GENERIC PROVING: REFLECTIONS ON SCOPE AND METHOD

URI LERON, ORIT ZASLAVSKY

A generic proof is, roughly, a proof carried out on a generic example. We introduce the term *generic proving* to denote any mathematical or educational activity surrounding a generic proof. The notions of generic example, generic proof, and proof by generic example have been discussed by a number of scholars (*e.g.,* Balacheff, 1988; Mason & Pimm, 1984; Rowland, 1998; Malek & Movshovitz-Hadar, 2011). All acknowledge the role of proof not only in terms of validating the conclusion of a theorem but, just as importantly, as a means to gain insights to why the theorem is true. In particular, we support and extend the argument made by Rowland (1998) that a generic proof does carry a substantial "proof power", and may in fact lie on the same continuum as the working mathematician's proof. In the same vein, we analyze possible ways that generic proof and proving may help in unpacking and making accessible to students at all levels *the main ideas* of a proof [1].

The article is organized as a reflection on three examples, or "mathematical case studies", which reveal increasingly more subtle facets of generic proving. The first mathematical case study is a simple and elementary theorem of numbers (also discussed in Rowland, 1998). The second example, a decomposition theorem on permutations, is still elementary in the sense of not requiring subject-matter knowledge beyond high school mathematics, but is more sophisticated in terms of the proof techniques required. The third example, Lagrange's theorem from elementary group theory, is more sophisticated both in terms of the proof techniques and the subject matter knowledge required. All the examples are introduced in a self-contained manner and all the terminology is explained and exemplified.

In the second part of the article, we reflect in more depth on the mathematical case studies of the first part, in an attempt to explicate some of the general features of generic proofs. For example, in an attempt to characterize the mathematical content of generic proofs, we look for commonalities with professional mathematicians' proofs as they appear in research journals and in university-level textbooks and lectures. For another example, we ask—and try to give some partial answers—about the *scope* of generic proving: what kind of proofs can be more or less helpfully approached via a generic version?

The article has been written in the form of a thought experiment. It is, however, solidly based in the experience of the authors in running many workshops with students and in international conferences on exactly these examples and ideas. Several researchers have previously discussed the more theoretical aspects of generic proofs. This research,

while relevant to the topic at hand, would take us away from our mathematical and pedagogical focus [2].

Mathematical case study 1: counting the factors of a perfect square

Theorem: *A natural number which is a perfect square (i.e., the square of another natural number) has an odd number of factors.*

For example, the number 16 (= 4^2) has 5 factors (namely: 1, 2, 4, 8, 16), and 25 (= 5^2) has 3 factors (namely: 1, 5, 25).

Generic Proof: Let us look at the perfect square 36 (= 6^2). We want to show that it has an odd number of factors. We list systematically all the factorizations of 36 as a product of two factors:

$$1 \times 36$$
$$2 \times 18$$
$$3 \times 12$$
$$4 \times 9$$
$$6 \times 6$$

All the factors of 36 appear in this list. (We could go on listing 9 × 4, 12 × 3, *etc.,* but because multiplication is commutative, this would just repeat the previous factorizations and would not produce new factors.) Counting the factors, we see that the factors appearing in all the products, except the last, come in pairs and are all different, thus totaling to an even number. Since the last product, 6 × 6, contributes only one factor to the count, we get, in total, an odd number of factors. Specifically, we have 2 × 4 + 1 = 9 factors.

From this first simple example we can already get an initial idea of what a generic proof is, and of some of its strengths and weaknesses. Obviously, our generic proof is not a complete proof, since the theorem has only been proved for the particular number 36. However, the number 36 was treated as *generic* in the sense that we did not make use of any of its specific properties except that it is a perfect square. In fact, all the important ideas of the general proof already appear in this generic proof, with the result that students could easily reproduce the proof for any other example. Indeed, they would most likely feel that they were carrying out *the same* proof. Thus, a generic proof serves as an easy introduction to the proof's *main ideas*.

Note the choice of 36 as our generic example. We felt that 36 could represent for the learner *any* perfect square, while 4, 16, 25 or even 169 (= 13^2) would have been too special to highlight the generalizability of the proof (*e.g.,* they would have too few factorizations). In Rowland's (1998) words, it is "small enough to be accessible with

mental arithmetic but with sufficient factors to be non-trivial" (p. 68). In Mason and Pimm's (1984) terms, it allows us "to see the general in the particular" (p. 277).

