from any point that lies on the circumference. While this characterization has been taken to indicate Euclid’s preference for precisely the sort of *construction procedure* we will endorse, in terms of the rotation of a straight line about a fixed endpoint (see Heath’s notes in Euclid 1956, p. 184), it is not clear that it captures the basic visual experience of circularity. Giaquinto (2011, p. 290) pursues a different tack, opting to specify a perceptual

⁶⁵ As is clear from our discussion of the common notions, however, Euclid himself assumes free mobility in any case, due to his reliance on the method of superposition associated with Common Notion 4, without which Proposition I.4 (or an equivalent) would have to be taken as a postulate, as is the case in Hilbert’s (1910) axiomatic system for Euclidean plane geometry.

concept for circles in terms of another distinctive property of circles: their reflection symmetry about *all* axes that pass through the center. We will take a different (though closely related) symmetry property as basic, while allowing that sophisticated visual understanding of circles will involve the treatment of alternative perceptual concepts as equivalent, given prior insights to the effect that the different symmetry-based characterizations converge uniquely on the same geometrical object. We will take continuous self-symmetry under rotation to be the basic defining property of circles: To see something as circular, on this account, is to see that rotating it continuously (about the center) leaves it in a state of coincidence with its original position. This perceptual concept does seem plausibly to capture the naïve experience of circularity, as well as the characteristic functional utility of the circular form (think of the *wheel* as well as the *lazy Susan*). Another nice feature of this choice is that it gives us an elegant way to intuit the fact that our construction procedure—rotating a straight line about a fixed endpoint—will deliver the intended result. Consider that the *static* Figure 4.4 can be interpreted (somewhat fancifully, but also with clarity and legitimacy) as an *animation* of a line being continuously rotated, while the subjective frame of reference rotates in the same direction and at the same rate. Keeping this interpretation in mind, consider Figure 4.7.

Figure 4.7: Circle with radius



This *static* figure can be similarly interpreted as an animation: The radius is continuously rotating counterclockwise, as is our subjective reference frame. The circle, which is *in fact* motionless on this interpretation, is moving clockwise relative to our subjective reference frame—this clockwise motion is invisible, however, which shows that the traced line is symmetrical under this rotation. If we visually understand the figure on this interpretation (or one of various other equivalent ones), we can grasp in an instant that when the straight line is rotated about a fixed endpoint, the other endpoint traces a closed curve that is symmetrical under continuous rotation—and which is therefore immediately recognizable as a circle, given the perceptual concept we have chosen.⁶⁶

This completes our discussion of Euclid’s construction postulates. In the next section, we turn finally to his fourth and fifth postulates, which serve, respectively, to

⁶⁶ Of course, much more explanation would be required to fully justify the claim that such a visual experience of Figure 4.7 can justify knowledge of the third postulate.

enforce the global spatial properties of homogeneity and flatness on the planar figures that are constructible by successive applications of postulates 1-3.

11 The structure of Euclidean space

Euclid's fourth postulate reads as follows: "That all right angles are equal to each other" (1956, p. 200). As we established earlier, the notion of equality, as applied to angles, is to be understood in terms of the definition implicitly provided by Common Notion 4: For two angles to be *equal* is for it to be possible that they should be made to coincide, if one is 'applied to' the other by means of the method of superposition. The notion of *right angle* invoked in the postulate is the one captured by Euclid's Definition 10:


When a straight line set up on a straight line makes the adjacent angles equal to one another, each of the equal angles is **right**, and the straight line standing on the other is called a **perpendicular** to that on which it stands. (1956, p. 153; **bold** in original)

In Figure 4.8, line AB has been 'set up' on CD in a way that makes angles ABC and ABD (apparently) equal—which we can visually grasp by rehearsing a reflection across the axis on which line AB lies, which maps these angles onto one another.⁶⁷

⁶⁷ Seen in a different way, this reflection maps line CD onto (collinearity with) itself. This reflection symmetry thereby gives us a basis on which we could specify

Putting together these interpretations of ‘right angle’ and ‘equal’, Euclid’s fourth postulate amounts to the claim that any two configurations that have the symmetry properties of Figure 4.8 can be made to coincide with each other by rigid motion (or ‘application of figures’).

Figure 4.8: Right angles



What is striking about Euclid’s fourth postulate is that it is actually redundant, in the sense that it could itself be proved by the method of superposition, in a manner analogous to the proof of the side-angle-side triangle congruence criterion that Euclid provides in Proposition I.4. In this latter proof, Euclid *assumes* the ability to apply one angle to another.⁶⁸ Now, he has already proven (in Proposition I.2) that a *straight*

perceptual concepts for right angles and perpendiculars, respectively. This won’t be spelled out here.

⁶⁸ More carefully: He assumes the ability to move a configuration of lines that ‘contain’ the angle, so that not only the angular magnitude but also the linear magnitudes are preserved.

line equal to a given one can be placed at a given point. He accomplishes this in his construction by composing what are, in effect, *rotations* of straight lines, made possible without explicit movement by the ‘scaffolding’ provided by circles.⁶⁹ But the construction of circles won’t permit Euclid to ‘move’ angles in the same way (that is, to construct an angle equal to a given one, at a given location), nor will any other postulate.⁷⁰ So the proof of Proposition I.4 has to be interpreted as appealing directly to the assumption of the application of angles.⁷¹ There is, then, no reason why this assumption should not be available for proving the fourth postulate itself, which could then be taken as a proposition. We won’t go through the full proof here, but the basic idea is that we apply the configuration ABCD in Figure 4.8 to its counterpart A’B’C’D’, which is constructed to have the same symmetry properties that are indicated by Definition 10, so that the angles of both configurations are right. We then apply the one configuration to the other so that B coincides with B’, and CD

⁶⁹ The fact that circles can be used to ‘scaffold’ rotations of straight lines rests, of course, on the third postulate, whose intuitive justification *itself* rests on the assumption that we can freely rotate straight lines about their endpoints.

⁷⁰ Of course, the fourth postulate, by asserting equality of right angles directly, underwrites this possibility in the special case of right angles. But this won’t get us to the general case that is needed for the proof of Proposition I.4, because Euclid’s system provides no basis for making precise comparisons in magnitude between right and arbitrary angles. It should also be noted that Euclid *does* eventually (in Proposition I.23) provide a general procedure for constructing at an arbitrary point an angle equal to a given one—but his proof depends crucially on Proposition I.8, which is itself proved by assuming that one angle can be rigidly applied to another.

⁷¹ Moreover, it cannot be proved without this assumption, which is why contemporary axiomatic treatments of Euclidean plane geometry take the side-angle-side congruence criterion—or a close relative, as in the system of Hilbert (1910)—as itself an axiom.

with $C'D'$ —until an endpoint is reached, since we are not assuming equal length—and so that AB and $A'B'$ are positioned on the same side. To demonstrate that the (right) angles are equal, we need to show that the lines containing them coincide, so it suffices to show that AB coincides with $A'B'$ (up to an endpoint), which can be done in different ways by appealing to the (configuration-internal) angular equalities (that is, to symmetry).

Given that the fourth postulate is in this sense redundant, why does Euclid bother to state it at all? Heath's comments on this point are worth quoting in full:

While this Postulate asserts the essential truth that a right angle is a *determinate magnitude* so that it really serves as an invariable standard by which other (acute and obtuse) angles may be measured, much more than this is implied, as will easily be seen from the following consideration. If the statement is to be *proved*, it can only be proved by the method of applying one pair of right angles to another and so arguing their equality. But this method would not be valid unless on the assumption of the *invariability of figures*, which would therefore have to be asserted as an antecedent postulate. Euclid preferred to assert as a postulate, directly, the fact that all right angles are equal; and hence his postulate must be taken as equivalent to the principle of *invariability of figures* or its equivalent, the *homogeneity of space*. (Euclid 1956, p. 200; emphasis in original)

Heath accurately observes that the homogeneity of space (the fact that it has ‘the same structure’ at all locations) is of a piece with the ‘invariability of figures’ under changes in position, so that they can be moved through planar space without deformation. This is important for our purposes, because it shows the manner in which the possibility of subjecting Euclidean objects to free rigid motion (which is a fundamental assumption on our view) serves implicitly to enforce the global property of homogeneity on the space through which figures are moved. Heath’s comments also underscore the appropriateness of Euclid’s *postulating* an equivalent to the assumption of free rigid motion. As we noted earlier, Common Notion 4 implicitly *defines* equality (in the sense of congruence) in terms of possible coincidence under rigid motion. As a definition, this can’t be taken to justify the *existence* of any actual equality—in general, justifying existence is the role of the postulates. The primary role of the fourth postulate, then, is presumably to establish the existence of equality between angles.⁷² The problem with Euclid’s formulation of the postulate is that it doesn’t succeed in establishing that non-right angles can be equal, and this is needed to prove Proposition I.4, on which much of the rest of Euclid’s geometry crucially depends. While Heath’s remark that the postulate “must be taken as equivalent to the principle of *invariability of figures*” is perhaps plausible as an interpretation Euclid’s intentions, it is therefore problematic in that that equality of right angles is *not* equivalent to the assumption of the invariability of figures. In order to properly

