

THOUGHT EXPERIMENTS IN MATHEMATICS

Irina Starikova¹ and Marcus Giaquinto

It is not news that we often make discoveries or find reasons for a mathematical proposition by thinking alone. But does any of this thinking count as conducting a thought experiment? The answer to that question is “yes”, but without refinement the question is uninteresting. Suppose you want to know whether the equation $[8x + 12y = 6]$ has a solution in the integers. You might mentally substitute some integer values for the variables and calculate. In that case you would be mentally trying something out, experimenting with particular integer values, in order to test the hypothesis that the equation has no solution in the integers. Not getting a solution first time, you might repeat the thought experiment with different integer inputs.

The fact that there are such mundane thought experiments is no surprise and does not answer the question we are really interested in.² The numerical thought experiment just given involves nothing more than applying mathematically prescribed rules (such as rules of substitution and calculation) to selected inputs. It would be more interesting if there were mathematical thought experiments in which the experimental thinking goes beyond application of mathematically prescribed rules, by using sensory imagination as a way of eliciting the benefits of past perceptual experience.³ In what follows we will try to show that there are such thought experiments and to assess their epistemic worth.

Our method will be to present some candidate thought experiments with what we hope is enough background explanation and in sufficient detail for you, the reader, to perform the relevant mental operations yourself; without this participation the paper will be neither convincing nor engaging. We have tried to avoid run of the mill

¹ I would like to thank the Brazilian Coordination for the Improvement of Higher Education Personnel (CAPES) and the Russian Foundation of Basic Research (RFBR).

² For this reason we find that the category of thought experiments as characterised by Jean-Paul Van Bendegem in “Thought experiments in mathematics: anything but proof” *Philosophica* **72**, (2003), pp. 9-33 to be too broad.

³ For a different focus, see Eduard Glas, “On the role of thought experiments in mathematical discovery” in J. Meheus and T. Nickles (eds.), *Models of Discovery and Creativity*, (Springer 2009). Glas says that “imagery, mental or experiential, is not essential” to the aspect of thinking that he counts as thought experiment (even when accompanied by imagery). For this reason, the kinds of thinking that we discuss in this paper do not fall under what Glas counts as thought experiment.

examples by staying out of universally familiar mathematical areas; but to keep the material accessible, the examples are mathematically quite simple, with something a bit more advanced reserved for the end. The paper has three main parts, corresponding to the mathematical areas from which the examples are drawn: knot theory, graph theory and geometric group theory. In the last two parts later exposition depends on earlier; so the material is best read in the order presented.

1. CANDIDATES FROM KNOT THEORY


Preliminaries

For the examples to be intelligible, some background about knots in mathematics is needed. Here it is with a minimum of technical detail.

A *knot* is a tame closed non-self-intersecting curve in Euclidean 3-space.

The word “tame” here stands for a property intended to rule out certain pathological cases, such as curves with infinitely nested knotting. Knots are just the tame curves in Euclidean 3-space which are homeomorphic to a circle.⁴ In Figure 1 on the left is a diagram of a knot and on the right a pathological case.

Figure 1



A knot has a specific geometric shape, size and axis-relative position, but if it is made of suitable material, such as flexible yarn that is stretchable and shrinkable, it can be transformed into other knots without cutting or gluing. Since our interest in a knot is the nature of its knottedness regardless of shape, size or axis-relative position, the real focus of interest is not just the knot but all its possible transforms. A way to think of this is to imagine a knot transforming continuously, so that every possible transform is realized at some time. Then the thing of central interest would be the object that persists over time in varying forms, with knots strictly so called being the things captured in each particular freeze frame. Mathematically, we represent the relevant entity as an equivalence class of knots.


Two knots are *equivalent* iff one can be smoothly deformed into the other by stretching, shrinking, twisting, flipping, repositioning or in any other way that

⁴ We are setting aside higher dimensional knot theory.

does not involve cutting, gluing, passing one strand through another or eliminating a knotted part by shrinking it down to a point.⁵

In practice equivalent knots are treated as the same, with a knot strictly so called regarded as just one of the forms a knot can take. This practice will be followed here. More precisely, the word 'knot' without the qualification 'strict' will be used to refer to an equivalence class of strict knots. Figure 2 presents diagrams of the same knot.

Figure 2



Diagrams like these are not merely illustrations; they also have an operational role in knot theory. But not any picture of a knot will do for this purpose. We need to specify:

A knot diagram is a regular projection of a strict knot onto a plane (as viewed from above) which, where there is a crossing, tells us which strand passes over the other.

Regularity here is a combination of conditions. In particular, regularity entails that not more than two points of the strict knot project to the same point on the plane, and that two points of the strict knot project to the same point on the plane only where there is a crossing.


A knot diagram with one or more crossings tells us at each crossing which strand passes over the other, but it does not tell us how far above the other it goes. So distinct strict knots can have the same knot diagram. But this does no harm, because strict knots with the same knot diagram are equivalent. This is all the background we need in order to proceed to examples.

⁵ There are mathematically precise definitions of knot-equivalence. It is clearly not enough to say that equivalent knots are homeomorphic, as all knots are homeomorphic to the circle hence to each other. They are equivalent iff there is an ambient isotopy taking one to the other. More about that shortly.

A thought experiment with knots


An important and obvious fact is that a knot has many knot diagrams. As we represent knots by knot diagrams, a major task of knot theory is to find ways of telling whether two knot diagrams are diagrams of the same knot. In particular we will want to know if a given knot diagram is a diagram of *the unknot*, which is the only knot representable by a knot diagram without crossings. To warm up, here are some exercises. Using your visual imagination on the two knot diagrams in Figure 3, see if you can tell whether either is a diagram of the unknot.

Figure 3




In fact it is not possible to deform the knot represented on the left so that the result is a diagram without crossings, but you will probably have no difficulty with the one on the right. Figure 4 indicates a simple way.