We partly concur with Movshovitz-Hadar's (1988) suggestion that a generic example should be "large enough to be considered a non-specific representative of the general case, yet small enough to serve as a concrete example" (p. 17). In general, however, "size" should be replaced by a measure of the *complexity* of the example. In the case considered here, complexity is measured by the number of factors, not the magnitude of the number; thus, 169 is less generic than 36, since the former is too special, having only 3 factors. In general, the example chosen should be "complex enough" to ensure that all the main ideas of the target proof will naturally surface in the context of the example.

Mathematical case study 2: decomposing a permutation into cycles

Theorem: *Every permutation has a unique decomposition as a product of disjoint cycles.* (These terms will be explained as the proof unfolds.)

In order to highlight both the mathematical and educational aspects of generic proving, we will present the theorem and its proof via a thought experiment of an idealized virtual classroom scenario [3]. In their previous lesson, the students in our scenario have already learned and practiced the definition of a permutation (a one-to-one mapping of the set $\{1, 2, \ldots, n\}$ onto itself) [4], and the 2-row notation for permutations [5]. They have also learned when two permutations are equal (when they are equal as functions), and the definition of multiplication for permutations (*i.e.,* perform the two mappings in succession, the same as composition of functions).

1. *Teacher:* Let's look at an example of a permutation, say the permutation σ below, and see if we can find anything interesting about its structure—how it can be constructed from simpler permutations (similarly to how numbers are constructed from their prime factors).

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 6 & 1 & 2 & 4 & 7 & 3 & 5 \end{pmatrix}$$

For example, let's start at 1, and follow its path as we apply the permutation σ over and over again, thus: $\sigma: 1 \to 6 \to \ldots$

[The students work in teams, continuing what the teacher has started: $\sigma: 1 \to 6 \to 3 \to 2 \to 1$]

2. *Alpha:* It came back to 1! There is no point going on, since it will just repeat the same numbers.

3. *Teacher:* Right. This part of the permutation is called *a cycle,* and is written (1 6 3 2). It is a special kind of permutation, in which each letter in the cycle notation goes to the next one on the right, except the last one, which goes back to the first. (The letters that don't appear in this notation are understood to be mapped to themselves; for example, in this cycle, $5 \to 5$.) Note that the same cycle

can also be written as (6 3 2 1), (3 2 1 6) or (2 1 6 3), since they are all equal as functions.

Let's see if we can find more cycles in our permutation. The letters 1, 2, 3 have already been used up, but 4 has not, so let's repeat the same game starting with 4.

[The students work in their teams to find the path of σ starting at 4.]

4. *Beta:* 4 goes to itself; we cannot construct a cycle.

5. *Teacher:* Since we see that $4 \to 4$, we write this as (4) and call this a *trivial cycle.* It is equal to the identity function, sending every letter to itself. What do we do next?

6. *Students:* Construct the cycle starting at 5 (the next unused letter).

[The students construct the path $5 \to 7 \to 5$ and the corresponding cycle (5 7).]

7. *Students:* Now all the numbers 1, 2, 3, 4, 5, 6, 7 have been used up—we can't construct any more cycles.

8. *Teacher:* Right. We can't and we needn't; we have now found all the cycles of our permutation. In fact, if we recall the definition of permutation product, we can see that our original permutation is actually equal to the product of the cycles we have found!

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 6 & 1 & 2 & 4 & 7 & 3 & 5 \end{pmatrix} = (1 \quad 6 \quad 3 \quad 2)(4)(5 \quad 7)$$

How do we know this? Take 1, for example. You can see that on both sides 1 goes to 6, and similarly for all other letters. (This is no coincidence: it's how we constructed the cycles.) Hence the permutations on the two sides are equal as functions. Notice that no number appears in two (or more) cycles on the right-hand side. The cycles are therefore said to be *disjoint.*

9. *Gamma:* Just a minute, if σ is a product of cycles, we must also take the other cycles into account when we calculate $\sigma(1)$ from this product.

10. *Delta:* Yes, but because the cycles are disjoint, 1 and 6 don't appear in any other cycle, which means that we can ignore them when calculating $\sigma(1)$.

11. *Teacher:* Right. We can summarize our work so far by saying that the permutation *has been decomposed as a product of disjoint cycles.*

You can also check that, unlike multiplication of permutations in general, the multiplication of disjoint cycles is commutative.

12. *Epsilon:* Can we always do this? Can we decompose *any* permutation as a product of disjoint cycles?