⁷² Of course, the fourth postulate as stated will only be able to fulfill this role given the existence of right angles themselves, which is demonstrated in Euclid’s Proposition I.11.

ground Euclid's system, it would be preferable to assert invariability of figures directly, rather than allowing the equality of right angles, which is only a special case, to implicitly justify the general method of superposition. But since angles in general do not have the distinctive symmetry properties that characterize right angles, we can't accomplish this by taking on a postulate to the effect that all 'similar' angles are equal—for there is no available equality-independent way to capture 'similarity' in the general case. The alternative seems to be to take on a postulate that, like the construction postulates, is framed in a way analogous to a *problem* instead of a *theorem*—something along the lines of: "To freely move any figure without deformation." Of course, this corresponds precisely to one of the core assumptions of the dynamic imagery account: that when a figure is apprehended *qua* object, rigid motions can be imagined applied to the figure, by means of dynamic transformational imagery. Not only, then, does the dynamic imagery account allow us to explain knowledge of Euclid's fourth postulate, as stated—it actually directly provides a more acceptable foundation for Euclid's edifice, one that Euclid is in any case committed to assuming.

We now turn, finally, to Euclid's infamous fifth postulate, whose statement is significantly longer than the first four:

That, if a straight line falling on two straight lines make the interior angles on the same side less than two right angles, the two straight lines, if produced

indefinitely, meet on that side on which are the angles less than two right angles. (1956, p. 155)

While this is often referred to as the ‘parallel postulate’, it notably does not mention parallel lines, nor are any lines parallel in the situation it describes. Rather, since the postulate states that the two straight lines *do* meet if produced indefinitely, they are implicitly postulated *not* to be parallel, according to Euclid’s Definition 23:

Parallel straight lines are straight lines which, being in the same plane and being produced indefinitely in both directions, do not meet one another in either direction (1956, p. 154).

This definition—the final one Euclid provides in Book I—departs in one critical respect from all the definitions that precede it: It defines its object with respect to a condition whose representation lies beyond the scope of any possible visual experience. While we noted previously that it is possible for visual experience to justify a judgment *that* a straight line can be produced indefinitely, the explanation of that justification was importantly free of any commitment to a visual experience *of* an indefinitely produced line. That would involve a problematic commitment to the visual experience of an *actual* (rather than potential) infinity. This is precisely the commitment we would have to take on if we attempted to specify a perceptual concept of parallels according to Euclid’s definition, for in order to perceptually


judge that two straight lines were parallel, we would need to be able to visually entertain the result of their indefinite extension so that we could judge on the basis of that visual representation that the indefinitely extended lines do not meet each other. This particular definition is therefore decidedly uncharacteristic of Euclid's definitions, which otherwise serve to indicate how a subject might visually experience a depicted object as having the property in question. It is, in fact, precisely on these grounds that Kant objected to Euclid's definition of parallels. Accordingly, we will attempt to characterize parallels in an alternative way, which does allow for the possibility of visual experiences that represent straight lines *as* parallel.

In particular, we will appeal to the notion of *parallel transport*: Roughly speaking, this is a translation of a straight line along an intersecting straight axis, which preserves the angle it makes with respect to that axis (and hence preserves orientation). On our view, one will see two straight lines as parallel just in case one visually represents their being made to coincide by means of a parallel transport. To be more precise, we will take parallel transport to be defined only in reference to a given pair of intersecting straight lines, and we will define the motion of parallel transport in terms of the translation along one of the lines of a special geometrical object that we term an *intersection*. This term refers to the roughly X-shaped figure that opens out locally around the point at which the lines intersect. (We allow that the T-shaped figures where one straight line meets another along its length, and the L-shaped figures where two lines meet at an endpoint can also be counted as intersections.) An intersection, in this technical sense, does not encompass the full

configuration of intersecting lines, but is rather composed only out of the *parts* of the lines that are ‘close to’ the point at which they intersect. As a rigid object, an intersection thereby serves as a kind of brace, which holds the lines at fixed angles to each other while one is translated along the axis provided by the other. The *parallel transport* of a straight line along another straight line, then, is visually grasped as an alternative aspectual interpretation of the motion of an *intersection* along one of the lines, which ‘forces’ the translation of the other line it ‘holds’ along that same linear axis.


The specification of a perceptual concept for parallel lines in this manner, in terms of the motion of parallel transport, enables us to explain intuitive belief in Euclid’s fifth postulate, provided that we accept the core assumption of the dynamic imagery account, that transformational and aspectual imagery can be employed in an integrated manner in visual experiences of drawn figures. Consider first the construction characterized in this postulate: a straight line, falling across two straight lines, making the interior angles on the same equal to less than two right angles. It is not obvious how we can specify the general construction of such a configuration, where the only restriction on the angles is that they be together less than two right angles. Suppose, then, that we restrict this characterization somewhat, so that we are considering only the range of cases where one of the internal angles is right (and the other is therefore acute). We can construct such configurations by applying a parallel transport to an initial configuration like the following:

Figure 4.9: Diagram for fifth postulate



This is merely an augmentation of the configuration in reference to which Euclid defines right angles (shown in Figure 4.8), with an additional line AB drawn internal to the angle CBE. It follows from construction that angle DBE is right, and also that angle ABC is acute (since it is a proper part of angle CBE, which is right). Suppose we now apply a parallel transport to AB along CD in the direction of D. The dotted line in Figure 4.10 indicates roughly where AB would be located after being transported along this axis by a certain amount.

Figure 4.10: Animated diagram for fifth postulate



If we visualize this transformation *as* a parallel transport of AB along CD, we can immediately see that it is also seeable as a translation of the *angles* AB makes where it meets CD, through a trivial application of aspectual imagery. Simply put, when a line is parallel-transported, the angles ride along for free. In particular, angle ABC has been translated, so that it now lies on the opposite side of BE. Since it remains acute and we know by construction that DBE is right, the result of applying this parallel transport to the initial configuration can now be seen to be an instance of the kind of configuration characterized by Euclid’s fifth postulate—for the straight line CD “falls upon” the straight lines BE and transported-AB, and the interior angles on the same side as E (the right angle DBE and the acute angle transported-ABC) are together equal to less than two right angles.

Now, what the fifth postulate claims about this figure is that if BE and transported-AB are produced indefinitely, they will meet on the same side of CD as

E. This can be seen to be the case, provided we use aspectual imagery to reappraise the parallel transport by which we construct the configuration to which the postulate refers. For note that in the initial configuration (in Figure 4.9), since AB intersects with *two* lines at point B, it is held by *two* distinct intersections (in our technical sense), and hence there are *two* axes along which it could be parallel-transported. So instead of imagining the parallel transport of AB along CD in the direction of D, we could alternatively imagine the parallel transport of AB along BE in the direction of E. The important thing to notice here is that it doesn't make any difference which one of these options is used to drive the transformational imagery that rehearses the parallel transport—the *appearance* of the visualized continuous motion of overall configuration is exactly the same in both cases. This means, in particular, that when we imagine AB to be parallel-transported in the manner that yields the configuration to which the fifth postulate refers, we can use aspectual imagery to see the configuration-in-motion *as* the parallel transport of AB along BE in the direction of E. On this latter interpretation, we will see the intersection that initially holds lines AB and BE as being translated continuously along BE in the direction of E. Moreover, we can see that this translation can be continued indefinitely, in the same way that we were able to grasp that a straight line can be continuously extended in a given direction, as explained previously. The upshot is that we can *see that* as we imagine AB to be parallel-transported along CD in the direction of D, that motion is *seeable* as the translation of the AB-BE intersection along BE in the direction of E. We can thereby visually grasp that the construction of the configuration to which the

fifth postulate applies guarantees that BE and transported-AB will meet on the same side of CD as E—for these two lines will remain held in place throughout by the AB-BE intersection. And this is just to see that what the fifth postulate asserts about this configuration is true.