Figure 4



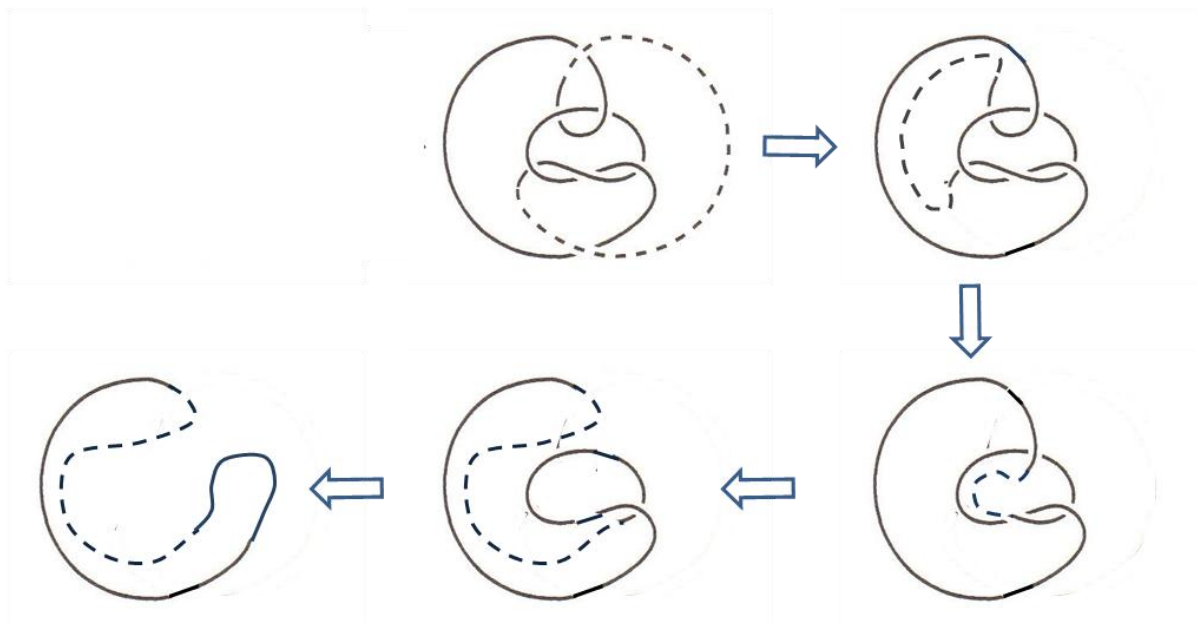
Before considering what you can reasonably conclude from the results of your efforts, try to visualize deforming the knot represented by this more complicated knot diagram, Figure 5, to get a diagram without crossings.

Figure 5



It can be done, but it is difficult without actually producing physical diagrams representing the knot at one or more intermediate stages of the complete deformation. To conduct this thought experiment one performs one or more trials, a trial being a finite sequence of steps, each of which consists of (a) visualizing a deformation in 3-space of a knot as represented by one seen diagram and (b) drawing (or otherwise producing) another knot diagram corresponding to the projection of the knot at the end of the visualized deformation so far. The experiment has a positive outcome when one of the trials ends with a diagram which has no crossing. Figure 6 illustrates the intermediate stages of a successful trial for this case. The dashed section of each diagram indicates the part about to be moved or the part just moved.

Figure 6



Assessment and objections

One charge laid against so-called thought experiments in physics is that they are not experiments at all, but ‘merely picturesque arguments’ for a claim that is already believed.⁶ Is that true of the process of visual imagining and diagram making just described? Picturesque it may be. But one may embark on the process in order to find out whether the knot represented by the initial diagram in Figure 6 (matching Figure 5) is an unknot, lacking any conviction either way; one may even doubt that it is the unknot. So one may really be experimenting, in the sense of performing a series of actions in order to test the hypothesis that the original diagram is a diagram of the unknot.

Granting that it is a genuine experiment, another worry concerns its epistemic value. Can we really make discoveries this way, relying heavily on visual imagination? Are we not simply replacing proper experiments by ‘fantasies of the imagination’⁷? There is an important distinction between veridical imagining and fantastical imagining.⁸ Veridical imagining is aimed at finding out the true answer to some question, and is constrained by the accumulated effects of past perceptual experience. Of course this does not make veridical imagining infallible; the adjective “veridical” is intended to describe an aim, not a result, of imagining. Fantastical imagining is not constrained in that way, as it is not aimed at answering a question, but serves psychological ends such as wish-fulfilment, horror thrill, sensory fascination and so on. Veridical imagining is common and useful. Wanting a desk for a particular room, you visit a furniture showroom; there you see an attractive desk, somewhat bigger than you had in mind. Would it fit reasonably well into the room with its other furniture? In this situation it is reasonable to visualize the room to reach a judgement. This is veridical imagining. Other examples readily come to mind: Can I prepare a tolerable evening meal from the ingredients I now have at my disposal? Can five normal adults sit comfortably in my car? We do in fact rely on sensory imagination to answer such questions, and we get correct answers frequently enough for this practice to persist.

⁶ Norton 1996. “Are thought experiments just what you thought?” *Canadian Journal of Philosophy* **26** (3), p.333-366.

⁷ Norton 1996.

⁸ Articulated by Paul Boghossian in a New York Institute of Philosophy workshop on the a priori in June 2013.

This still leaves open the question of epistemic value in *this* case. Here are two questions we need to answer:

- (1) Can our visual imagination be sufficiently reliable here?
- (2) How do we reliably reach a conclusion about a mathematical question from information about physical situations?

The question of reliability is a serious one when trying to imagine deformations starting from a complicated knot diagram. But in our case the complexity is quite small and there is no real worry. Let us call a maximal part of a knot diagram between undercrossings an *arc*. Then the first step in Figure 6 involves flipping the rightmost arc over the central part of the diagram and shrinking it until it falls just within the leftmost arc. This clearly preserves knot identity (up to equivalence, of course). The remaining steps are clearly identity preserving atomic moves known as Reidemeister moves.⁹ When one makes a non-atomic move in getting from one diagram to another, as in the first step of Figure 6, one can check its permissibility by seeing if one can break it down into a sequence of Reidemeister moves. Returning to the rightmost diagram of Figure 3, it is easy to see that (with two Reidemeister moves) it can be turned into a diagram without crossings, as shown in Figure 4. So the answer to the first question is: yes, a person's visual imagination can be (and often is) sufficiently reliable for this task. Your own experience in following the examples should provide you with supporting evidence.

In these cases we are visualizing a physical possibility, at least partly based on experience with string, yarn, cotton thread or suchlike. How do we get from an empirical discovery of a physical possibility to a mathematical possibility? The highly informal way in which the subject has been presented here hides the conceptual distance between the physical thought and the corresponding mathematical proposition. In these cases the conclusion is that a strict knot which projects the starting diagram is ambient isotopic to a strict knot which projects a diagram without crossings. But what does "ambient isotopic" mean? To get a sense of the full mathematical content of such a claim, note first that a strict knot is mathematically identified with a homeomorphism γ from the unit circle S^1 into R^3 (not the image of S^1

⁹ Every introductory text on knot theory and some more advanced texts define the Reidemeister moves, and they can be readily found on the web. For an introduction see Colin Adam's *The Knot Book*, American Mathematical Society 2001.

under the homeomorphism), with an additional condition on the homeomorphism to rule out wild knots¹⁰. Ambient isotopy is defined as follows:

Strict knots γ_0 and γ_1 are *ambient isotopic* iff there is a continuous map

$F: \mathbb{R}^3 \times [0,1] \rightarrow \mathbb{R}^3$ such that for each r in $[0,1]$, $F(x, r)$ is a homeomorphism of \mathbb{R}^3 , $F(x, 0)$ is the identity map on \mathbb{R}^3 , and $F(x, 1) \cdot \gamma_0 = \gamma_1$.