13. *Teacher* (to the class): Well, what do you think?

14. *Epsilon:* Why shouldn't we just repeat the same process for any permutation?

15. *Alpha:* Wait a minute! What if this procedure didn't work? We were lucky that 2 went back to 1 in the first cycle, but what if it didn't? What if it went back to 6 for example? Then we wouldn't have a *cycle*.

16. *Teacher:* If we had 2 going to 6, and earlier we also had 1 going to 6, then we would have both $\sigma(1) = 6$ and $\sigma(2) = 6$. Is this possible?

17. *Beta:* No, this is impossible, because a permutation is a one-to-one function so we can't have $\sigma(1) = \sigma(2)$.

18. *Teacher:* That's correct, therefore our procedure will always yield *cycles*. For the same reason, we can't have the same letter appearing in two different cycles, because this too would violate the one-to-one property of the permutation [6]. This guarantees that our procedure will generate *disjoint* cycles.

19. *Teacher* (summarizes): We have jointly constructed a generic proof of the theorem: *Every permutation can be decomposed as a product of disjoint cycles.*

20. *Teacher* (moving on): As an optional homework exercise you can try to formalize and generalize this generic proof to show that the conclusion holds for any permutation [7].

[The teacher and students proceed to establish the uniqueness of the decomposition, but we shall skip this part because of space limitations.]

Now we can state our full theorem: *Every permutation has a unique decomposition as a product of disjoint cycles.*

Note again the choice of generic example for the cycle decomposition theorem: a permutation on 7 letters, having cycles of lengths 1, 2 and 4. A shorter permutation on 6 letters would have been possible, with cycle lengths of 1, 2 and 3, but this orderly sequence looked to us a bit too special and possibly misleading. Thus, again, we have chosen the simplest example that would still be complex enough to represent the general case.

This mathematical case study also demonstrates a subtle pitfall that lurks behind generic proofs. In the example, some phenomena *just happen,* automatically, but would require a proof in the general case. Thus, the cycles *just happened* to close back on the initial letter, and they also *just happened* to be disjoint. The fact that this phenomenon *just happened* in the example might conceal the need for proof in the general case, thus bypassing some of the important ideas of the general proof. In fact, this is the only place where we are using the crucial property of permutations as a *one-to-one* function. In our idealized Lakatosian dialogue, the bright

students have brought up this issue themselves, but under more realistic conditions it is more likely that the teacher would have to raise this point.

Mathematical case study 3: Lagrange's theorem

Before we move on with a generic proof, we bring a brief mathematical introduction of Lagrange's theorem. We do this not by presenting a crash course in group theory, but by limiting our explanations to the context of the examples used. In effect we are preceding the generic proof of Lagrange's theorem by a *generic introduction* to the elements of group theory used in the proof.

The entire discussion of the generic proof occurs within the group Z_{12}, consisting of the set $\{0, 1, 2, \ldots, 11\}$ and the operation of *addition modulo* 12, denoted by $+_{12}$. For example, $2 +_{12} 3 = 5$, $5 +_{12} 7 = 0$, $5 +_{12} 8 = 1$, and, in general, $a +_{12} b$ is defined to be the remainder of the usual sum $a + b$ on division by 12. Z_{12} is a *group* in the sense that it contains 0 and is *closed* under addition modulo 12, *i.e.,* if a and b are in Z_{12}, then so is $a +_{12} b$ [8]. Furthermore, Z_{12} is a *finite* group since it contains a finite number of elements.

A *subgroup* of Z_{12} is a subset of $\{0, 1, 2, \ldots, 11\}$ which is in itself a group under the operation defined in Z_{12}. For example, it can be checked that the subset $H = \{0, 3, 6, 9\}$ is a subgroup of Z_{12}, since it contains 0 and is closed under $+_{12}$. (For example, $6 +_{12} 9 = 3$, which is again a member of H.)

The *order* of a finite group G is the number of its elements, and is denoted $o(G)$. Thus, $o(Z_{12}) = 12$ and $o(H) = 4$. We note that 4 divides 12. As it turns out, this is not a coincidence, and is in fact an example of the following theorem, which is probably the most important theorem in elementary group theory.

Lagrange's theorem: *If H is a subgroup of a finite group G, then the order H divides the order of G.*

A generic proof of Lagrange's theorem: We will carry out the proof of the theorem for the group $G = Z_{12}$ and the subgroup $H = \{0, 3, 6, 9\}$. It may seem odd to prove that $o(H)$ divides $o(G)$ for this case, for obviously no proof is necessary for the fact that 4 divides 12. However, while we already know that $o(H)$ divides $o(G)$ in our example, we do not know *why* this is so, and to this end we do need the generic proof. The generic proof will demonstrate the general *process* which serves to carry out the proof in general.