Now, this doesn't quite give us a visually justified belief in the fifth postulate as stated, because we haven't considered the full range of cases to which it applies, in which it may be that neither of the interior angles is right, so long as they are together less than two right angles. This doesn't pose an insuperable difficulty, however—in the next chapter we will see, in connection with our discussion of Euclid's Proposition I.13, how we could capture the general case. In any case, we have in a significant sense already achieved our target, for in the technique just used to justify intuitive belief in the restricted cases, we have already laid down the foundations that are necessary for the visual apprehension of the truth of those propositions that depend, in Euclid's system, on his fifth postulate. The assumption—one that follows directly from our account's core postulation of the *integrated* employment of transformational and aspectual imagery—that we can *reappraise* a parallel transport along one line *as* a parallel transport along a different line is, in fact, equivalent to Euclid's fifth postulate, in the context of the first four postulates. It is all we will need in order to impose the distinctive Euclidean *flatness* on the configurations that are constructible in accordance with postulates 1-3. With this result, then, we complete our discussion of Euclid's postulates, and with that, our account of basic knowledge of Euclidean plane geometry.

12 Conclusion

In the last two chapters, we have taken the first and most significant step forward in explaining how the core psychological posits of the dynamic imagery account are able to support justified intuitive belief in the claims of Euclidean plane geometry, on the basis of visual experiences of drawn figures. We began by considering Giaquinto's account of basic geometrical knowledge, and we then applied his central idea of perceptual-geometrical concepts to the most fundamental claims of Euclid's geometry, through a careful reading of the definitions, postulates, and common notions that appear at the beginning of Book I of the *Elements*. The upshot of this investigation is that Euclid's fundamental starting points are themselves based in more fundamental assumptions, which constitute the true 'psychological axioms' underlying Euclidean plane geometry. These consist in the perceiver's ability to employ, in an integrated fashion, the two forms of dynamic visual imagery postulated by our account. In the next chapter, we will consider how this account can be extended to explain general knowledge of Euclidean theorems.

Chapter 5: The Generality Problem

1 Introduction

In this chapter, I address an outstanding challenge to the account of geometrical knowledge put forward across the previous chapter: the generality problem. This problem arises largely when one attempts to extend that account beyond the domain of basic geometrical knowledge to the understanding of geometrical theorems.

2 The problem


The generality problem arises due to a contrast between the inherent particularity of the geometric objects that can be depicted by diagrams, on the one hand, and the generality of the theorems those diagrams are intended to support, on the other. It should be clear that the problem does not arise with the same force for geometric objects of all kinds. Any particular circle or square, for instance, is geometrically similar to all others, in the sense that any such pair of the same kind can be made congruent by means of uniform scaling. Here, the problem of generality is relatively trivial, since ensuring the generality of a conclusion reached in reference to a depicted circle or square requires nothing more than recognizing that the result persists across uniform scaling. A different order of difficulty confronts the attempt to generalize a truth grasped about a particular depicted triangle, so that the result can be legitimately held to apply to *any* triangle whatsoever. In such cases, the generality of the intended conclusion appears to outstrip the scope of what can be depicted in a single diagram,

or indeed, in any finite set thereof. The reason, of course, is that the class of all triangles encompasses an indefinite range of minute variations in shape, so that any particular depicted triangle will necessarily possess accidental spatial properties that are not characteristic of the class as a whole. The problem is that this ‘particularizing’ character of triangles would seem to render theorems concerning triangles *in general* unknowable by diagram-based means. At best, it would seem that generalizations could be drawn across a sample of special cases, which might provide inductive support for the general theorem. If this were the best that could be achieved through visual understanding, then that would considerably undermine the claim of the latter to be a genuine source of geometrical knowledge.

Discussions of the generality problem go back at least to Proclus’ fifth-century commentary on the *Elements*, and continue in the modern period with Locke, Berkeley, and Kant. Interestingly, these discussions all seem to take as their common example a single theorem from Euclidean plane geometry: the *angle-sum theorem*, Proposition 32 from Book I of the *Elements*. This choice is perhaps unsurprising, since this theorem is perhaps the most basic theorem of Euclidean geometry that establishes a truly significant result about triangles in general. While I will not rehearse Euclid’s proof in full, it will be instructive to acquaint ourselves with the basic character of Euclid’s diagram-based reasoning. To begin with, the theorem claims that the sum of the interior angles of any triangle is equal to two right angles (equivalently, to a ‘straight angle’). The proof consists of an initial construction phase, followed by a demonstration performed in reference to the constructed

diagram. We are first asked to construct (i.e., to draw) an arbitrary triangle ABC , and then to extend line BC to a point D , and to draw a line CE , parallel to AB , on the same side of BD as A . This yields a diagram like the one in Figure 5.1 (though of course, any actual construction will necessarily introduce accidental features as regards the particular shape of the triangle).

Figure 5.1: Diagram for Euclid's Proposition I.32



In order to see that the theorem holds true for the particular triangle depicted, it suffices to attend to two facts in concert: first, that the three angles about point C (on the 'upper' side of line BD) are together equal to two right angles; second, that these three angles are equal to the interior angles of the triangle itself. This second fact is grasped by considering the angles in turn: First, we see that $\angle ACB$ is common to both the triangle and the arrangement about point C ; second, we see that $\angle ECD = \angle ABC$

(since these angles are, respectively, ‘exterior’ and ‘opposite-interior’ in relation to parallel lines AB and CE); third, we see that $\angle ACE = \angle BAC$ (since these angles are ‘alternate’ in relation to the Z-shaped arrangement of parallel lines AB and CE, joined by the transversal AC). When these various spatial relations are grasped in combination, as sharing the space of a common figure, the conclusion—that the interior angles of (the particular, depicted) $\triangle ABC$ are equal to two right angles—becomes intuitively apparent.

In Euclid’s own presentation, the steps are individuated more finely than this, and are arranged linearly, in deductive sequence—all of which might suggest to the contemporary reader that something like a *formal* deduction is being provided. This impression would be misleading, however, for the argument depends at almost every step on what can be seen in the diagram. For instance, the fact that angles $\angle ACB$, $\angle ACE$, and $\angle ECD$ collectively fill (without remainder) the space about point C on the ‘upper’ side of line BD is something that must be seen by inspection of the diagram, as must facts about the spatial arrangements of the angles, such as that $\angle ACE$ and $\angle BAC$ are ‘alternate’. In fact, even the *existence* of these angles is not something that follows deductively from any postulates or prior theorems set out by Euclid; they must rather be seen to ‘arise’ spatially out of the specific placement of lines that is prescribed in the construction phase, as their ‘negative space’. It follows, then, that regardless of whether in understanding Euclid’s proof we place primary emphasis on the diagram or on the accompanying text, our conclusion appears to be limited in scope to the special case our diagram depicts. On the first option, where we

read the text simply as a guide to the way spatial relations within the diagram must be appreciated in combination, this limitation is obvious, because the theorem has been seen to hold only for this particular depicted triangle. On the second (more traditional) reading, where we identify the *proof* primarily with the deductive argument set out in the text itself, we note that this argument contains a number of premises (some explicit, some unstated) that are justified only by what is observed to hold of the geometric structure depicted by the diagram; in this case, the scope of our conclusion will be limited to that for which these premises have been established, which is to say, to the depicted special case.

To be sure, there is much more to say about the visual processes by which we can derive diagram-based knowledge of the *special* case of the angle-sum theorem, but even this brief consideration suggests that there is no essential bar to providing such an account. In contrast, it remains utterly mysterious how we can attain knowledge of the *general* case by diagram-based means. Repeating the proof for multiple variations of the diagrammatic construction may be *practically* sufficient to convince us of the truth of the general claim, but it is the principle that matters here, for we will still remain in the realm of special cases, of which there are indefinitely many—in this case, our justification for believing the general claim will be merely inductive, and this falls short of an acceptable standard for mathematical knowledge.

In a particularly influential critique of the diagram-based view, Kitcher (1984) argues forcefully that this problem—which he terms the *irrelevance problem*—is in fact fatal to the whole enterprise. Understanding Kitcher's critique will prove useful

for appreciating the deficiencies of the attempted solutions to be considered shortly. Kitcher frames his criticism in response to Kant's version of the diagram-based view. Briefly, Kant argues that our geometric knowledge is grounded in the constraints our mind necessarily imposes on any possible spatial experience—including our perceptual experience of constructed geometric figures. The constructed figure, as perceived, reflects not only properties that are contained in the *concept* of triangle (such as being three-sided), but also properties that are necessarily imposed on triangular constructions by the nature of spatiality in general (such as the angle-sum property). Our *synthetic a priori* knowledge of geometry arises out of our intuitive discernment of properties of the latter sort (which in Kant's view, take us 'beyond' our concepts themselves). Kitcher points out, however, that the spatial constraints imposed on the construction of a triangular figure necessarily introduce a *third* sort of property, namely the *accidental* properties of specific shape that arise as an inevitable by-product of spatial construction (for instance, that a triangle is acute rather than obtuse). The problem is that there is no apparent diagram-based means for distinguishing between the (relevant) properties that necessarily result from spatiality in general, and the (irrelevant) properties that *accidentally* result from the particular spatial construction. Given that the particular figure has the angle-sum property *as well as* (say) the acute-angled property, in order to derive a suitable general conclusion from our perception of the diagram, we need to be able to (legitimately) generalize over the former sort of property, while avoiding (illicitly) generalizing over the latter sort. But that would seem to require that we have antecedent

knowledge of which properties are spatially necessary, and which are accidental—precisely the knowledge, Kitcher observes, that the diagram-based view is supposed to provide. Kitcher’s negative conclusion is that putative solutions to the ‘irrelevance problem’ can be presumed to be circular, tacitly assuming possession of the very knowledge they purport to explain. As we shall see shortly, this is in fact a fair assessment of a number of the solutions that have been offered to the problem.