This definition makes clear that visualizing *alone* does not enable us to discover a full mathematical fact expressed in saying, of two strict knots, that they are ambient isotopic. This is because one cannot tell by visualizing alone that there is a continuous map fulfilling the stated conditions. But we have been assuming that visualizing can make it reasonable to believe the mathematical claim and lead to discovery. How is this possible?

The answer is that these visual thought experiments take place in the context of background knowledge about the links between the mathematical definitions and idealised physical objects and transformations that can be visualized. These links belong to what is referred to as the foundational aspect of knot theory, and often expositions of the foundations reveal that the mathematical definitions are tailored to represent the intended kind of visualizable objects and transformations. Sometimes promising definitions are put forward only for the sake of showing their inadequacy for representing the intended visualizable material, before proper definitions are given.¹¹ Moreover, mathematically inequivalent definitions of tame knots are given in different texts, but it is known that each adequately represents what is intended, much as real numbers can be defined as Dedekind cuts of rational numbers or as Cauchy equivalence classes of Cauchy convergent sequences of rational numbers.

For foundational purposes there needs to be some way of fixing the subject matter in mathematical terms, so that the correctness of basic assumptions and methods can be proven. But once that job has been done, we may proceed without adverting to our foundational definitions. This is the situation with regard to basic knot theory. The

¹⁰ Even here there is some oversimplification. First, the homeomorphism is usually required to preserve orientation (which has not been defined here), to avoid identifying chiral knots with their mirror images. Also, for the advantages of operating in a compact space the co-domain of the homeomorphism is usually taken to be S^3 instead of \mathbb{R}^3 .


¹¹ See for example Josh Greene, Combinatorial Methods in Knot Theory, Lecture 1: Foundations, January 2013. <https://www2.bc.edu/joshua-e-greene/MT885S13/Lecture%201.pdf>

foundational definitions are needed for proving Reidemeister's Theorem: two strict knots are equivalent if and only if there is a finite sequence of Reidemeister moves taking a diagram of one to a diagram of the other. Once that has been established we can go a long way with visual thought experiments.

Another example: the trefoil and the unknot

Let us return to the leftmost diagram of Figure 3, reproduced in Figure 7.

Figure 7



This is a trefoil knot. If you have been visualizing properly your attempts to visualize a deformation of this trefoil so that it projects a diagram without crossings will have been unsuccessful. After a few trials you may have become convinced that this trefoil is not the unknot. The diagram is cognitively quite simple. So, unless your visual imagination is poor, a few negative trials provides evidence that the trefoil is not the unknot.

For more conclusive evidence, we can use a knot invariant known as colourability¹².

A knot diagram is *colourable* if and only if each of its arcs can be coloured one of three different colours so that (a) at least two colours are used and (b) at each crossing the three arcs are all coloured the same or all coloured differently.

Colourability is a knot invariant in the sense that if one diagram of a knot is colourable every diagram of that knot is colourable.¹³ This fact can be proved using

¹² Colourability is sometimes called 'tricolourability'.

¹³ There is a combinatorial version of colourability. If instead of colouring the arcs one labels them 0, 1 or 2, the colourability conditions together have a numerical equivalent: (a) at least two of the numerical labels are used and (b) at each crossing if x is the value of the overcrossing arc and y and z are the values of the other two arcs, $2x - y - z = 0 \pmod{3}$. This mod 3 labelling readily generalises to other invariants known as 'mod p labelling', where p is an odd prime.

Reidemeister's theorem. Since any diagram of a knot can be reached from any other diagram of that knot by a finite sequence of Reidemeister moves, to prove the invariance of colourability it suffices to show that if a Reidemeister move is performed on a colourable knot diagram the resulting diagram is again colourable.

A standard diagram of the unknot, a diagram without crossings, is clearly not colourable because it has only one arc (the whole thing) and two colours cannot be used. So in order to show that the trefoil is distinct from the unknot, it suffices to show that the trefoil diagram is colourable. So here is a thought experiment to test the hypothesis that the trefoil represented in Figure 7 is colourable: while looking at the diagram, visualize each of the arcs as coloured red, green or blue using at least two colours; alternatively, when looking at the diagram mentally label each arc with one of the words "red", "green" or "blue" using at least two of them.¹⁴ Then check that at each crossing all three arcs have the same colour or all three have different colours.

Because the trefoil diagram of Figure 7 is visually so simple, this thought experiment can be carried out reliably, thereby giving the thinker very strong reason to believe that the trefoil is colourable, as in fact it is, hence not equivalent to the unknot. With more complicated diagrams, it is difficult to hold the relevant information in visual imagination, and one is forced to colour or label arcs on the page or screen and then check that the conditions are met.

¹⁴ This can be done either by visually imagining a written colour word placed next to an arc or, just as easily, by aurally imagining uttering a colouring word as a label of an arc one is visually attending to.

2. EXAMPLES WITH GRAPHS

Cycle graphs

We often represent mathematical objects by a configuration of dots connected by line segments, such as a tree or a cycle. This gives rise to the algebraic notion of a graph G which consists of a set V_G , the 'vertices' of G , and a set E_G of pairs of members of V_G , the 'edges' of G .¹⁵ We will be concerned with cycle graphs:

G is a *cycle graph* iff $V_G = \{v_1, v_2, \dots, v_k\}$ for $k \geq 3$ and every edge in E_G occurs just once in the sequence $\{v_1, v_2\}, \{v_2, v_3\}, \dots, \{v_n, v_{n+1}\}, \dots, \{v_k, v_1\}$.

A cycle graph has an obvious representation as a regular polygon; there are just as many edges as vertices. The spatial representation of graphs makes us notice not only kinds of graphs, but also various graph-theoretic properties and relations. The following are relevant examples.

A *path* between vertices u and v is a non-empty sequence of edges $\{y_1, y_2\}, \{y_2, y_3\}, \dots, \{y_{n-2}, y_{n-1}\}, \{y_{n-1}, y_n\}$, with the y_j distinct, and $u = y_1$ and $v = y_n$.

A graph is *connected* iff between any two of its vertices there is a path.

The *length* of a path is the number of edges in the path.

For connected graphs we have the following notions of distance and diameter:

This *distance* between two vertices u and v , $d(u, v)$ = the length of a shortest path between u and v .

The *diameter* of a graph is the maximum distance between vertices, i.e.

$$\max \{d(x, y) : x, y \in V_G\}.$$

With these definitions at our disposal, we can proceed to our initial example.