The main idea of the proof is that by creating "shifts" of $H,$ we obtain a partition of G into disjoint subsets having the same cardinality as $H,$ from which we can calculate $o(G)$ from $o(H)$. Specifically, given an element g in G, we define *the coset of g and H in G* as the following subset of G:

$$H +_{12} g = \{h +_{12} g : h \text{ in } H\}$$

We calculate the cosets of H in G in our example.

$$H +_{12} 0 = \{0 +_{12} 0, 3 +_{12} 0, 6 +_{12} 0, 9 +_{12} 0\} = \{0, 3, 6, 9\},$$
(*i.e., H* itself).

Similarly:

$$H +_{12} 1 = \{1, 4, 7, 10\}$$
$$H +_{12} 2 = \{2, 5, 8, 11\}$$
$$H +_{12} 3 = \ldots$$

We leave it for the reader (or for the class in the Lakatosian scenario) to show that from now on we are not getting new cosets but are only repeating the old ones.

The number of distinct cosets of any subgroup H of a finite group G is called the *index of H in G,* and is denoted $i_G(H)$. In our example, where $G = Z_{12}$, we have $i_G(H) = 3$.

We can see that all the cosets have the same number of elements as H (4 in our example), and that each element in G appears in one and only one of the cosets of H in G. Since G is now presented as the disjoint union of the different cosets of H in G, we can conclude that the order of G is the number of distinct cosets of H in G times the order of H, namely:

$$o(Z_{12}) = o(H) \times i_G(H)$$

which is an even a stronger statement than what we had to prove.

As in the permutation decomposition theorem, the virtual class would have noticed that some of the relations that in the example "just happened", would require a proof that they should *always happen* in the general case. This includes the facts that all cosets have the same number of elements (the order of H), and that distinct cosets are disjoint. In fact, as in the permutations example, proving these "lemmas" is where we actually use the group and subgroup definitions.

The main contribution of this mathematical case study to the general discussion of generic proving is the slippery nature of the question, "What is a good generic example in the context of a generic proof?" Indeed, the group Z_{12}, being a *cyclic* group (generated by a single element), is the simplest kind of group imaginable, and thus definitely *not* a generic example of a finite group. If someone asked you for an example of a "typical" finite group, you would definitely not think of giving them this example. Still it does a fair job of exemplifying the main ideas of the proof of Lagrange's theorem. Moreover, choosing a more "generic" example of a group (say, the so-called *symmetric group* S_3 with 6 elements or S_4 with 24 elements), would have made the generic proof computationally much more complicated, and what we might have gained in generality would have been lost in simplicity and learnability. Thus, a delicate balance between generality and simplicity is needed in making this didactical choice. When the mathematical objects making up the proof are simple (numbers, permutations, Eulerian circuits in graphs), then the balance leans towards the generality of the generic example. When the objects become more complicated and abstract (groups, subgroups, cosets), the balance leans towards simplicity rather than generality.

Finally, this mathematical case study also highlights the fact that the test of genericity should be applied not to the example itself (Z_{12} is not a good generic example of a group), but rather to the proving process that this example generates: the process of partitioning G by its cosets, and the properties of this partition, are quite general, though the group to which we are applying this process is not.

Reflections on scope and method

Reflecting and generalizing from the above examples brings up some important mathematical and educational issues regarding generic proofs. We list these issues as questions with tentative partial answers.

What are the strengths of generic proofs?

In learning proofs, one can distinguish two different types of activities: understanding a proof (presented by a book or teacher) and creating a proof (given the theorem). Generic examples can substantially help teachers and students in the pursuit of both goals.

First, generic proofs can help understanding by enabling students to engage with the main ideas of the complete proof in an intuitive and familiar context, temporarily suspending the formidable issues of full generality, formalism and symbolism. While a complete formal proof may be beyond the reach of almost all school children (*e.g.,* Healy & Hoyles, 2000; Stylianides, 2007), we could imagine a classroom activity whereby even elementary school children learn the generic proof of the perfect square theorem (our mathematical case study 1), and produce their own versions for other examples. Indeed, they would most likely feel that they were carrying out *the same* proof.