The plan for the rest of the chapter is as follows. In the following two sections, I consider, respectively, historical and contemporary solutions to the generality problem, concluding that none are satisfactory. Nonetheless, I suggest that Kant’s attempted solution provides an indication of the lines along which a successful solution might be provided. This task is then taken up in the final two sections. I first briefly sketch some of the key components of the dynamic imagery account that has been developed over the course of the last two chapter. This account is then applied to the specific case of the angle-sum theorem, yielding a novel solution to the generality problem.

3 Historical solutions

I will consider three different historical solutions to the generality problem: those offered by Locke, Berkeley, and Kant. While I will argue that each of these solutions end up failing, they collectively provide a useful map of the available conceptual terrain, with Locke and Berkeley occupying opposite horns of a dilemma that Kant attempts to straddle.

Locke's solution is effectively to deny that the diagrammatic proof acquaints us in the first place with a mere special case. Euclid's proof, on his view, is performed not on a particular triangle, but rather on the "general Idea of a Triangle", which is "neither Oblique, nor Rectangle, neither Equilateral, Equicrural, nor Scalene; but all and none of these at once" (1975, 596). This general Idea of a Triangle is derived by abstraction from particular perceived instances: that is, by "leaving out but those particulars wherein they differ, and retaining only those wherein they agree" (412). The problem with this proposal is that, as we saw in the previous section, Euclid's proof depends on relations that are observed to obtain in the perceived figure, and insofar as one has abstracted away from the determinate shape of that figure (which possesses the accidental properties noted by Kitcher), it is unclear in what sense there remains anything like a figure to perceive. Observing that two angles are 'alternate', for instance, seems to require that one be observing some concrete figure in which they appear in such a spatial configuration. If Locke's Idea of Triangle is meant to be something of which we can have an intuitive, perception-like acquaintance, then it is something with which we have no familiarity in ordinary experience, for there is no obvious answer to what such a general Idea *looks like*. On the other hand, if Locke does *not* intend that we should have at least a quasi-perceptual intuition of his Idea of a Triangle, then it is unclear how he can maintain that Euclid's proof is demonstrated for such an Idea, since the proof patently depends on something like the perceptual observation of a figure. Locke's solution, then, fails even to get off the ground.

Berkeley's solution, in contrast, steadfastly affirms the particularity of the figure in relation to which the theorem is proved; he simply denies that this poses a problem for the subsequent inference to the general conclusion:

To which I answer that, though the idea I have in view whilst I make the demonstration be, for instance, that of an isosceles rectangular triangle whose sides are of a determinate length, I may nevertheless be certain it extends to all other rectilinear triangles, of what sort or bigness soever. And that because neither the right triangle, nor the equality, nor determinate length of the sides are at all concerned in the demonstration. It is true the diagram I have in view includes all these particulars, but then there is not the least mention made of them in the proof of the proposition. It is not said the three angles are equal to two right ones, because one of them is a right angle, or because the sides comprehending it are of the same length. Which sufficiently shows that the right angle might have been oblique, and the sides unequal, and for all that the demonstration have held good. And for that reason it is that I conclude that to be true of any obliquangular or scalenon which I had demonstrated of a particular right-angled equicrural triangle, and not because I demonstrated the proposition of the abstract idea of a triangle. (1988, Introduction §13)

The basic idea in this passage seems to be that, since “there is not the least mention” of any property not universal to all triangles in Euclid's proof, therefore the proof

does not depend on the triangle possessing any such particular property; in that case, the conclusion will be known to hold for triangles in general.

This line of reasoning faces an obvious problem, however: It does not follow from the fact that a proof fails to explicitly *mention* that the triangle possesses a particular property, that the proof does not implicitly *depend* on the triangle possessing that property. In the case of a purely *formal* proof this would indeed be a valid inference to draw. As seen in the previous section, however, Euclid's proof is *not* purely formal; rather, it depends on a number of premises that are supported only by inspection of the diagram, and that diagram necessarily depicts only a special case. Now of course, Berkeley is correct that the truth of these premises depends only on properties that are universal to all triangles. What is in question, however, is how we are in a position to *know* that this is so, without actually observing that the relevant spatial relation holds for each of the indefinitely many possible constructions. Berkeley's solution, then, appears to succumb to precisely the sort of vicious circularity Kitcher warned about in connection with the 'irrelevance problem'. In justifying his claim that Euclid's argument does not depend on properties that are peculiar only to certain triangles, Berkeley is helping himself to the claim that the diagram-supported premises depend only on properties shared by *all* triangles. This is *true*, but the question of how we are able to *know* it is true remains unanswered. Berkeley's solution fails.

Finally, we turn to Kant's solution, which can be viewed as an attempt to carve out a middle ground between the Scylla of Locke's solution and the Charybdis

of Berkeley's. Consider that Locke's solution was unconvincing in its attempt to import generality into the initial intuition of the geometric figure; Berkeley's solution, in contrast, fully acknowledged the particularity of the intuited figure, and then failed to explain how the deliverances of this intuition can be generalized beyond the particular case. Kant's solution attempts to straddle the gap between generality and particularity by appealing to *schemata*, which seem to serve as something like an intermediary between general concepts and particular intuitions. Kant's account seems to me to be unsuccessful as well, though I think it does point the way towards a genuine solution to the generality problem.

On Kant's view, the judgment that the general theorem is true rests on our possession of a general *schema* for constructing the figures that depict particular instances of triangles:

In fact it is not the images of objects but schemata which ground our pure sensible concepts. No image of a triangle would ever be adequate to the concept of it. For it would not attain the generality of the concept, which makes this valid for all triangles, right or acute, etc., but would always be limited to one part of this sphere. (1998, A140/B180-A141/B181)

By 'schemata' here Kant is referring to the "representation of a general procedure of the imagination for providing a concept with its image" (A140/B180), e.g., the general rule according to which a triangular figure can be constructed (on paper or in

the imagination). Here Kant is taking seriously the presence of the initial ‘construction phase’ in the Euclidean proof. There are various general procedures for constructing arbitrary triangles—the simplest is to pick three arbitrary points in the plane (not all collinear), and to connect them. What is difficult to see is how Kant can appeal to such a schema in order to solve Kitcher’s irrelevance problem. For the schema of triangularity *itself*, as Kant notes, “can never exist anywhere except in pure thought”, and in particular, “is something that can never be brought to an image at all” (A141/B181). Rather, the schema embodies the *procedure* by which images of particular triangles (which necessarily have accidental spatial properties) can be constructed. Now, if we attend only to the perceivable result of *applying* the schema, we are unable to discern which qualities are common to all triangles constructed on that basis. On the other hand, if we try to reason about the general constraints on what is so constructible by entertaining the schema itself in “pure thought”, we lose the synthetic character of our judgment, which is clearly not Kant’s intent. Kant responds to this dilemma as follows:

Thus I construct a triangle by exhibiting an object corresponding to this concept, either through mere imagination, in pure intuition, or on paper, in empirical intuition, but in both cases completely a priori, without having to have borrowed the pattern for it from any experience. The individual drawn figure is empirical, and nevertheless serves to express the concept, without damage to its universality, *for in the case of this empirical intuition we have*

taken account only of the action of constructing the concept, to which many determinations, e.g., those of the magnitude of the sides and the angles, are entirely indifferent, and thus we have abstracted from these differences, which do not alter the concept of the triangle. (1998, A713/B741-A714/B742, my emphasis)

So Kant's attempted solution to the generality problem depends on the claim that we perceive the particular constructed triangle *qua* (and only *qua*) its status as a product of the general schema or procedure for constructing all such triangular figures. What Kant highlights here is that the 'drawing' of the figure is an action we undertake deliberately, and as such, in subsequently perceiving the figure, we can 'take account' only of what we put into it by enacting this procedure.