Suppose we want to express the diameter of a cycle graph with n vertices in terms of n . A thought experiment can help us here. Imagine the vertices of the graph to be small but heavy pearls of equal size and weight, adjacent pearls connected by a fixed unit length of strong flexible thread, like a necklace. Then imagine holding the necklace by any one pearl, letting the rest of it go. What will happen? The rest will

¹⁵ Strictly speaking, E_G is a multiset, so that an element $\{u, v\}$ can occur more than once, the number of occurrences being the number of edges with endpoints u and v . Those edges with endpoints u and v are said to be *parallel* to one another. Also, there can be one or more edges $\{u, u\}$, known as *loops*. Cycle graphs are *simple* graphs, in the sense that they have no loops and no two edges are parallel.

fall as far as the thread will let it; so the maximum distance between the top pearl (the held pearl) and any other will be the number of units of thread (representing edges) between the top pearl and a lowest pearl. What if we hold the necklace by any other pearl, say k units of thread further on? As the necklace is a cycle, by visualizing what happens each time we rotate the necklace by a unit, we can tell that the configuration made by the hanging necklace remains unchanged, and so the distance between the new top pearl and the new lowest pearl (or pearls) will be the same. So the number of units of thread between the top pearl and a lowest pearl is the maximum distance between pearls. This represents the diameter of the graph.

But what is this number, for a given number n of unit threads in the whole necklace? The thought experiment continues. We now visualize the dangling necklace with fine-grained attention to discover its form. At the top is a single pearl with its two neighbouring pearls hanging next to each other at a distance of one unit below the top; if there are at least two more pearls, the next pair of pearls will hang at distance of two units from the top; if there are at least two more, the next pair will hang at a distance of three from the top, and so on. That will be the same whether the number of pearls is even or odd. Now visualize the lowest few pearls of the dangling necklace. How will they be arranged? If the number is odd, below the top pearl the remaining pearls will hang in pairs, the lowest pair having a unit of thread connecting them, illustrated on the left in Figure 8. If the number is even, at one unit below the top pearl will be one pair of pearls, at one unit below them another pair of pearls, and so on until we run out of pairs and just one pearl remains (as the total number of pearls is even). By visualizing attentively the bottom of the image necklace in this situation, one can discern that this lowest single pearl will be connected by unit threads to each of the lowest pair of pearls just above it, as on the right in Figure 8.

Figure 8



To make use of these results, we reason as follows. Let the total number of unit threads (edges) be n . If n is even, we notice that there are two equal length paths from top to bottom; so we merely need to divide by 2. If n is odd, noticing that the thread between the bottom pair of pearls does not belong to any path between top and bottom pearls, we subtract that one thread from the total; then we can notice that there are two equal length paths from the top to either bottom pearl; so we only need to divide the remaining $n-1$ edges by 2. Either way we get $\lfloor n/2 \rfloor$, the greatest integer $\leq n/2$. So here is a discovery one can make with the help of a thought experiment: the diameter of a cycle graph with exactly n vertices is $\lfloor n/2 \rfloor$.

To assess this candidate thought experiment, the two questions we need to answer are: (1) Is what we have called a thought experiment in this case really an experiment, as opposed to a picturesque argument for a claim already believed? (2) Given that it is an experiment with certain outcomes, is it a reliable way of getting those outcomes?

The relevant mental actions here are (a) visualizing the cycle graph as a physical object, the necklace of pearls, and then visually imagining the result of holding the necklace by one pearl while letting go of the rest of it, (b) visualizing what happens to the configuration as we change top pearls, going from one pearl to an adjacent pearl, and (c) visualizing the spatial forms of the result of letting the necklace dangle, for odd and even numbers of pearls, with special attention to the top and the bottom. It seems right to say that parts (a) and (b) are unrevealing: we already know that the necklace will dangle, pulled down by gravity as far as the connecting thread will allow, and that changing the pearl by which the necklace is held will leave the configuration unchanged. Part (c), however, may be revealing. Finer-grained imagery results from taking into account the parity information. The forms of the lower end of the necklace in the two cases are revealed to us by the visualizing. So (a) followed by (b) and then (c) constitutes a series of mental actions, not to test a hypothesis, but to find the forms the necklace would take. That is a thought experiment. The results of the experiment are the forms we find; we use them as input to further thinking leading to our mathematical conclusion.

Is our visual imagination reliable here? If you agree that we would get the same results if we performed the experiment physically with actual necklaces matching the description in the thought experiment, you should accept that this use of visual imagination is reliable.¹⁶ This is not surprising: our visual experience of physical situations relevantly similar to the described situation is sufficiently extensive to produce reliable dispositions for veridical imagining. Mathematically the result is quite trivial. For something a bit more interesting mathematically we focus on Cayley graphs.

Cayley graphs

Cayley graphs are representations of groups with a finite set of generators. Recall that a *group* is a set G together with a binary function $x \bullet y$ satisfying exactly these conditions:

Closure: For all x, y in G , $x \bullet y$ is in G .

Associativity: For all x, y, z in G , $(x \bullet y) \bullet z = x \bullet (y \bullet z)$.

Identity: For some z in G , for every x in G , $x \bullet z = x = z \bullet x$

Inverse: For any z in G satisfying the identity condition, for every x in G there is a y in G such that $x \bullet y = z = y \bullet x$.

It is easy to prove from the identity condition that there is just one identity, often denoted e . It is easy to prove from the inverse condition that each member x of G has just one inverse, denoted x^{-1} . As \bullet is associative, we can omit brackets and the function symbol and use juxtaposition instead. This improves legibility.¹⁷

Let S be a finite subset of G . S is a set of *generators for* G iff every member of G is the product of a finite sequence of members of S or their inverses. More formally, putting S^{-1} for the set of inverses of members of S , this is:

For every x in G , there are y_i ($1 \leq i \leq n$) in $S \cup S^{-1}$ such that $x = y_1 y_2 \dots y_n$.

¹⁶ We clearly could perform this experiment physically as well as in thought; the same goes for the thought experiments on knots illustrated in Figures 4 and 5. This refutes Buzzoni's claim that a mathematical thought experiment 'leaves no room for a separate real performance of the experiment.' 'On Mathematical Thought Experiments' *Epistemologia* XXXIV (2011), pp. 61-88.

¹⁷ So for example, we write $ab^{-1}aab$ for $(a \bullet (b^{-1} \bullet a)) \bullet (a \bullet b)$.