In more complicated proofs, such as our permutations example, it is possible to build up the complexity gradually, via a chain of successively more elaborate *partial* generic proofs, each highlighting finer points of the proof that were not salient in previous steps. Thus, while the complete formal proof may be beyond reach for most high school students, they can still get a good view of the main ideas via a generic proof. Even for college-level students, preceding the complete proof by a generic version may help in highlighting the main ideas of the proof, separating them from the technicalities of formalism and notation.

Through the process of generic proving, teachers can help students move beyond empirical proof schemes (Harel & Sowder, 1998, 2007) and raise their need for proof (Zaslavsky *et al.,* 2012). Finally, generic proving can help students create the complete proof by serving as a graded sequence of hints in a guided discovery process. This aspect will be discussed more fully below.

What are the weaknesses of generic proofs?

The main weakness of a generic proof is, obviously, that it does not really prove the theorem. The "fussiness" of the full, formal, deductive proof is necessary to ensure that the theorem's conclusion infallibly follows from its premises. In fact, some of the more subtle points of a proof are prone to be glossed over in the context of the generic proof: some steps which "just happen" in the example, may require a special argument in the complete proof to explain *why* they happen, and to ensure that they will *always* happen. In the generic proof of the cycle decomposition theorem, for example, we have seen that cycles just "turn out" to close back to their first element, and that cycles just "turn out" to be disjoint. In fact, if we had not been careful, we could have completed the generic proof without ever utilizing the crucial one-to-one property of permutations. Since these essential issues do not naturally come up in the course of generic proving, the teacher's initiative here is crucial. Similarly, in the generic proof of Lagrange's theorem, students who see that all cosets have the same number of elements, may not notice that this fact requires proof and, in fact, depends crucially on the group definition.

We can capture this important difference by saying that in the generic proof some facts are simply *observed*, whereas in the complete proof they have to be *derived*. Thus, the fact that all cosets have the same number of elements (or even, for that matter, Lagrange's theorem itself) is simply observed in the case of our example, but must be derived in the general proof. This issue presents a challenge to anyone who teaches with generic proofs: the teacher needs to motivate the students to learn these additional parts of the proof, which may appear unnecessary in the context of the example [9].

When we have presented a class with a generic proof, what have students learned? What have we proved?

First, as we have pointed out before, students have learned how to carry out the proof on *any* example, not just the one we demonstrated. They have also learned (at the level of the example) the *main ideas* of the proof. But what have we actually proved? What is the mathematical status of a generic proof? We know that a proof carried out on an example does not count as proof, but does this mean that we have actually proved *nothing?* Well, it must be admitted that at the formal level we have indeed proved nothing: no mathematical journal would accept for publication a generic proof of a new theorem (though even mathematical journals and textbooks occasionally indulge in isolated generic proofs for some propositions; see the example two paragraphs below). Still, the feeling persists that we did go a long way towards presenting the complete proof. How can we capture more explicitly the source of this feeling?

We could start with an observation: given a good generic proof, any professional mathematician (say, a specialist in the relevant topic area) could easily generalize and formalize it into a full formal proof. If asked about the difficulty of the task, she would likely describe this as a "technical exercise" (or even "trivial exercise"): it could require quite a bit of technical work, but hardly any additional insight, discovery or creativity. (The fact that this "technical exercise" might be beyond the mathematical powers of most undergraduate mathematics majors is not relevant to the present theoretical discussion.)

We bring one example to support this observation. In his undergraduate textbook, *Abstract Algebra,* Herstein (1986) introduces the cycle decomposition theorem via a generic example (paralleling lines 1 – 11 in our virtual classroom scenario), concluding with the following remark:

> There is nothing special about the permutation [in our example] that made the argument we gave go through. The same argument would hold for *any* permutation […]. We leave the formal writing down of the proof to the reader. (Herstein, 1986, pp. 132-133)

Remarkably, writing down the complete proof appears as Exercise 4 on p. 136, under the subheading "easier problems" (the other categories are "middle-level" and "harder" problems). This reference does not establish that writing down the complete proof (given a generic proof) is an easy exercise for most undergraduate students—in our experience

it is not; it does support our claim that a professional mathematician would view this as just a technical exercise.