While I think that this account is on the right track, it leaves key questions unanswered that are crucial for solving the generality problem. For what is missing is precisely an explanation of how the 'free' character of any particular active construction (based in the fact that choices regarding, e.g., the lengths of the sides of the triangle, are arbitrary with respect to the general schema) is linked to the perception of the specific arbitrariness that obtains in the particular figure that emerges as a product. What is required is that, in grasping the figure as a product of one's action, one *sees that* one could have drawn it differently, and that the important relations one observes as holding in the particular figure would remain invariant. But how is one to *see* that this is so, if not by actually enacting all these (indefinitely

many) alternative constructions as well, and observing that the same truth holds across the entire set? In order for Kant's suggestion to constitute a viable solution to the generality problem, it is necessary that one be able to perceive the particular figure as in some sense 'modally animated', that is, situated amongst alternative variants, in such a way that one simultaneously apprehends the full range of possible variation (encompassing all particular triangles), together with the important invariants that obtain across this range (that all such triangles have certain geometric properties). This is the line of thought I take up later in the paper, in developing my own positive account. First, however, I briefly consider some contemporary solutions to the generality problem.

4 Contemporary solutions

Here I'll consider two representative solutions to the generality problem from contemporary philosophers. The discussion will be relatively routine, however, as these contemporary solutions are both recognizable as variants of Berkeley's solution, which we have already discussed.

The first contemporary solution is put forward by Norman (2006, 158-9), toward the end of his book-length case study of the epistemology of the angle-sum theorem. Norman's book contains a number of useful insights about the nature of diagram-based inferences, some of which are embraced by the account I will put forward in the following section. Here, however, we are concerned just with Norman's solution to the generality problem. Norman approaches the problem by first

distinguishing between generic properties (shared by all triangles) and non-generic properties (peculiar to only some triangles), and framing the question at issue as one of whether the reasoner can be justified in believing that Euclid's proof does not depend on any non-generic properties of the triangle depicted. Norman now considers the three components of construction phase of Euclid's proof: the drawing of an arbitrary triangle ABC, the extension of line BC to point D, and the drawing of line BE parallel to AB (and on the same side of BD). Norman notes, quite correctly, that each of these construction steps is warranted on the basis of some prior postulate or theorem, and that none of them requires that the triangle possess any particular non-generic property. Specifically, none of them explicitly mentions, nor depends upon, any property not common to all triangles. Norman now says:

But this is all the reasoner needs to believe in order to make the generalization with justification. Recall that, by the rule of Universal Generalization, if a claim about a given but arbitrary object rests on no prior assumption about the object in question, then it may be generalized into a claim about all such objects. The geometrical analogue of this claim here would be that if a claim about one or more given but arbitrary triangles rests on no prior assumption about the triangles in question, then it may be generalized into a claim about all such triangles. The antecedent of this conditional is satisfied here. No assumption has been made as to the triangles represented by the diagram.... So the reasoner can make the general claim with justification. (158-9)

This is, clearly enough, merely a variation on Berkeley's solution, which held that generalization can be justified by the observation that Euclid's argument fails to mention any non-generic property of the triangle in question. Here Norman goes a step farther, claiming explicitly that the proof does not *depend* on the assumption of any non-generic property. In fact, Norman has shown no such thing. He has shown only that the general *construction procedure* does not depend on the triangle possessing any non-generic property—this much does indeed follow from the justification of the various construction steps by prior postulates and theorems. But in order for the proof actually to be carried out, the construction procedure must be *performed* so as to yield a particular constructed figure. It is this *particular figure* the observation of which will serve to support the subsequent inferences in Euclid's deductive argument. And Norman has offered no reason to think that the spatial relations seen to hold in the constructed figure do not depend on any non-generic properties. After all, the figure itself will necessarily possess such properties, as Kitcher has stressed. The analogy Norman draws to the formal inference rule of Universal Generalization is a poor one, because the diagram-based proof of the angle-sum theorem is not a formal proof. It rather depends on constructing a particular figure in which certain relations are seen to obtain. Universal Generalization only applies where one has assumed nothing in particular about the object under consideration. In Euclid's proof, the 'particularization' that takes place when one actually *enacts* the general construction procedure so as to yield a concrete figure

violates this constraint, since subsequent inferences are justified by what is seen to be true of the particular figure itself. Norman has merely obscured this point by drawing attention to the general construction procedure, rather than the figure constructed on its basis.

Our next candidate solution is provided by Giaquinto (2007), whose account of basic geometrical knowledge we considered in detail in Chapter 3. Giaquinto's solution to the generality problem, in relation to cases like the angle-sum theorem, is essentially the same as Norman's, and fails for the same reason. In discussing the question of valid generalization from diagrams, Giaquinto considers a different example from the angle-sum theorem (Thales' theorem), but the essential points are all the same, since in this example too, we have inferential steps that are licensed by observation of a diagram depicting a geometric object with non-generic properties. Like Norman, Giaquinto appeals to the formal inference rule of Universal Generalization, and considers whether the reasoner can be justified in believing that the proof of the special case depends on no non-generic features, and hence can be justified in applying this generalization rule. Here is the crucial passage:

The question we need to answer is whether this kind of valid generalizing can occur in an argument that uses a diagram in a non-superfluous way to reach an intermediate statement about an arbitrary instance. The worry is that in using a diagram to reason about an arbitrary instance c of a class K , we will be using some feature of c represented in the diagram that is not common to all

instances of the class K.... Is there [such an] error in our example...? ... What property of instance not shared by all members of the class might we be unwittingly relying on when following the argument? Well, there is the size of the semicircle... and there is the orientation of the semicircle. But it is clear that the argument does not depend in any way on these features. There is also the position of the chords' meeting point... [and other properties—these details need not concern us]. A simple step by step inspection makes it clear that none of these properties is relied on in the argument.... As there is no threat from any other properties that the figure is represented as having, the argument contains no violation of the conditions for deductively valid generalizing. I conclude that valid generalizing can indeed occur in an argument that uses a diagram in a non-superfluous way to reach an intermediate statement about an arbitrary instance. (79-80)

Let's consider Giaquinto's argument in relation to our own example, the angle-sum theorem (our considerations here will apply in the same way to his example of Thales' theorem). Giaquinto proposes that, in order to know that the diagram-based proof of the special case is independent of the assumption of any non-generic properties, it suffices to perform a "step by step inspection" of all such properties, checking each to ensure that our proof depends on none. But how exactly is this inspection to be performed? Consider that the proof of (a given special case of) the angle-sum theorem depends on observing in the diagram properties like the

following: that *these* two angles occupy ‘alternate’ positions in the spatial configuration, that *those* three angles fill (without remainder) the space about a given point on one side of the line on which it is placed, and so forth. These facts are seen to hold in the particular constructed diagram, which as we have seen, is not ‘arbitrary’, but rather necessarily possesses some or other non-generic properties. How can we tell ‘by inspection’ that *these* properties are indeed generic ones, given that our ‘inspections’ are restricted to the special case depicted by the diagram? In order to convince ourselves that these properties are generic, we might consider an alternative construction, to see if they still obtain. Even if they do, however, this will only count as inductive evidence in support of the *hypothesis* of their generic status. In order to establish the truth of this hypothesis with certainty, we will need to inspect *all* possible constructions. What Giaquinto has assumed here is that we are able to *see* that a given property will hold of *all* possible constructions—but an explanation of this supposed ability is precisely what is at issue in the attempt to solve the problem of generality. Kitcher’s warning about circularity has proven prescient: Giaquinto has ended up presupposing precisely the knowledge he set out to explain. I conclude that both of these contemporary solutions fail, for essentially the same reason that Berkeley’s solution fails.

5 Review of the dynamic imagery account

In this section I briefly review the main features of the account of geometrical understanding that has been developed across the previous two chapters. I will restrict

myself to sketching some of the key features of the account that are relevant for addressing the generality problem as it arises for the angle-sum theorem. In the next section I go on to provide a novel solution to this problem based on the account.

The basic idea at the core of the account is that diagram-based insight into mathematical truth arises by means of the ‘animation’ of a visual diagram, accomplished by means of projecting ‘dynamic’ visual imagery into the static form. This dynamic imagery comes in two basic varieties. First, *dynamic aspectual imagery* captures the ability to appreciate, in an integrated fashion, the manner in which alternative ways of parsing or structuring a perceived form converge on a common spatial figure, in such a way that the mapping between these different ‘aspectual contents’ becomes intuitively apparent. A simple example is provided by the way that we can come to appreciate the way the ‘duck’ and ‘rabbit’ interpretations of the ambiguous duck-rabbit figure coexist as possible ‘takes’ on the same raw visual form. By attending to our alternating perceptual interpretation, we can grasp the way the duck’s bill ‘maps onto’ the rabbit’s ears. Second, *dynamic transformational imagery* actively represents the possible salient transformations of a given spatial form (or parts thereof), including, for example, translations, rotations, and deformations of lines and figures composed out of lines. This latter form of imagery allows us to imagine rigid translations of lines onto their parallels, for instance, or the continuous dilations and contractions of angles that result from rotations of the lines bounding them. The claim is that by bringing these two varieties of dynamic imagery to bear in concert, in relation to an appropriate geometric diagram, we can intuitively grasp not

only *why* a theorem is true in the depicted special case, but can also *see* that the spatial basis for its truth in the depicted case will extend to the non-depicted cases as well. This will constitute a solution to the generality problem.