In this case, $((G, \bullet), S)$, usually written simply (G, S) , is a *generated group*. Here are some examples of finitely generated groups:

- Let n be an integer greater than 2. The domain of the group is the set C_n of rotations of a regular n -sided polygon about its centre by $k2\pi/n$ radians for integers k . The function \bullet is composition. Let anticlockwise rotation by $2\pi/n$ be the sole generator.
- The set of integers \mathbb{Z} under addition, with generator 1.
- The set S_3 of permutations of a triple $\{a, b, c\}$ under composition, with generators $\{r, f\}$, where r (for 'rotation') takes $\langle a, b, c \rangle$ to $\langle c, a, b \rangle$, and f (for 'flip') takes $\langle a, b, c \rangle$ to $\langle c, b, a \rangle$.¹⁸

Some groups are not finitely generated. An example is the set \mathbb{Q} of rationals under addition.¹⁹ Many finitely generated groups have different sets of generators. For example, C_5 is generated by $\{2\pi/5\}$; it is also generated by $\{4\pi/5\}$. The group $(\mathbb{Z}, +)$ is generated by $\{1\}$; it is also generated by $\{2, 3\}$.²⁰

Cayley graphs represent finitely generated groups in the following way: each group member is represented by a unique vertex, and each vertex represents exactly one group member; for any group member g and generator s there is a directed edge from the vertex representing g to the vertex representing sg .²¹

Let us look at some examples. A graph for C_6 is suggested by the general geometrical description of C_n given above. Put c for the generator, anticlockwise rotation by $2\pi/6$ radians; for $k > 0$, put c^k for this operation repeated k times, that is, anticlockwise rotation by $k2\pi/6$ radians, and c^{-k} for clockwise rotation by $k2\pi/6$

¹⁸ S_3 is also the group of symmetries of an equilateral triangle. If we take a, b , and c to be vertices of an equilateral triangle, r (rotation) and f (flip) are obvious operations (symmetries) of the triangle.


¹⁹ Take any finite set of rationals with denominators d_1, d_2, \dots, d_n . Any sum of those rationals could be expressed as a rational n/m with denominator $m = d_1 \times d_2 \times \dots \times d_n$. But not all rationals could be so expressed: consider p/q where p and q are primes and $q > m$.

²⁰ To see that $\{4\pi/5\}$ generates C_5 , note that a sequence of three anticlockwise rotations by $4\pi/5$ = anti-clockwise rotation by $2\pi/5$; two anticlockwise rotations by $4\pi/5$ = *clockwise* rotation by $2\pi/5$. To see that $\{2, 3\}$ generates \mathbb{Z} note that $+3+(-2) = +1$ and $+2+(-3) = -1$.

²¹ As the group operation is so often function composition, we maintain the convention that s applied to g (i.e. s after g) is sg , so that an edge goes from g to sg .


radians. Put e for the identity, that is, anticlockwise rotation by 2π . We can represent this generated group as in Figure 9.

Figure 9



As the generated group $(\mathbb{Z}, +, \{1\})$ is infinite, we can only show part of its graph, as in Figure 10; but it is obvious how it continues. The identity e is 0, the sole generator is 1 and any integer n results from adding 1 or -1 $|n|$ times.

Figure 10



Groups with two or more generators have more complicated structures than their single generator counterparts. This can be seen by comparing diagrams of their Cayley graphs. Here are a couple of examples. Figure 11 is a diagram of the Cayley graph of $(\mathbb{Z}, +, \{2, 3\})$. We use different colours for composition with the different generators, black for an edge from n to $n+2$ and red for an edge from n to $n+3$.

Figure 11






Figure 12 depicts the graph of S_3 with generators r and f . We use red for an edge from x to rx and black for an edge from x to fx . Also, for each generator s which is its own inverse, there will be two edges between adjacent vertices, one from x to sx , the other from sx to x , as $x = ssx$. In this case it is visually easier to read the image if we merge the two edges using arrowheads both ways. We do this for edges between x and fx , as f (= flip) is its own inverse.

Figure 12



Although in practice we often ignore the difference between the visual image and the graph, they are not the same, as there can be visually divergent images of the same graph. The lines could be arcs of a circle, for instance; or the shape, size and positioning of polygonal faces can be changed without changing the graph. Figure 13 for example shows the graph of S_3 with generators $\{r, f\}$ as a prism (without labels).

Figure 13



The Cayley graph of a finitely generated group represented by these diverse images is a graph-theoretic object, not a drawing. We can be precise about this. Let G be a

group generated by a finite set S . The *Cayley graph* of (G, S) is the graph (V, E) with $V = G$ and $E =$ the set of ordered²² pairs $\langle x, sx \rangle$ for x in G and s in S .

Why bother with the visible diagrams of Cayley graphs? Because they can help us grasp the nature of the Cayley graphs they represent; they can help us reason about them; they can suggest to us hypotheses about them; they can help us to discover or explain facts about them.²³

Here is a simple example. Recall that a single line segment with arrowheads both ways represents two edges with opposite directions between the same pair of vertices. Then inspection of the visual representations of graphs so far will show that all the vertices of a graph have the same number of edges coming into them and the same number leaving them. Is this true for all Cayley graphs? A little reflection shows that it is. Let v be any vertex of the Cayley graph of (G, S) . Then for each s in S , $\langle v, sv \rangle$ is an edge leaving v ; moreover, every edge leaving v is $\langle v, sv \rangle$ for some s in S . So the total number of edges leaving v is $|S|$. Again, for each s in S , $\langle s^{-1}v, v \rangle$ is an edge into v (as $v = ss^{-1}v$) and all edges into v are of this form. So the total number of edges into v is $|S|$. So all vertices of the graph have the same number of edges coming in and the same number leaving: in the terminology of graph theory, every Cayley graph is regular. In this case, visual inspection of some visual graphs (the standard visual representations of Cayley graphs) led to a general conjecture about Cayley graphs (the mathematical entities), a conjecture that is confirmed by reasoning.

Thought experiments with Cayley graphs: vertex transitivity

In the case just considered, attentive visual *inspection* of the visual graphs suggested a conjecture. Now we claim that operations in visual *imagination* can do

²² The pairs are ordered because all edges of a Cayley graph have a direction. Notice that edges can run in both directions between a given pair of vertices: both $\langle e, f \rangle$ and $\langle f, e \rangle$ are edges in the Cayley graph represented by Figure 12.

²³ For more on the roles of diagrams of Cayley graphs see the following articles by Irina Starikova: "Why Do Mathematicians Need Different Ways to Present Mathematical Objects? The Case of Cayley Graphs", *Topoi* **29**(1), (2010), pp. 41-51; "From Practice to New Concepts: Geometric Properties of Groups", *Philosophia Scientiae*, **16**(1), (2012), pp. 129-151.

the same kind of work. Looking at the visual graph for C_6 with anticlockwise rotation through $2\pi/6$ as generator, it is clear that we can move any vertex to any other by a transformation of the whole configuration that maps vertices one-to-one onto vertices, in such a way that edge relations are preserved. Putting g for the mapping, this means that $\langle v, w \rangle$ is an edge if and only if $\langle g(v), g(w) \rangle$ is an edge. The transformation is simply a rotation of the whole about the centre by as much as is required to take v to w , and this is made obvious to us by visual imagination.

A one-to-one mapping of the vertices of a graph onto those vertices (i.e. a permutation of the vertices) which preserves edge relations is said to be an *automorphism* of the graph. So the property of the Cayley graph of $(C_6, \{r\})$ which visual imagination revealed to us is this: for any of its vertices v and w there is an automorphism which maps v to w . A graph with this property is said to be *vertex transitive*.