One way to clarify the status of a generic proof is to reconceptualize it as *a recipe for the learner on how to construct the complete proof* (a kind of closely-guided discovery learning). We could once again imagine a classroom thought experiment, presenting the following step-wise teaching activity. The teacher (T) asks a student (S) to prove the cycle-decomposition theorem. S tries for a while and asks for help. T: ok, I'll give you a hint. T shows S how the proof is carried out on an example, something like lines 1 – 11 in our classroom scenario, and once again asks S to try to prove the theorem. If she still can't complete the proof, T gives her a few more hints (like the subsequent steps in our classroom scenario), until she says: now I have all the ingredients for constructing the full proof. (Let us assume for simplicity that she does have the technical expertise to deal with the formalization itself.)

On the face of it, this generic proving activity seems good only for didactical purposes and is not related to real proofs as conceived by working mathematicians, but actually there may be a stronger connection than first meets the eye. The reason is that even the working mathematician's proofs, as they appear in research journals, are far from being full formal proofs, and you might well ask (just as for generic proofs), what have they actually proved. The answer is not simple, but we could approximate it by saying that the working mathematician's proof is still a recipe for how to write the "ideal" complete proof. The unofficial implicit rules of mathematical discourse require only that you explicate the details of the proof to the extent that it can convince another expert in your field that given enough time and motivation, your sketch could be fleshed out into a full "ideal" proof.

Several mathematicians have expressed closely related views on the nature of a "working mathematician's proof". Here are two examples:

> Proving a claim is, for a mathematician, an act of producing, for an audience of peer experts, an argument to convince them that a proof of the claim exists. […] The convinced listener feels empowered by the argument, given sufficient time, resources, and incentive, to actually construct a formal proof. (Hyman Bass, personal communication [10])

> To be sure, in practice no one actually bothers to write out such formal proofs. In practice, a proof is a sketch, in sufficient detail to make possible a routine translation of this sketch into formal proof. (Mac Lane, 1986, p. 377)

In view of the above quotations, since both the professional mathematician and the students do not write a complete formal proof, but only a recipe that convinces someone else that writing such a proof would be "merely a technical exercise" (or, in Mac Lane's words above, "a routine translation"), it is possible to view the difference between a generic proof and a mathematician's proof as a matter of degree rather than kind.

Not all proofs are equally amenable to a genuine generic version. Can we characterize the proofs (or parts thereof) that are so amenable?

This fascinating and difficult question raises a host of mathematical and educational issues. An answer would likely involve the form and structure of the proof (a mathematical aspect), but the effectiveness of a generic version is expressed in terms of its ability to render the main ideas of the general proof accessible to a particular audience (an educational aspect).

One observation that stands out of the mathematical case studies is that if a proof involves an act of construction (of a mathematical object or process), then this construction can be effectively presented via a generic example. Thus, we have shown how to construct all factorizations of the given number in our first mathematical case study, a cycle decomposition in our second mathematical case study, and a partition of the group into its cosets in the third. Significantly, whether a proof does or does not involve an act of construction may depend on how we choose to formulate the proof. Often an act of construction in a proof is hidden by the linear mathematical formalism, but may be revealed by "structuring" the proof (Leron, 1983, 1985a), whence a generic version of the proof may become accessible.

Some proofs may not seem on the surface to be amenable to a generic version because of their structure or logical form, or the nature of the mathematical objects involved; for example, proof by contradiction or proofs involving infinite objects. But even in such cases, we can often isolate some constructive element that can be presented via a generic example. We mention three such examples.

1. *Euclid's proof of the infinitude of prime numbers.* The basic construction here (given any finite set of primes, construct a new prime not in the set) can be presented via a generic example. In fact, Euclid himself does this in his *Elements*, Book IX, Proposition 20 (Reid & Knipping, 2010, p. 135; see also Leron, 1985b). Reid and Knipping point out that Euclid had no choice but to base his proof on a generic example, since he lacked the notation to discuss *any* number.

2. *Cantor's proof that the real numbers are uncountable, where given any list of real numbers, a new real number is constructed by the diagonal method.* The diagonal method itself (given a rectangular table of numbers, construct a row different from all the rows in the table) can be first introduced via small finite generic examples and then gradually extended to the infinite case (*e.g.,* Leron & Moran, 1983).

3. *Lagrange's theorem on finite groups: the order of a subgroup divides the order of the group.* Since this theorem concerns a relation on the collection of all finite groups, it might be hard to see at first glance how it could be helped by a generic example. But since the *proof* involves a construction (a partition of the group into cosets) it is possible to devise classroom activities that demonstrate the main ideas of the proof on a generic example as

we show in our third mathematical case study.