I'll now discuss both types of dynamic imagery as they apply specifically to the case of the angle-sum theorem. Dynamic aspectual imagery is important, in the first place, in order to appreciate the interrelations of lines and their associated angles. As noted earlier, angles are not explicitly mentioned in the construction phase of Euclid's proof; they rather arise as by-products of the placements of lines, as, in effect, their 'negative space'. It is crucial for grasping even the truth of the special case that the observer be able to see a given figure alternately as depicting two lines that share an endpoint, and as depicting the angle 'contained' by these lines; moreover, that the observer be able to appreciate the 'mapping' in virtue of which these alternative aspectual contents are mutually constraining. Indeed, the importance of this instance of 'seeing a figure two ways' is directly tied to what on Kant's account appears as the *synthetic a priori* character of mathematical judgments. While the philosopher "will never produce anything new" with the concept of a triangle, Kant says, the geometer

extends one side of his triangle, and *obtains two adjacent angles* that are together equal to two right ones. Now he divides the external one of these angles by drawing a line that is parallel to the opposite side of the triangle, and

sees that *here there arises an external adjacent angle* which is equal to the internal one, etc. (1998, A716/B744, my emphasis)

The ‘synthesis’ of the angles that are said to “arise” out of the drawing of a line in this passage is the result of the observer’s appreciation of a figure-ground shift, which is based in the employment of dynamic aspectual imagery. In the construction phase, the line is apprehended thematically, as a sort of object, placed within an empty space; the drawing as a whole is then reappraised in a way that regards the pie-slice-shaped regions (the ‘angles’) as objects, whereupon the line is viewed negatively, as the boundary they share. In addition to facilitating the grasp of line-angle interdependencies, dynamic aspectual imagery will serve to mediate among alternative groupings of angles. For instance, the angle $\angle ACB$ will need to be seen alternately as belonging to the group of interior angles of the triangles, and then as belonging to the group of angles about point C, which fill the space on the upper side of line BD.

Dynamic aspectual imagery can also operate at a higher level, when applied in reference to a figure already animated by dynamic transformational imagery. In this case, the spatial dynamics of an imagined transformation can be interpreted in alternative ways, in such a way that the alternative interpretations are mapped to each other. For a concrete example, consider the rotating helixes traditionally found outside barbershops. The movement of the helical lines can be seen in the ‘veridical’ manner, as a horizontal rotation about the axis, but the lines are also easily interpreted

as moving in a vertical direction. By progressively switching back and forth between these alternate ‘takes’ on the helical movement, one is able to arrive at an appreciation of the manner in which the perceived form supports both interpretations. This ‘higher-level’ use of dynamical aspectual imagery will be important for grasping the general truth of the angle-sum theorem.

Dynamic transformational imagery is the crucial component for explaining how one can grasp that the *basis* of truth in the special case extends to the other cases not depicted by the static diagram. The basic idea here is that one imagines continuous spatial transformations that would map the perceived figure onto the full range of its counterparts, which collectively realize the totality of figures that can be drawn in accordance with the relevant construction procedure. It is crucial that as one imagines this range of transformation, one simultaneously attend to the spatial *basis* for the theorem’s truth in the case depicted. Since attending to the full set of relations can be taxing for visual attention, it is helpful that the different components that constitute this basis can be considered separately. The intuitive apprehension that the relation will generalize appropriately consists in each case in the recognition that as the figure is gradually transformed, the spatial features that constitute the *reason why*, for instance, one angle is equal to another will remain constant. In general, this ‘reason why’ (in the special case alone) may itself be grasped only by virtue of the combined use of both kinds of dynamic imagery.

As that comment implies, dynamic transformational imagery has an important role to play before the question of generalizing from the special case even arises. In

the case of the angle-sum theorem, its specific use in this capacity concerns the perception of lines AB and CE as parallel. According to the dynamic imagery account, one important way in which we can see two lines *as* parallel consists in our imagining the rigid translation of one onto the other. Given that the line CE is stipulated to be parallel to AB in the construction phase of Euclid's proof, this use of dynamic transformational imagery will be implicated in our understanding of the constructed figure from the outset. In seeing CE as parallel to AB, we will be poised to imagine AB as 'sliding' along the line BC (which functions somewhat like a rail), until it coincides with CE. In addition, repeatedly enacting this translation in imagination will form the basis for our grasp of both of the crucial angle equalities.


6 A novel solution

To a significant degree, our diagram-based grasp of the general truth of the angle-sum theorem rests on a prior grasp of the general truth of two preliminary results (proved by Euclid in Propositions 13 and 15 of Book I). Therefore, a complete account of our diagram-based knowledge of the general angle-sum theorem will begin by addressing these preliminary results.

First is the theorem that the sum of the angles that fill a straight line (without remainder) about a point is equal to two right angles. This theorem is a trivial generalization from the claim that any straight line set up on a given horizontal will form angles equal to two right angles. In order to grasp *this* result, it suffices to

consider the diagram in Figure 5.2, which depicts a line placed upright on a straight line so that it is perpendicular.

Figure 5.2: Demonstration of Euclid's Proposition I.13

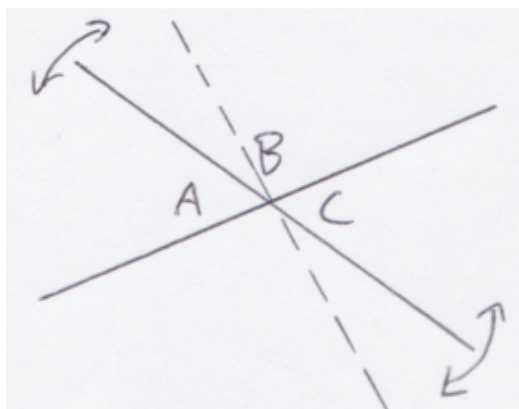


By Euclid's definition, the angles that arise on either side of this line in virtue of its placement will both be right angles. In order to grasp that an *arbitrary* line set up on a given straight line will form angles equal to two right angles, one must first employ dynamic transformational imagery in order to imagine the rotation of the initial (perpendicular) line about the point at which it meets the horizontal line. (It is easy to see that extensions and contractions of the line are irrelevant, since imagining these transformations leaves the angles visibly the same.) One then has to recognize that the full range of this rotation is adequate to capture the full range of possibilities permitted by the construction procedure. Now, as one imaginatively enacts this

rotation, one concurrently employs dynamic aspectual imagery, alternating between parsing the figure as a line rotating through the empty space above the horizontal line, and parsing the figure as two angles that continuously change in size, with one dilating and the other contracting, depending on the direction of imagined rotation. On the former parsing, it is obvious that the total space within which the line is rotating (the ‘straight-angled’ region about the meeting point on the top of the horizontal) remains invariant throughout; on the latter parsing, it is not immediately clear how the sum of the two angles might change across the visualized transformation. The crucial insight occurs when one maps these two alternative interpretations onto each other, thereby mapping the total space within which the line rotates onto the sum of the two angles. Immediately, one grasps that throughout the course of transformation, the sum of the two angles remains equal to the angle of the total space within which the line is rotating. Since in the initial placement of the line, the two angles are both right, it follows that the angles formed by *any* line set up on a horizontal are together equal to two right angles.

Second, the proof of the angle-sum theorem presupposes the result that in any intersection of two straight lines, the opposite (or ‘vertical’) angles will be equal. Here let’s begin with the special case depicted in Figure 5.3.

Figure 5.3: Demonstration of Euclid's Proposition I.15



We want to show that angles A and C are equal. One way to achieve this (which corresponds to Euclid's own proof) is by means of dynamic aspectual imagery. First we consider the grouping of angles A and B, noting that together they fill a straight line (and hence, by the previous result, are together equal to two right angles). Then we regroup, considering angles B and C together. Again, these angles together fill a straight line. By mapping these alternate parsing schemes to each other, and noting the occurrence of B in both groupings, we immediately see that A and C are equal in the depicted case. We now use dynamic transformational imagery to rotate one of the intersecting lines (it doesn't matter which one), confirming that the full range of rotation indeed yields the total set of figures that can be constructed under the description 'intersecting straight lines'. As we rotate the line, we continue to track our alternative groupings, and notice that across this range of transformation, the spatial relations that underwrite our conclusion in the special case continue to obtain. In particular, even as the angles dilate and contract, angle B continues to fill a straight

line with both angles A and C, for the reason that the intersecting lines remain straight throughout rotation. Again, by employing the two varieties of dynamic imagery in concert, a general theorem is grasped.