For the finite cyclic groups $(C_n, \{r\})$ and the infinite cyclic group $(\mathbb{Z}, \{1\})$ finding an automorphism is very easy. Take any v and w in the group. If $v = w$, the identity function does the job. Otherwise, there will be some non-zero integer k such that $w = r^k v$. This is k rotations through $2\pi/n$ (or k unit translations) with direction depending on whether k is negative or not. But this same operation applied to all members of the group will preserve edge relations of the Cayley graph, as can be recognised from visualizing the operation on the graph as a whole.

What about finitely generated groups with more than one generator? Let us look at the Cayley graph for S_3 with generators r and f as depicted in Figure 13, the prism. A red directed edge represents one application of r , i.e. a step from a vertex v to a vertex rv ; a black edge is the merging of two edges with opposite directions, each representing one application of f . Let v and w be any distinct vertices. The thought experiments involve visualizing spatial operations on the whole prism until one finds one (or a sequence of them) which maps v to w and takes edges to edges without exception. There are three possibilities to consider. In each case we describe a visualizable operation (sequence) which clearly does the job.

(1) v and w belong to the same triangle, that is, $w = rv$ or $w = r^2v$. Visualize a rotation of the whole prism about the axis through the centre of both triangles (the 'horizontal' axis) by one or two thirds of a revolution in the direction of the red edge from v . For

example, let v and w be as in Figure 14. Anti-clockwise rotation of the whole prism by two thirds of a revolution maps v to w leaving edge relations undisturbed.²⁴


Figure 14



(2) v and w lie at opposite ends of the same black edge. Then reflection in a plane parallel to the triangles cutting the prism in half maps each vertex with its counterpart at the other end of a black edge and preserves edge relations. This is the mapping that takes each x to fx . An alternative is to rotate anticlockwise about the horizontal axis until the black edge between v and w is at the top, then rotate about the vertical axis through the centre by half a revolution as in Figure 15. If v and w start at the bottom left edge, this is the mapping that takes each x to fr^2x .

²⁴ We should be careful here when specifying the mapping mathematically, because the whole-prism rotation does not coincide with rotation r of the group: for x in the near triangle the mapping takes x to rx , but for x in the far triangle it takes x to $r^{-1}x$.

Figure 15



(3) v and w do *not* lie at opposite ends of the same black edge and do *not* belong to the red same triangle. It is easiest to consider two cases. (i) Let v and w be diagonally opposite vertices on the bottom face of the prism. Then a half revolution about the vertical axis as in Figure 15 takes v to w and leaves edge relations undisturbed. (ii) Let just one of v and w be a vertex at the top of a triangle. Then half a revolution about the vertical axis followed by one or two thirds of an anti-clockwise revolution, will take v to w and preserve edge relations.

We can visualize these whole-prism operations using a picture of the prism with vertices appropriately labelled v and w , and in so doing we readily discern that vertices are mapped one-to-one onto vertices and that edges are taken to edges. This is not surprising. Each whole-prism revolution we have mentioned is a bijection of vertices onto vertices which preserves edge relations; so the composite operation of performing one of these revolutions after another is a bijection which preserves edge relations: a composition of automorphisms is an automorphism.

We can conclude that for any v and w in the Cayley graph of S_3 there is an automorphism taking v to w : the Cayley graph is vertex transitive. What if we replace the triangular faces of the prism with matching regular polygons of more than three sides? The same three cases for pairs of distinct vertices need to be considered (with reference to triangles replaced by reference to n -gons), and using a power of schematic imagination it is not difficult to discern that the same kinds of

transformation will provide the needed automorphisms. The only difference is that in this case, the whole-prism rotations about the central horizontal axis which are available to us are k/n of a revolution for each integer k such that $1 \leq k \leq n$, instead of $1/3$ or $2/3$ of a revolution. As this works regardless of the number of polygon sides, we have a way of telling that the Cayley graphs of an infinite class of groups with two generators (the dihedral groups D_n) are vertex transitive.

Against the background knowledge that all groups with a single generator (the cyclic groups) have vertex transitive graphs, this finding raises the questions: Is the Cayley graph of every group with two generators vertex transitive? Is every Cayley graph vertex transitive? There are many kinds of groups with two generators that we have not considered; so it would be wrong to regard the thought experiments described here as providing significant evidence for the hypothesis that all 2-generated groups have vertex transitive graphs. *A fortiori* our thought experiments do not provide much evidence for the hypothesis that all Cayley graphs are vertex transitive. But the outcomes of our thought experiments make these hypotheses worthy of investigation, by trying to find a proof or a counterexample. In fact there is a fairly straightforward proof that every Cayley graph is vertex transitive.

What we have shown is that in some cases one can use one's visual imagination to find the required automorphisms, without already knowing that there are any, hence in a truly experimental way. This active use of visual imagery, first studied by cognitive scientists in the 1970s,²⁵ is a useful part of the toolkit of mathematicians and students of mathematics, though the results are usually recorded symbolically, without trace of the mental experimentation which led to them. The utility of visual imagination depends on confining our efforts to images and image transformations which are simple enough for us to manipulate reliably in imagination. But the variety of images and image transformations that we can handle reliably suffices to make visual imagination a potent instrument of mental experimentation in mathematics.

²⁵ Shepard, R. and Cooper, L. (ed.s), *Mental Images and Their Transformations*. Cambridge, Mass.: MIT Press 1982.

3. A CASE FROM GEOMETRIC GROUP THEORY

The example we will now present is the first step of a revolutionary advance in geometric group theory due to Russian mathematician Mikhail Gromov. To keep the exposition short and digestible, we omit some of the technical details.

If S and T are distinct finite subsets of a group G and both generate G , the Cayley graphs (G, S) and (G, T) will not in general be isomorphic. For example, the Cayley graphs of $(\mathbb{Z}, \{1\})$ and $(\mathbb{Z}, \{2, 3\})$, illustrated by Figures 10 and 11, are not isomorphic. How, if at all, can we use Cayley graphs of a group to discover properties of the group itself, that is, properties which are invariant with respect to generating sets?

The seminal thought is that we may be able to find group properties which do not depend on the generating set by ignoring the fine-grained local features of the different Cayley graphs of a given group and attending only to the coarse, global features shared by all the group's Cayley graphs. But how, given a particular Cayley graph of a group, can we tell what its coarse global features are?

A Cayley graph of an infinite (finitely generated) group is an infinite graph; so only finite portions of it can be visually represented. But we can imagine viewing ever larger portions of the graph in the hope that large scale features of the group may emerge. We can give this idea mathematical articulation by regarding a Cayley graph as a metric space, as follows.