In contrast to these examples, we may consider the Heine-Borel theorem from analysis:

4. *A subset of R^n is compact if and only if it is closed and bounded.* The mathematical concept of compact (or closed or bounded) space can of course be exemplified, but since the theorem and its proof deal with complex logical relations between infinite collections of infinite objects, it is hard to imagine how it could be effectively demonstrated through an example.

Conclusion

At the core of this article is the method of generic proving: accessing a complicated proof by a chain of intermediate steps, where each step highlights some of the ideas of the proof by performing them on a generic example, and where details, refinements and complications are gradually added as we progress along the chain. Ideally, this should enable the learner to reach even a complicated and "unnatural" proof via a sequence of relatively easy and natural steps.

We conclude with a few possible directions for further theoretical and empirical study.

- Elaborate on methods of using generic proofs in actual classrooms, both for understanding and for generating proofs, and test their efficacy empirically.

- Study the work of college students to investigate the hypothesis that preceding a complex formal proof by a generic version would enhance their understanding of the main ideas behind the proof.

- Interview professional mathematicians to investigate the role of generic proofs in their research and teaching.

- Interview professional mathematicians to investigate the hypothesis that proofs are stored in their long-term memory in the form of generic examples. On theoretical grounds, two supporting arguments can be given for this hypothesis. First, this is simply a more economical way to store such proofs in memory: one only memorizes the main ideas, leaving out the technical details which can be readily reconstructed (by an expert) when needed. Second, this is an efficient way of storing past experiences so that one may recognize similar problem-solving situations in the future and re-use the same methods [11].

Notes

[1] A comprehensive discussion of research on proofs in mathematics education, including generic proofs, is beyond the scope of this article. The reader is referred to the excellent recent book by Reid and Knipping (2010) for such a survey.

[2] Readers are referred, for example, to Balacheff's (1988) distinctions between pragmatic vs. conceptual proofs, and between generic example vs. thought experiment and to Herbst's (2004) discussion of students' interaction with drawings and diagrams in geometrical proofs, and the teacher's role in supporting this interaction.

[3] In adopting this format we have obviously been influenced by the imag-

inary classroom scenario in Lakatos's (1976) seminal *Proofs and Refutations,* as well as the dialog with "the ideal mathematician" in Davis and Hersh (1981). As in Lakatos (1976), the scenario is taking place in an *idealized class* of bright and highly motivated students. A realistic class is a noisy situation (pun intended) and we prefer to describe our ideas first in a simplified and idealized setting. A realistic scenario would be much lengthier and meandering, with lots of false starts and dead ends and, most significantly, with the teacher having to shoulder a larger portion of the classroom discussion, rather than eliciting most of the insights from the students.

[4] The numbers on which the permutation operates could be any *n* symbols and are therefore referred to as letters.

[5] In this graphical representation of the permutation, the numbers in the bottom row are the images of the corresponding numbers in the top row. Thus, for the permutation σ defined next, $\sigma(1) = 6$, $\sigma(2) = 1$, $\sigma(3) = 2$, *etc.*

[6] A little trick is needed here, in fact a masked application of mathematical induction: assume on the contrary that there is a letter that appears in two different cycles, and take *the first* such occurrence. Then the letters appearing immediately before these two occurrences must be different, which cannot happen in a one-to-one function.

[7] For one such proof see Gallian (1990, p. 88).

[8] The general definition of a group includes more requirements, namely associativity and the existence of inverses. However, we do not need to worry about these here because, in general, associativity for addition mod *n* can be shown to be "inherited" from the associativity of the usual addition of integers, and the existence of inverses can be shown, in the finite case, to follow from the other group properties.

[9] This is similar to the teacher's problem when proving a theorem in geometry, which appears obvious from the accompanying figure or from dynamic geometry activities.

[10] From a draft manuscript by Bass for a book in preparation under the editorship of Peter Casazza, Steven G. Krarb and Randi D. Ruden currently entitled "I, Mathematician."

[11] This last argument sits well with Minsky's K-theory (1985, Chapter 8), according to which the most efficient way for people to store their memories for future problem solving is at middle-level abstraction: not too concrete but also not too abstract. It also sits well with Rosch *et al.'s* (1976) theory of basic level categories: "[Categories] within taxonomies of concrete objects are structured such that there is generally one level of abstraction at which the most basic category cuts can be made. In general, the basic level of abstraction in a taxonomy is the level at which categories carry the most information, possess the highest cue validity, and are, thus, the most differentiated from one another" (p. 383).