Now we are in a position to tackle the angle-sum theorem itself. It is best to begin by considering the special case depicted in Figure 5.1. As indicated previously, there are four separate spatial relations that one must attend to. In order to genuinely *perceive* the truth of the special case, one must cyclically consider these relations in turn, progressively mastering their interconnections, until one is able to grasp their combined force in (something approaching) a single intuition. I'll consider them in increasing order of difficulty. First, we must grasp that the three angles about point C on the upper side of line BD sum to two right angles. Once one has internalized the first preliminary result discussed above, this becomes perceptually obvious, since it is obvious that these three angles fill the space about a point on a straight line without remainder. Second, we must grasp that $\angle ACB$ remains an element across both of the relevant groupings of angles: the interior angles of the triangle ABC, on the one hand, and the angles that fill the space above the straight line BD about point C, on the other. This is accomplished by means of a relatively trivial application of dynamic aspectual imagery.

Third, we must grasp that angle $\angle ECD = \angle ABC$. In Euclid's proof, this follows from a prior theorem (Proposition 29). However, it is possible to perceive this relation in the diagram for angle-sum by employing the two kinds of dynamic imagery in combination. First, recall that seeing lines AB and CE as parallel consists in

rehearsing the rigid translation of line AB onto CE, as it slides along the ‘rail’ of line BC. If we now apply dynamic aspectual imagery to this imagined translation, we can reappraise the translation of the *line* AB (onto CE as its ‘destination’) as a translation of the *angle* $\angle ABC$ (onto the *angle* $\angle ECD$ as *its* ‘destination’), by flipping figure and ground as we imaginatively enact this translation. Note that this does not involve ‘eyeballing’ the two angles to see that they (approximately) correspond. Since it is stipulated (by construction) that lines AB and CE are parallel, we are justified in imagining the translation of the former into precise coincidence with the latter. The recognition that the angles correspond due to a corresponding rigid translation is parasitic on our confidence that the lines do; it is merely a result of reframing the initial translation of lines, and attending to the ‘negative space’.

Finally, we must grasp that $\angle ACE = \angle BAC$. Again, Euclid here appeals to Proposition 29 to establish this, but we can see that the relation holds just by considering the diagram for angle-sum, provided we have internalized the second of the preliminary results discussed above, so that opposite angles of intersections (of straight lines) have come to *look* equal, in a way that we feel confident trusting. Again, grasping the equality turns on our rehearsing the rigid translation of AB onto its parallel CE. In this case (in order to reduce ‘noise’), it will be useful to consider a diagram that displays just the relevant parts of the appropriately ‘animated’ figure (see Figure 5.4).

Figure 5.4: Animating the diagram for Euclid's Proposition I.32



The figure depicts the distinctive Z-shaped configuration formed by the parallel lines AB and CE, joined (at top and bottom, respectively) by the transversal AC. It also displays the translation of line AB onto CE, in virtue of which these lines are perceived as parallel. In order to grasp the relevant equality of angles, we can apply dynamic aspectual imagery to the range of 'transitional' figures that arise as the line AB, mid-translation, occupies intermediate positions at which it intersects with the transversal AC. If we attend to the angles formed around this intersection as the translation is imaginatively rehearsed, we see that the translation of the line AB along the 'rail' of BC can be reappraised as a translation of the *angle* $\angle BAC$ along the 'rail' of AC. (It is this reappraisal that is reminiscent of the ambiguous movement of the barbershop poles.) Since this interpretation results merely from reframing the translation of the line AB, we can be sure that the angle-in-translation remains constant throughout. Moreover, we can see that as soon as the translation from AB to CE begins, an angle arises that is *opposite* to the angle-in-translation $\angle BAC$ and

therefore (assuming the second preliminary result) equal to it. As the translation is enacted in visual imagination, this opposite angle remains entrained to $\angle BAC$, as the intersection as a whole translates along AC in the direction of point C . Finally, at the point where the translation is completed, and the intersection disappears into a coincidence with point C , the angle that was opposite the angle-in-translation $\angle BAC$ is seen to slide into coincidence with $\angle ACE$. Once this is observed, it becomes clear that angles $\angle BAC$ and $\angle ACE$ are equal.

Attending to these four relations in combination, then, is sufficient for grasping the truth of the angle-sum theorem for the special case depicted. In order to grasp that the result applies to the general case, one must consider the diagram of Figure 1 in light of the general construction procedure that gave rise to it. In order to capture the full range of figures that can be constructed according to this procedure, it suffices to imagine point A being moved freely throughout the space above line BC , while holding this line fixed. The resulting deformations of triangle ABC will yield a range of figures geometrically similar to all possible triangles. Since line CD is merely an extension of BC , which is held fixed, it too will remain stationary. Line CE , since it is stipulated to be parallel to AB , will ‘shadow’ the movements of the latter line as the entire figure undergoes deformations. In order to grasp that the result seen to hold for the special case is truly general, it will be necessary to attend to the structural bases for each of the four components that determine this result, as the range of deformations of triangle ABC is imagined by means of imagining the free movement of point A . In particular, one will need to observe that each of these

structural bases remains invariant across this range of deformation. Fortunately, it is perfectly acceptable to consider these piecemeal.

Observing that the first two relations hold across the relevant range of deformations is actually fairly trivial, since it suffices to observe that the overall structural topology of the figure remains invariant. In order to grasp the constancy of the first relation (the three angles filling the space above the horizontal line BD about point C), it suffices to observe that the line BD remains straight (which is trivial, since it is unaffected by the deformation), and that lines AC and CE continue to divide the space above BD into three separate angles. As one plays with the available positions of point A, noting the effect this has on the shape of the overall figure, it becomes clear that this latter condition only becomes threatened when point A approaches the axis along which lies line BC. As point A approaches this axis on the left side of point C, angle $\angle ACB$ threatens to disappear; as it approaches the axis on the right side of point C, angles $\angle ACE$ and $\angle ECD$ threaten to disappear. But since the construction procedure does not permit point A to meet this axis, the condition remains secure. This observation, in conjunction with the necessary (stipulated) persistence of triangle ABC across deformation, is similarly sufficient to guarantee that $\angle ACB$ remains available to be grouped alternately with this same set of angles, as well as with the interior angles of triangle ABC, so the second condition holds too.

The third relation, the equality of $\angle ECD$ and $\angle ABC$, can also be seen to persist across deformation, because it is grasped in the first place as a side-effect of the available translation of AB onto CE, along the 'rail' of BC. Now, BC remains

unaffected by deformation, as does the parallelism between AB and CE (since CE ‘shadows’ the movements of AB across the available deformations). The final relation is the equality of $\angle ACE$ and $\angle BAC$. As we saw above, this equality is grasped in virtue of observing the manner in which the translation of AB onto CE can be reappraised as the translation of an intersection, between the line-in-translation AB and the transversal AC, along the ‘rail’ of line AC in the direction of C—an intersection that resolves into the points A and C at the opposite resting points of the translation. If one repeatedly rehearses this translation in imagination (attending to the mapping between the alternative appraisals), it is not difficult to perceive that the successful ‘passing’ of one angle onto the other (*via* ‘reflection’ onto its opposing angle in mid-translation) rests on the structural feature of the Z-shaped configuration formed by parallel lines AB and CE, connected by the transversal AC, which meets AB at its top and CE at its bottom. For it is this Z-shaped configuration that underlies the passage of the intersection from points A to C as AB is translated to CE. By imaginatively playing with the position of point A, and observing how this deforms the configuration of lines AB, AC, and CE, it becomes clear that the Z-shape remains constant across the available range of deformation, since the lines AB and CE remain entrained in parallelism throughout, and also remain joined to the (shifting) transversal AC at the same endpoints (A, which is mobile, and C, which remains fixed). However this configuration is deformed, it can be seen to support the same intersection-based reappraisal of the translation of AB onto CE, which remains constantly available due to the ‘shadowing’ of AB by CE.

By attending to the way that these four spatial relations, whose joint appreciation is sufficient to grasp the truth of the angle-sum theorem in the special case, remain invariant across the range of deformation sufficient to capture all the possible triangles that can be depicted by figures drawn in accordance with the construction procedure, it is thereby possible to perceptually apprehend that the angle-sum theorem holds for *all* triangles, merely by a concerted deployment of the two kinds of dynamic imagery posited by our account, in relation to the diagram that lies at the core of Euclid's proof. I conclude that, by appealing to the imagistic resources of the dynamic imagery account, it is possible to explain how the generality problem can be solved for Euclid's angle-sum theorem.