For every pair g, h of members of a generated group (G, S) there is at least one path in its Cayley graph from g to h . The length of a path is just the length of the sequence of consecutively adjacent edges which constitutes the path, and the distance between g and h , denoted $d_S(g, h)$, is the length of a shortest path starting at g and ending at h . This distance function is the *shortest path metric*.²⁶ Viewing ever larger portions of the Cayley graph amounts to successively viewing diagrams representing the parts of the Cayley graph containing vertices at most n units away

²⁶ This *shortest path metric* d_S is the same as the *word metric* for (G, S) . Suppose each element of S is assigned a unique name of the form " s_k ". Let a symbol of the form " s_k^{-1} " name the inverse of what is named by " s_k ". A *word in S* is a finite sequence of elements of the form " s_k " or " s_k^{-1} ", so that a word denotes a product of members of G . The word metric for (G, S) is defined: $d_S(g, h) =$ least length of a word w in S such that w denotes hg^{-1} . Note that left multiplying g by hg^{-1} takes g to h .

from the identity e , for increasing n .²⁷ We call this kind of visual transformation ‘zooming out’.

Now let G be any infinite group with different finite sets of generators S , T and maybe others. From a visual representation of the Cayley graph of (G, S) or (G, T) – it does not matter which – we can try zooming out in visual imagination *so far* that the fine details of the Cayley graph are lost and features of the large scale geometry (or ‘coarse’ geometry) of the object now come into view. The hope is that the large scale geometry is the same for whichever Cayley graph we start with, that is, regardless of generating set. If this works, then we might find that algebraic properties of the group G itself, properties which are invariant with respect to generating set, can be discovered by attending to the coarse geometry of the object we reach by zooming out. Here is how Gromov put it:

This space [the space of the Cayley graph of $\Gamma = (G, S)$ with the shortest path metric] may appear boring and uneventful to a geometer’s eye since it is discrete and the traditional local (e.g. topological and infinitesimal) machinery does not run in Γ . To regain the geometric perspective one has to change his/her position and move the observation point far away from Γ . Then the metric in Γ seen from a distance d becomes the original distance divided by d and for $d \rightarrow \infty$ the points in Γ coalesce into a connected continuous solid unity which occupies the visual horizon without any gaps or holes and fills the geometer’s heart with joy.²⁸

To get Gromov’s point one must bear in mind that a Cayley graph is not a geometric object: its ‘edges’ are just pairs of vertices and so contain no points between endpoints. A Cayley graph with shortest path metric is a metric space, but the metric (the distance function) is discrete, as the distance between any two vertices is a non-negative integer. By ‘moving the observation point far away from’ the Cayley graph metric space (that is, by zooming out from it), the discrete object is transformed in appearance into a space with a dense and continuous metric, having (non-negative) real values.

²⁷ In customary terms, we bring into view (representations of) the n -balls for (G, S) for increasing n , where the n -ball = $\{g \in G: d_S(e, g) \leq n\}$.

²⁸ Gromov, M. ‘Asymptotic invariants of infinite groups’ In Niblo and Roller (ed.s), *Geometric group theory*. Volume 2 London Mathematical Society Lecture Notes **182**, (Cambridge University Press 1993).

What happens if we imagine zooming out from a visual presentation of a Cayley graph? If, actually looking at one, we zoomed out perceptually far enough the whole thing would disappear from view. To avoid this, an idealization is involved in our mental exercise: we suppose that while distances between vertices shrink as we zoom out, the vertices themselves do not fade at all. The experimental question is: what would happen to one's view of a standard diagram of a Cayley graph as one moved the observation point away by distances without upper bound, if vertices remained in view like points of starlight? What kind of space would emerge as a result? The answer depends on the Cayley graph one starts with, and is obtained by a combination of visual imagination and physical reasoning.

What happens then to standard diagrams of Cayley graphs of $(\mathbb{Z}, \{1\})$ and $(\mathbb{Z}, \{2, 3\})$, shown partially in Figures 10 and 11? Both become indistinguishable from the traditional representation of the real numbers as a single uninterrupted line without ends. In this case at least, differences due to different generating sets have been wiped out, as desired. What happens to the integer points of the plane $(\mathbb{Z} \times \mathbb{Z}, \{\langle 1, 0 \rangle, \langle 0, 1 \rangle\})$? The vertices in each horizontal string coalesce, and at the same time the vertices in each vertical string coalesce; that is, spaces between the points shrink and disappear, resulting in a continuous plane. What happens to the appearance of a *finite* Cayley graph as we imagine zooming out? Eventually its vertices coalesce to a single tiny dot. Does this nullify the whole exercise? Not at all: we are looking for *asymptotic* properties, properties that emerge at the limit of zooming out or properties that emerge at a late stage and persist, and so our focus is naturally on infinite groups (with finite generating sets.)

There is no reason to think that what we have described as an idealized mental operation of zooming out in visual imagination is really a disguised argument for something we already believe. For the question would then arise how we came to believe it, if not by the kind of thinking we describe. It is true that there is more to the mental operation than a simple transformation in visual imagination, for we add conditions. We ask how the appearance of the diagram of the Cayley graph of $(\mathbb{Z}, \{2, 3\})$, for instance, would be transformed by zooming out under the condition that vertices remained visible, though not necessarily distinguishable. This indicates that the cognitive processes involved are complicated and probably also not fully open to introspection. But it is clear that visual imagination of a spatial change is

involved and that the thinking as a whole does not reduce to the application of mathematically prescribed rules. Reliability is going to be limited by the fact that the kind of spaces we can easily visualize are Euclidean or embeddable in a Euclidean space. But the examples given fall within these limits.

The direct outcomes of these experiments do not count as mathematical results and, as just mentioned, the outcomes are limited. This is not a problem, because the real rewards of the zooming-out thought experiments are not their direct outcomes, but their effects in suggesting three mathematical possibilities. First, zooming-out suggests that there is a way of filtering out differences due to the different generating sets of the same group. Secondly, zooming out suggests that there is a way of thinking of an infinite generated group in terms of a metric space with a continuous metric (so that a group may have properties determined by geometric properties of the continuous metric space). Thirdly, zooming out suggests that we will sometimes get the same continuous metric space from distinct groups (not just the same group with different generating sets), perhaps giving us an equivalence relation on groups.

To benefit mathematically from these effects, we need to find a mathematically precise account of a suitable relation that holds between the Cayley graph of a finitely generated group (G, S) with the shortest path metric – call it $\Gamma(G, S)$ – and the continuous metric space we arrive at by zooming out from $\Gamma(G, S)$. The relation is suitable only if for any infinite group G generated by finite subsets S and T , $\Gamma(G, S)$ and $\Gamma(G, T)$ stand in this relation to the same continuous metric space.