References

Balacheff, N. (1988) Aspects of proof in pupils' practice of school mathematics. In Pimm, D. (Ed.) *Mathematics, Teachers and Children,* pp. 216-235. London, UK: Hodder and Stoughton.

Davis, P. J. & Hersh, R. (1981) *The Mathematical Experience.* Boston, MA: Houghton Mifflin.

Gallian, J. A. (1990) *Contemporary Abstract Algebra* (2nd edition). Lexington, MA: D. C. Heath & Co.

Harel, G. & Sowder, L. (1998) Students' proof schemes: results from exploratory studies. In Schoenfeld, A. H., Kaput J. & Dubinsky, E. (Eds.) *Research in Collegiate Mathematics Education III,* pp. 234-283. Providence, RI: American Mathematical Society/Mathematical Association of America.

Harel, G. & Sowder, L. (2007) Toward comprehensive perspectives on the learning and teaching of proof. In Lester, F. (Ed.) *Second Handbook of Research on Mathematics Teaching and Learning,* pp. 805–842. Charlotte, NC: Information Age Publishing.

Healy, L. & Hoyles, C. (2000) A study of proof conceptions in algebra. *Journal for Research in Mathematics Education* **31**(4), 396-428.

Herbst, P. (2004) Interactions with diagrams and the making of reasoned conjectures in geometry. *ZDM* **36**(5), 129-139.

Herstein, I. N. (1986) *Abstract Algebra.* New York, NY: Macmillan.

Lakatos, I. (1976) *Proofs and Refutations: The Logic of Mathematical Discovery.* Cambridge, UK: Cambridge University Press.

Leron, U. (1983) Structuring mathematical proofs. *American Mathematical Monthly* **90**(3), 174-185.

Leron, U. (1985a) Heuristic presentations: the role of structuring. *For the Learning of Mathematics* **5**(3), 7-13.

Leron, U. (1985b) A direct approach to indirect proofs. *Educational Studies in Mathematics* **16**(3), 321-325.

Leron, U. & Moran, G. (1983) The diagonal method. *The Mathematics Teacher* **76**(9), 674-676.

Mac Lane, S. (1986) *Mathematics Form and Function.* New York, NY: Springer-Verlag.

Malek, A. & Movshovitz-Hadar, N. (2011) The effect of using transparent pseudo-proofs in linear algebra. *Research in Mathematics Education* **13**(1), 33-58.

Mason, J. & Pimm, D. (1984) Generic examples: seeing the general in the particular. *Educational Studies in Mathematics* **15**(3), 277-289.

Minsky, M. (1985) *The Society of Mind.* New York, NY: Simon & Schuster.

Movshovitz-Hadar, N. (1988) Stimulating presentations of theorems followed by responsive proofs. *For the Learning of Mathematics* **8**(2), 12-19, 30.

Reid, D. A. & Knipping, C. (2010) *Proofs in Mathematics Education: Research, Learning and Teaching.* Rotterdam, The Netherlands: Sense Publishers.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M. & Boyes-Braem, P. (1976) Basic objects in natural categories. *Cognitive Psychology* **8**(3), 382-439.

Rowland, T. (1998) Conviction, explanation and generic examples. In Olivier, A. & Newstead, K. (Eds.) *Proceedings of the 22nd Conference of the International Group for the Psychology of Mathematics Education,* vol. 4, pp. 65-72. Stellenbosch, South Africa: University of Stellenbosch.

Stylianides, A. J. (2007) Proof and proving in school mathematics. *Journal for Research in Mathematics Education* **38**(3), 289-321.

Zaslavsky, O., Nickerson, S. D., Stylianides, A. J., Kidron, I. & Winicki-Landman, G. (2012) The need for proof and proving: mathematical and pedagogical perspectives. In Hanna, G. & de Villiers, M. (Eds.) *Proof and Proving in Mathematics Education,* pp. 215-229. Dordrecht, Germany: Springer.

And if this were not enough, there is the question of the discourse into which the story is woven and the two aspects of story […]: the *fabula* and the *sjuzet*, the timeless and the sequenced. Which is constrained, and in what ways? […]

I think we would do well with as loose fitting a constraint as we can manage concerning what a story must "be" to be a story. And the one that strikes me as most serviceable is the one with which we began: narrative deals with the vicissitudes of intention.

Bruner, J. (1986) *Actual Minds, Possible Worlds,* pp. 16-17. Cambridge, MA: Harvard University Press.