References

- Appelle, S. (1972). Perception and discrimination as function of stimulus orientation. *Psychological Bulletin* 78: 266-278.
- Arnheim, R. (1974). *Art and Visual Perception*. 2nd edition. University of California Press.
- Avigad, J., Dean, E., and Mumma, J. (2009). A Formal System for Euclid's Elements. *Review of Symbolic Logic* 2: 700-768.
- Azzouni, J. (2013). That We See That Some Diagrammatic Proofs Are Perfectly Rigorous. *Philosophia Mathematica* (III) 21, 323–338.
- Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy* 70: 661-679.
- Berkeley, G. (1988). *A Treatise Concerning the Principles of Human Knowledge*. Penguin.
- Brown, J. (1999). *Philosophy of Mathematics: An Introduction to the World of Proofs and Pictures*. Routledge.
- Brown, J. (2010). Seeing things in pictures. In C. Abell and K. Bantinaki (eds.), *Philosophical Perspectives on Depiction*, Oxford University Press.
- Burgess, J. (2015). *Rigor and Structure*. Oxford University Press.
- Chudnoff, E. (2012). Presentational phenomenology. In Miguens & Preyer (eds.), *Consciousness and Subjectivity*. Ontos Verlag.
- Dehaene, S., Izard, V., Pica, P., and Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science* 311: 381-384.

- Dennett, D. (1991). *Consciousness Explained*. Back Bay Books.
- Dillon, M., Huang, Y., Spelke, E. (2013). Core foundations of abstract geometry. *Proceedings of the National Academy of Sciences* 110 (35): 14191-14195.
- Elder, J. and Goldberg, R. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision* 2(4), 324–353.
- Elgin, C. (2007). Understanding and the facts. *Philosophical Studies* 132, 33-42.
- Euclid. (1956). *The Thirteen Books of the Elements*. 2nd edition. Trans. T. Heath. Dover Publications.
- Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research* 37(20): 2835–2848.
- Ferrante, D., Gerbino, W., and Rock, I. (1997). The right angle. In I. Rock (ed.), *Indirect Perception*, MIT Press.
- Fischbein, H. (2002). *Intuition in Science and Mathematics*. Springer.
- Gardiner, G. (2012). Understanding, integration, and epistemic value. *Acta Analytica* 27 (2): 163-181.
- Giaquinto, M. (1993). Diagrams: Socrates and Meno's slave. *International Journal of Philosophical Studies* 1: 81-97.
- Giaquinto, M. (1998). Epistemology of the obvious: a geometrical case. *Philosophical Studies* 92: 181-204.
- Giaquinto, M. (2005). Symmetry perception and basic geometrical knowledge. In K. Joergensen and P. Mancosu (eds.), *Visualization, Explanation and Reasoning Styles in Mathematics*, Kluwer Academic Publishers.

- Giaquinto, M. (2007). *Visual Thinking in Mathematics: An Epistemological Study*.
Oxford University Press.
- Giaquinto, M. (2011). Crossing curves: a limit to the use of diagrams in proofs.
Philosophia Mathematica (III) 19(3): 281-307.
- Gilovich, T., Griffin, D., and Kahneman, D., Eds. (2002). *Heuristics and Biases: The
Psychology of Intuitive Judgment*. Cambridge University Press.
- Goldmeier, E. (1937). Similarity in visually perceived forms. *Psychological Issues* 8:
monograph 29.
- Grimm, S. (2012). The value of understanding. *Philosophy Compass* 7: 103-117.
- Hahn, H. (1933). The Crisis in Intuition. In *Empiricism, Logic and Mathematics*
(1980). Springer.
- Hilbert, D. (1910). *Foundations of Geometry*. (E. Townsend, trans.) Open Court
Publishing.
- Izard, V. and Spelke, E. (2009). Development of sensitivity to geometry in visual
forms. *Human Evolution* 23(3): 213-248.
- Izard, V., Pica, P., Spelke, E., and Dehaene, S. (2011). Flexible intuitions of
Euclidean geometry in an Amazonian indigene group. *Proceedings of the
National Academy of Sciences* 108 (24): 9782-9787.
- Kahneman, D. and Treisman, A. (1984). Changing views of attention and
automaticity. In R. Parasuraman and D. Davies (eds.), *Varieties of Attention*,
Academic Press.

- Kahneman, D., Treisman, A., and Gibbs, B. (1992). The reviewing of object files: object-specific integration of information. *Cognitive Psychology* 24: 174-219.
- Kant, I. (1998). *Critique of Pure Reason*. (P. Guyer and A. Wood, trans.) Cambridge University Press.
- Kitcher, P. (1984). *The Nature of Mathematical Knowledge*. Oxford University Press.
- Kosslyn, S. (1994). *Image and Brain: The Resolution of the Imagery Debate*. MIT Press.
- Leibniz, G. (1704). *New Essays on Human Understanding*.
- Locke, J. (1979). *An Essay Concerning Human Understanding*. Oxford University Press.
- Macbeth, D. (2010). Diagrammatic reasoning in Euclid's Elements. In B. Kerkhove, J. Vuyst, and J. Bendegem (eds.), *Philosophical Perspectives on Mathematical Practice*. College Publications.
- Mach, E. (1897). *The Analysis of Sensations*.
- Mancosu, P. (2005). Visualization in logic and mathematics. In P. Mancosu, K. Jorgensen, and S. Pederson (eds.), *Visualization, Explanation and Reasoning Styles in Mathematics*. Synthese.
- Manders, K. (2008). The Euclidean diagram. In P. Mancosu (ed.) *The Philosophy of Mathematical Practice*. Oxford University Press.
- Norman, J. (2006). *After Euclid: Visual Reasoning and the Epistemology of Diagrams*. CSLI.

- Palmer, S. (1983). The psychology of perceptual organization: a transformational approach. In J. Beck, B. Hope, and A. Rosenfeld (eds.), *Human and Machine Vision*. Academic Press.
- Palmer, S. (1985). The role of symmetry in shape perception. *Acta Psychologica* 59: 67–90.
- Palmer, S. and Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance* 4: 691-702.
- Peacocke, C. (1992). *A Study of Concepts*. MIT Press.
- Plato. (1992). *Republic*. Trans. G.M.A. Grube. Hackett.
- Plato. (2002). *Five Dialogues*. Trans. G.M.A. Grube. Hackett.
- Rock, I. (1973). *Orientation and Form*. Academic Press.
- Rock, I. (1997). Orientation and form. In I. Rock (ed.), *Indirect Perception*, MIT Press.
- Rock, I. and Leaman, R. (1997). In I. Rock (ed.), *Indirect Perception*, MIT Press.
- Scott, D. (2005). *Plato's Meno*. Cambridge University Press.
- Shabel, L. (2003). *Mathematics in Kant's Critical Philosophy: Reflections on Mathematical Practice*. Routledge.
- Shepard, R. and Metzler, D. (1971). Mental rotation of three-dimensional objects. *Science* 171: 701-3.
- Shusterman, A., Lee, S., and Spelke, E. (2008). Young children's spontaneous use of geometry in maps. *Developmental Science* 11 (2): F1-F7.

- Spelke, E. and Lee, S. (2012). Core systems of geometry in animal minds. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367: 2784-2793.
- Spelke, E., Lee, S., and Izard, V. (2010). Beyond core knowledge: natural geometry. *Cognitive Science* 34: 863-884.
- Suppes, P. (1977). Is visual space Euclidean? *Synthese* 35: 397-421.
- Treder, M. and van der Helm, P. (2007). There is no symmetry like orthogonal symmetry. *Vision* 7 (9): article 764.
- Tversky, A., and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, New Series, Vol. 185, No. 4157, 1124-1131.
- Von Koch, H. (1906). Une méthode géométrique élémentaire pour l'étude de certaines questions de la théorie des courbes planes. *Acta Mathematica* 30: 145–174.
- Wagner, M. (2006). *The Geometries of Visual Space*. Lawrence Erlbaum.
- Walton, K. (1970). Categories of art. *Philosophical Review* 79 (3): 334-367.
- Wenderoth, P. (1994). The salience of vertical symmetry. *Perception* 23 (2): 221-236.
- Wollheim, R. (1980). Seeing-as, seeing-in, and pictorial representation. In *Art and Its Objects*, 2nd edition. Cambridge University Press.
- Yuille, A., Fang, F., Schrater, P., and Kersten, D. (2004). Human and ideal observers for detecting image curves. In S. Thrun, L. Saul, and B. Schoelkopf (eds.), *Advances in neural information processing systems* (vol. 16), pp. 59–70. MIT Press.

Zwicky, J. (2006). Mathematical analogy and metaphorical insight. *Mathematical
Intelligencer* 28: 4-9.

Zwicky, J. (2009). *Plato as Artist*. Gaspereau Press.