The mathematization of the intuitive relation meeting these requirements is so neat that we present it now. An isometric mapping between metric spaces is one that preserves distances; a quasi-isometric mapping is one that preserves distances to within fixed linear bounds:

A map f from (S, d) to (S', d') is a *quasi-isometric mapping* iff there are real constants $K \geq 0$ and $L > 0$ such that for all x, y in S

$$d(x, y)/L - K \leq d'(f(x), f(y)) \leq L \cdot d(x, y) + K.$$

Quasi-isometric mappings are not in general surjective on the intended target space, and we will fail to capture the intuitive relation if we impose surjectivity as an extra

condition.²⁹ But we would like to find an equivalence relation on metric spaces which is a suitable loosening of isometry. So some extra condition is needed. The condition is that the mapping be surjective to within a fixed bound. Precisely put, the mapping f from (S, d) to (S', d') is *quasi-surjective* on S' iff there is a real constant $M \geq 0$ such that every point of S' is no further than M away from some point in the image of S under f . Putting these together, we define:

A map f from (S, d) to (S', d') is a *quasi-isometry* iff f is a quasi-surjective quasi-isometric mapping from (S, d) to (S', d') .

(S, d) is *quasi-isometric to* (S', d') iff there is a quasi-isometry from (S, d) to (S', d') .

This is an equivalence relation, which works as intended. First, a discrete space can be quasi-isometric to a dense continuous space. The inclusion (identity) mapping from $\Gamma(\mathbb{Z}, \{1\})$ to \mathbb{R} , with constants $L = 1$ and $K = 0$, is a quasi-isometric mapping; and it is quasi-surjective as every real number is at most $1/2$ a unit distance away from an integer. So $\Gamma(\mathbb{Z}, \{1\})$ is quasi-isometric to \mathbb{R} with standard distance metric.³⁰

Moreover, $\Gamma(\mathbb{Z}, \{2, 3\})$ also is quasi-isometric to \mathbb{R} with standard distance metric; so $\Gamma(\mathbb{Z}, \{1\})$ and $\Gamma(\mathbb{Z}, \{2, 3\})$ are quasi-isometric spaces. This fact generalizes: for any infinite group G with finite generating sets S and T , $\Gamma(G, S)$ and $\Gamma(G, T)$ are quasi-isometric spaces, as intended. This means that properties of G which are quasi-isometric invariants will be independent of the choice of generating set, and therefore informative about the group itself.

Furthermore, for some *different* groups G and H with generating sets S and S' respectively $\Gamma(G, S)$ and $\Gamma(H, S')$ are quasi-isometric; in this case the groups G and H are said to be quasi-isometric. So finitely generated infinite groups fall into equivalence classes modulo quasi-isometry.

Finally, an immensely rewarding outcome: there are some kinds K of *geometric* space such that groups with Cayley graph spaces which are quasi-isometric to a space of kind K (though not necessarily the *same* space) share significant algebraic

²⁹ This is because the intuitive relation holds between $\Gamma(\mathbb{Z}, \{1\})$ and \mathbb{R} with the normal distance metric, but $|\mathbb{R}| > |\mathbb{Z}|$; so there is no surjection from \mathbb{Z} to \mathbb{R} .

³⁰ For a quasi-isometry from \mathbb{R} to \mathbb{Z} , the mapping which takes each real number r to the nearest integer or, if r is half-way between integers, to the greatest integer less than r , is a quasi-isometric mapping (with $L = 1$ and $K = 1$) and is surjective, hence trivially quasi-surjective.

properties. This turns out to be the case for groups which are quasi-isometric to hyperbolic geodesic spaces, but that is a story for another occasion.³¹

4. SUMMARY AND DEFENCE

We have presented examples from knot theory, graph theory and geometric group theory of a kind of thinking which involves active use of visual imagination and goes beyond the application of mathematically prescribed rules, as a way of answering questions or overcoming obstacles. Is a trefoil knot equivalent to the unknot? What is the diameter of a cyclic graph in terms of the number of its edges? Is the Cayley graph of $(S_3, \{f, r\})$ vertex transitive? What spatial representations enable us to discover properties of a finitely generated group which are invariant with respect to generating sets?

Our impression is that the role (or roles) of this kind of experimental thinking in the advance of mathematical knowledge is under-appreciated, though we have not justified that opinion here. Our aim has been merely to substantiate the view that there are thought experiments in mathematics which involve visualization of physical situations or transformations, often with an idealised aspect.

These visual thought experiments neither are, nor serve in place of, mathematical *proofs* of the conclusions reached, even when those conclusions are true and the thought experiments are reliable ways of reaching them. But we hope that the cases we have presented support our view that the thought experiments can give the thinker good reason to believe them.

This raises a general philosophical worry. If visual thought experiments of the kinds we have described can provide reasons for mathematical beliefs, they would provide empirical reasons. But mathematics, as opposed to the application of mathematics to non-mathematical subject matter, is an *a priori* science. How then can there be empirical evidence for a mathematical fact?

The main problem here lies with the dictum that mathematics is an *a priori* science. It is ambiguous. If it means that any knowable mathematical truth can be known

³¹ See Starikova, I. "From Practice to New Concepts: Geometric Properties of Groups", *Philosophia Scientiae*, 16 (1), (2012), pp. 129-151.

without empirical justification, it is consistent with the claim that we can have empirical reasons for believing a mathematical proposition. But it may mean something much stronger, ruling out the possibility that we can have empirical reasons for believing a mathematical proposition. If the dictum has this stronger meaning, so much the worse for the dictum. Here is a simple example. How many vertices does a cube have? Your background knowledge includes the facts that cubes do not vary in shape and that material cubes will not differ from geometrically perfect cubes in number of vertices. To find out the answer one can inspect a material cube and count its vertices. (Or you can visualize a cube to find four at the top surface and four at the bottom.) The visual experience in this case provides evidence for your conclusion that a cube has 8 vertices. “But is this really a *mathematical fact*?” Why not? It is a very simple fact, but we can extend the problem: Do all platonic solids have the same Euler characteristic? Surely the answer to *that* is a mathematical fact. And it can be verified in the same way. Physical models of each of the five platonic solids can be visually inspected to find out whether $V - E + F$ is the same for all of them. The visual inspection provides empirical evidence for a positive answer. There are plenty of other examples, and a good case can be made that our initial knowledge of some single-digit addition facts is acquired empirically, from experiences of counting.

Empirical evidence has a much larger role in the epistemology of actual mathematical belief acquisition than is often thought. Even so, one may resist the idea that visual imagination is a way of providing us with empirical evidence. But visual imagination is not just a way of indulging in fantasy. It is also a way of harnessing the amalgamated memories of past experiences of visual perception to come to conclusions about physical situations. In this role it provides empirical evidence. Of course there is always the question, for any particular use of visual imagination to answer a question, whether it is reliable. There is no general test for reliability, but in the context of mathematics we have a way of resolving doubts: we look for a proof.

For these reasons we see no general bar to accepting that full blooded thought experiments are instruments, alongside proofs, for the advancement of mathematical knowledge.